

THÈSE DE DOCTORAT DE

L'UNIVERSITÉ DE RENNES 1

ÉCOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*
Spécialité : *Mathématiques et leurs interactions*

Par

Lucien GRILLET

La Conjecture de Smith en faible régularité

Thèse présentée et soutenue à Rennes, le 1er décembre 2022
Unité de recherche : IRMAR

Rapporteurs avant soutenance :

Louis FUNAR Directeur de recherche au CNRS, Université de Grenoble
Frédéric LE ROUX Professeur des universités, Sorbonne Université

Composition du Jury :

Présidente :	Françoise DAL'BO	Enseignante-chercheuse, Université de Rennes 1
Examineurs :	Laurent BESSIERES	Professeur des Universités, Université de Bordeaux
	Louis FUNAR	Directeur de recherche au CNRS, Université de Grenoble
	Frédéric LE ROUX	Professeur des universités, Sorbonne Université
Dir. de thèse :	Juan SOUTO	Directeur de recherche CNRS, Université de Rennes 1

REMERCIEMENTS

Dans le travail en grande partie solitaire que représente une thèse de doctorat, il est bon de ne pas oublier que cette dernière s'inscrit dans un contexte plus global l'ayant rendu possible. Les idées ne pouvant germer que dans des conditions propices et étant le fruit de l'échange, de nombreuses personnes peuvent donc se sentir concernées par l'aboutissement de la présente thèse :

- Juan, pour ses apports évidents, son aide, ses conseils et son soutien tout au long de ma thèse. Il a su trouver un sujet me correspondant parfaitement et a toujours été à l'écoute pour m'accompagner au mieux dans tous mes choix.
- Louis et Frédéric, pour avoir accepté le rôle de rapporteur en prenant le temps de se plonger profondément dans mes travaux. Merci pour leurs nombreuses remarques pertinentes qui ont permis de parfaire ce manuscrit.
- Laurent et Françoise pour leur travail et leur temps dans leur rôle de Jury.
- Toutes les équipes administratives, techniques et d'entretien, absolument nécessaires au bon fonctionnement du laboratoire.
- Toutes les personnes ayant participé aux projets tels que LibGen ou Sci-Hub, outils indispensables à tous les chercheurs, et remplissant un rôle d'accès à la connaissance dont les états se sont rendus incapables.
- Tous les mathématiciens et mathématiciennes avec qui j'ai pu échanger à Rennes et ailleurs, et particulièrement Chloé et les membres de l'équipe de Théorie Ergodique pour leur accueil chaleureux.
- Ma famille et tous mes amis qui m'ont toujours soutenu, que je n'oserais pas essayer de lister, par peur de trop en oublier.

Merci à tout ce beau monde.

TABLE OF CONTENTS

Présentation générale de la thèse	9
I The Smith conjecture in low regularity	21
1 Topological tameness	27
1.1 Introduction	27
1.1.1 The Jordan-Brouwer separation theorem	27
1.1.2 Generalizations to higher dimensions	28
1.1.3 The Alexander horned sphere	29
1.2 Wildness and tameness	32
1.2.1 Wild and Tame embeddings	32
1.2.2 Wild and tame manifolds	33
1.2.3 The Fox-Artin sphere	36
1.2.4 Some objects seem wild but are not	39
1.2.5 Local flatness	41
1.3 Shrinkings of manifolds	42
1.3.1 Cell-like maps and Moore's Theorem	44
1.4 A tameness criterion	45
1.4.1 First case : Σ is a surface	46
1.4.2 Second case : Σ is a circle	54
2 The Smith conjecture	57
2.1 Introduction	57
2.1.1 Smith theory	58
2.1.2 The Bing Involution	59
2.2 The Lipschitz Smith Conjecture	60
2.2.1 Early results	60
2.2.2 Some results on Lipschitz vector fields	62
2.2.3 The Smith Conjecture for C^1 maps	63

TABLE OF CONTENTS

2.2.4	Reducing Theorem 1 to two propositions	64
2.3	Proof of Proposition 2	67
2.3.1	A vector field near the fixed set	67
2.3.2	Proof of Lemma 15 and Lemma 16 in a flat setting	70
2.3.3	Reduction to the flat case	75
II	Manifold learning	77
3	Effective estimation of the dimension of a manifold from random sam- ples	81
3.1	Introduction	81
3.2	The estimator	87
3.3	Some Geometry	89
3.3.1	Reach-1 manifolds	90
3.3.2	Injectivity radius	90
3.3.3	Volumes of balls	91
3.3.4	Volume of thick diagonal	94
3.3.5	The gap	95
3.4	Sampling the thick diagonal	96
3.4.1	Some probability	96
3.4.2	The function we care about	98
3.4.3	Some technical results	100
3.4.4	Searching decent scales	102
3.5	Heuristics	105
3.6	Comparison with other estimators	112
III	3D printing	121
	Bibliography	129

Présentation générale de la thèse

LA CONJECTURE DE SMITH EN FAIBLE RÉGULARITÉ

La Conjecture de Smith

La conjecture de Smith s'inscrit dans l'étude et la classification des homéomorphismes d'ordre fini des variétés de dimension 3.

Soit σ un homéomorphisme d'ordre fini de S^3 dans lui-même avec des points fixes et préservant l'orientation. En 1939, Paul Althaus Smith démontra [Smi39, 7.3 Theorem 4] que l'ensemble des points fixes d'une telle application était nécessairement homéomorphe à un cercle S^1 . Ses résultats ne renseignent cependant que sur la topologie intrinsèque de l'ensemble des points fixes, mais ne donnent pas d'information sur la manière dont ce cercle est plongé dans S^3 . P. A. Smith demanda donc [Eil49, Problem 36] si ce cercle de points fixes peut être noué, ou s'il est nécessairement isotope au plongement standard du cercle. C'est ce que l'on appelle la conjecture de Smith.

Conjecture (Conjecture de Smith). *Les points fixes d'un homéomorphisme non trivial d'ordre fini de S^3 dans lui-même avec des points fixes et préservant l'orientation ne peuvent pas former un cercle noué.*

Plus généralement, cette conjecture adresse la question de savoir si tous les homéomorphismes d'ordre fini de S^3 sont conjugués à un élément du groupe des transformations orthogonales O_4 .

En l'état, la réponse à cette conjecture est négative. En 1952, R. H. Bing donna un exemple [Bin52] d'une involution de S^3 renversant l'orientation et dont l'ensemble des points fixes est homéomorphe à une sphère S^2 plongée de manière "sauvage" dans S^3 . Cet exemple fut modifié peu après par Montgomery et Zippin [MZ54], démontrant l'existence d'un contre-exemple à la Conjecture de Smith dans lequel les points fixes forment un cercle "sauvagement noué".

Dans ce contexte, le terme "sauvage" désigne un objet qui ne peut pas être envoyé sur un polyèdre par le moyen d'un homéomorphisme ambiant de S^3 . Ce caractère sauvage

peut apparaître car la continuité est la seule régularité demandée sur l'application σ .

Sous des hypothèses plus fortes, la conjecture de Smith peut cependant avoir une réponse positive. Ainsi, s'il on suppose que l'application σ est lisse, la conjecture est démontrée.

Théorème (Conjecture de Smith lisse). *Les points fixes d'un difféomorphisme non trivial d'ordre fini de S^3 dans lui-même avec des points fixes et préservant l'orientation ne peuvent pas former un cercle noué.*

Ce théorème fut démontré en 1984 par John Morgan et Hyman Bass grâce à d'importantes avancées dans l'étude des variétés de dimension 3 par de nombreux mathématiciens. Aujourd'hui, ce résultat découle plus généralement des travaux de géométrisation de William Thurston et Grigori Perelman.

Nous pouvons alors nous demander ce qu'il advient de la Conjecture de Smith dans un cadre où l'application σ a davantage de régularité qu'un homéomorphisme, sans être lisse.

Dans une série de conférences données en 2013 à Santa Barbara [Fre13, Conjecture 3.21], Michael Freedman énonce la conjecture suivante.

Conjecture (Freedman). *Toute action bilipschitzienne d'un groupe fini sur une variété compacte de dimension 3 est conjugué à une action lisse.*

La présente thèse se propose de répondre partiellement à cette question. Nous nous intéresseront au cas d'une application σ d'ordre fini et $(1 + \varepsilon)$ -bilipschitzienne. Plus précisément, notre résultat principal est le théorème suivant.

Théorème. *Pour $\varepsilon = \frac{1}{4000}$, toute action finie d'un groupe cyclique par homéomorphismes $(1 + \varepsilon)$ -bilipschitziens sur une variété compacte de dimension 3 est conjuguée à une action lisse.*

Nous démontrons ce théorème grâce à la notion de plongement sage.

Plongements sages

Soit N une partie d'une variété M . Si N est une variété topologique pour la topologie induite par M , on dit que N est une sous-variété topologique de M . S'il existe un homéomorphisme ambiant de M envoyant N sur un polyèdre, on dit que le plongement de N dans M est un plongement sage. Sinon, ce plongement sera qualifié de sauvage.

Un des plongements sauvages les plus célèbres est la sphère cornue d'Alexander, représenté dans la Figure 1. Il s'agit d'un plongement sauvage de la sphère S^2 dans la 3-sphère S^3 . En effet, cette partie de S^3 est homéomorphe à S^2 , mais une composante de son complémentaire (la partie extérieure sur la Figure 1) n'est pas simplement connexe. Il est donc impossible de l'envoyer sur une sphère sage via un homéomorphisme ambiant.




Figure 1 – La sphère cornue d'Alexander.

Notons W cette composante non simplement connexe du complémentaire de la sphère cornue d'Alexander ∂W et notons $\overline{W} = W \cup \partial W$ son union avec la sphère cornue. Bing démontra que le recollement $\overline{W} \sqcup_{\partial W} \overline{W}$ de deux copies de \overline{W} selon leur bord ∂W est homéomorphe à la 3-sphère S^3 .

L'échange de ces deux copies de \overline{W} produit donc une involution i : l'involution de Bing. Cette involution de S^3 a la propriété d'avoir un ensemble de points fixes sauvagement plongé dans S^3 , ce qui montre que cette involution n'est pas conjuguée à une involution lisse.

$$\text{L'involution de Bing : } \overline{W} \sqcup_{\partial W} \overline{W} \xrightarrow{i} S^3$$

La théorie de Smith nous assure que les points fixes d'un homéomorphisme d'ordre fini sur une 3-variété compacte forment une sous-variété topologique compacte, mais nous souhaitons éviter les cas pathologiques similaires à l'involution de Bing. En fait, il suffit de s'assurer que l'ensemble des points fixes est sagement plongé pour montrer que l'homéomorphisme qui nous intéresse est conjuguée à un difféomorphisme. Plus précisément, nous avons le résultat suivant.

Théorème. [KL88, Corollary 2.3] *Une action par homéomorphismes d'un groupe cyclique G sur une 3-variété compacte M est lissable si et seulement si, pour tout sous-groupe H de G , l'ensemble des points fixes M^H est sagement plongé dans M .*

Pour utiliser ce théorème, nous démontrerons et utiliserons le critère de sagesse suivant.

Proposition. *Soit Σ un sous-variété topologique compacte d'une 3-variété compacte M . Si son complémentaire $M \setminus \Sigma$ est homéomorphe à l'intérieur d'une variété compacte X à bord et que l'inclusion $i : M \setminus \Sigma \rightarrow M$ s'étend continûment en une application de X dans M*

$$\begin{array}{ccc} & X & \\ \uparrow & \nearrow & \\ M \setminus \Sigma & \xrightarrow{i} & M \end{array}$$

alors Σ est sagement plongé dans M .

En particulier, s'il existe un flot dans un voisinage de Σ pour lequel toutes les lignes de flot convergent uniformément vers un point de Σ , cela démontrera que Σ est sagement plongé dans M , concluant la preuve de notre résultat principal.

Les homéomorphismes bilipschitziens

Soit $\sigma : M \rightarrow M$ un homéomorphisme bilipschitzien d'ordre fini sur une 3-variété riemannienne compacte. Notre but est d'utiliser le critère précédent pour démontrer la conjecture formulée par Freedman.

En fait, nous démontrons la proposition suivante :

Proposition. *Pour $\varepsilon = \frac{1}{4000}$ et pour toute action $(1 + \varepsilon)$ -bilipschitzienne d'un groupe fini G sur une variété Riemannienne compacte M , il existe un flot dans un voisinage de l'ensemble des points fixes M^G pour lequel toutes les lignes de flot convergent uniformément vers un point de M^G .*

Le flot ainsi construit va nous permettre d'identifier la topologie du complémentaire de M^G dans un voisinage de M^G et de montrer comment une compactification de cet ouvert peut s'étendre au voisinage de M^G tout entier.

Pour définir ce flot, nous raisonnons par analogie avec les isométries de \mathbb{R}^n : soit x un point de \mathbb{R}^n et σ une isométrie d'ordre fini sur \mathbb{R}^n . Le barycentre de l'orbite de x par σ est alors un point fixe pour σ . Comme une application $(1 + \varepsilon)$ -bilipschitzienne est presque une isométrie, cette analogie nous indique une direction vers laquelle on peut souhaiter voir se diriger notre flot.

Ainsi, nous définissons un champ de vecteurs \mathbf{v} dans un voisinage de M^G dont l'expression dans une carte est :

$$\mathbf{v}(x) = B(x) - x$$

où $B(x)$ est le barycentre de l'orbite de x par G .

$$B(x) = \frac{1}{n} \sum_{g \in G} (x)$$

Les applications $g \mapsto gx$ étant bilipschitziennes, l'application B hérite de ce caractère bilipschitzien. Le champ de vecteurs \mathbf{v} obtient donc également une régularité bilipschitzienne. On peut donc intégrer le champ de vecteurs \mathbf{v} pour en obtenir les lignes de flot. De plus, cette régularité permet de montrer que ce champ de vecteurs induit un flot φ agissant par homoméomorphismes.

La borne sur ε va nous permettre de démontrer que, pour tout point x suffisamment proche de M^G , la ligne de flot $\mathbb{R}^+ \rightarrow M : t \mapsto \varphi_t(x)$ converge vers un point de M^G . Plus précisément, nous démontrons le Lemme suivant.

Lemme. *Il existe des constantes $\tau > 0$ et $k < 1$ ainsi qu'un voisinage V de M^G tels que nous avons l'inégalité*

$$\|\mathbf{v}(\varphi_\tau(x))\| \leq k \|\mathbf{v}(x)\|$$

pour tout x de V .

En particulier, la convergence du flot vers M^G est localement uniforme.

Nous définissons ensuite une surface X intersectant chaque ligne de flot en un unique point. Par l'action du flot φ , le voisinage de M^G privé de M^G est homéomorphe à $X \times [0, 1[$.

$$\begin{aligned} X \times [0, 1[&\rightarrow M \\ (x, t) &\mapsto \varphi_{\frac{1}{1-t}}(x) \end{aligned}$$

Cet homéomorphisme peut alors être étendu en une application de $X \times [0, 1]$ dans un voisinage entier de M^G .

$$X \times [0, 1] \rightarrow M$$
$$(x, t) \mapsto \begin{cases} \varphi_{\frac{1}{1-t}} & \text{si } t < 1 \\ \lim_{t \rightarrow \infty} \varphi_t(x) & \text{si } t = 1 \end{cases}$$

Nous pouvons alors appliquer le critère présenté dans la section précédente, ce qui montre que M^G est sagement plongé dans M , terminant la preuve de notre théorème principal.

APPRENTISSAGE STATISTIQUE DES VARIÉTÉS

Considérons un jeu de données X représenté par une suite de n points dans l'espace à s dimension \mathbb{R}^s . Un des buts de la science des données est d'extraire des informations de ce jeu de données pour mieux le comprendre ou pour extrapoler de futur résultats.

Dans de nombreuses situations, les données ont tendance à s'aligner selon certaines directions particulières dans \mathbb{R}^s ou à s'organiser autour de structures aux propriétés géométriques remarquables. En particulier, de nombreux jeux de données semblent avoir été tirés sur des sous-variétés de dimension d de \mathbb{R}^s . Si s est très grand devant d , il peut être intéressant de projeter nos données dans un espace à d dimensions pour faciliter les calculs. Certains algorithmes donnent également de meilleurs résultats si la dimension intrinsèque d des données est connue. Dès lors, développer des méthodes permettant de déterminer cette dimension intrinsèque devient un problème important dans le domaine de l'apprentissage statistique des variétés (*manifold learning*).

Estimation de la dimension intrinsèque

Soit M une sous-variété de dimension d de \mathbb{R}^s et soit X une suite de n points dans M . Notre but est d'estimer d depuis la seule connaissance de X .

Il convient tout d'abord de remarquer que, sans plus de contraintes, il n'est pas possible de résoudre ce problème. En effet, pour tout nuage de points X et pour tout entier $\tilde{d} \leq s$, il existe toujours une sous-variété de dimension \tilde{d} passant par tous les points de X . Ceci n'est cependant plus toujours possible si l'on impose une borne inférieure sur la courbure de M . Cela revient en fait à fixer une échelle ϵ à laquelle on étudie la sous-variété M . Les objets géométriques




Figure 2 – Quelle est la dimension d'un grillage ?

peuvent en effet faire apparaître des dimensions différentes selon la taille caractéristique que l'on considère. Par exemple un grillage peut paraître de dimension 0, 2, 1 ou 3 selon qu'il soit observé à l'échelle du kilomètre, du mètre, du centimètre ou du micromètre.

Nous définissons alors la portée τ de M comme le plus grand nombre tel que tout point à distance au plus τ de M ait un unique plus proche point sur M . Pour fixer une distance caractéristique, nous pouvons donc limiter notre étude aux variétés pour lesquelles τ est plus grand qu'une valeur donnée.

Méthodes d'approximations

Il existe de nombreux algorithmes dans la littérature pour estimer la dimension intrinsèque d d'un jeu de données X . Ces méthodes utilisent différentes informations extraites de X , comme par exemple :

- **La distance entre les points**, utilisée par exemple pour le calcul de la dimension de corrélation. Le nombre points dans une boule de rayon r étant de l'ordre de $C \cdot r^d$ pour une constante C , il est possible d'estimer la dimension d en comptant le nombre de points dans différentes boules de rayons différents. Par exemple, s'il y a N_1 paires de points à distance au plus ϵ_1 et N_2 paires de points à distance au plus ϵ_2 , la dimension de X peut être approximée grâce à la formule

$$d \simeq \text{Arrondi} \left(\frac{\log(N_1) - \log(N_2)}{\log(\epsilon_1) - \log(\epsilon_2)} \right). \quad (1)$$

- **L'angle entre les points**, il est possible d'estimer la dimension de X en calculant la variance des angles entre les points proches, grâce à une méthode décrite dans [DQV19a].
- **Les directions de variance maximale**, utilisée dans l'analyse en composantes principales (PCA). La dimension est estimée en comptant le nombre de valeurs propres élevées dans la matrice de covariance du jeu de données.
- **Le type d'homotopie global de M** peut être calculé si l'on dispose de suffisamment de points. On peut alors en extraire la dimension de X . Cette méthode est étudiée dans [NSW08a].

Taille minimale du jeu de données

La question consistant à déterminer la taille minimale que doit avoir un jeu de données pour pouvoir estimer sa dimension grâce à ces méthodes a été relativement peu étudiée dans la littérature.

Dans [NSW08a], les auteurs proposent une estimation du nombre de points nécessaires pour retrouver avec une probabilité donnée le type d'homotopie de M depuis un échantillon X de points tirés indépendamment et uniformément sur M . Dans [Bre+18a], les auteurs estiment que pour une sous-variété de dimension 4, de volume 1000 avec $\tau > 1$, si on souhaite retrouver le type d'homotopie de M avec une probabilité de 90%, le nombre de points nécessaires donné par la méthode de [NSW08a], doit être d'au moins 1 592 570 365. En précisant ce calcul, il est possible de réduire cette estimation à 21 415 600 points.

Il n'est cependant pas nécessaire de déterminer complètement le type d'homotopie de M pour connaître sa dimension. L'estimation proposée dans [NSW08a] calcule en fait le nombre de points nécessaires pour que M soit totalement recouvert par un ϵ -voisinage de X avec une certaine probabilité. La dimension d'une sous-variété étant une caractéristique locale, il n'est pas nécessaire d'avoir des points sur la totalité de cette sous-variété pour en estimer la dimension. Un algorithme calculant la dimension de corrélation grâce à l'estimateur (1) devrait donc nécessiter beaucoup moins de points pour estimer la dimension de M avec la même probabilité.

Résultats

Nous nous proposons donc de fournir une estimation du nombre de points nécessaires pour retrouver la dimension intrinsèque de X en calculant sa dimension de corrélation grâce à l'estimateur (1).

Notre premier résultat est une borne supérieure sur le nombre de point minimal nécessaire pour obtenir la bonne dimension avec probabilité d'au moins 90%. Ce résultat se base sur l'étude des propriétés géométriques des variétés de portée 1 et sur les propriétés statistiques de l'estimateur (1).

Il est important de noter que cet estimateur nécessite le choix de deux échelles de calcul ϵ_1 et ϵ_2 . Les performances de l'estimateur dépendant de ce choix, nous avons dû optimiser ces valeurs pour obtenir les bornes les plus basses possibles. Cette optimisation a été réalisée de manière numérique, et nous obtenons le théorème suivant.

Théorème. Pour $d = 1, \dots, 10$ posons ϵ_1 et ϵ_2 fournis dans le tableau suivant. Soit M une sous-variété compacte de dimension d dans \mathbb{R}^s , de portée $\tau > 1$ et de volume $\text{vol}(M)$. Posons n fourni par le tableau. Si nous tirons n points indépendamment et uniformément sur M , alors l'estimateur (1) donnera la dimension d avec probabilité d'au moins 90%.

d	ϵ_1	ϵ_2	n
1	1.5	0.19	$9 + 21 \cdot \text{vol}(M)^{\frac{1}{2}}$
2	0.78	0.2	$94 + 58 \cdot \text{vol}(M)^{\frac{1}{2}}$
3	0.63	0.23	$635 + 146 \cdot \text{vol}(M)^{\frac{1}{2}}$
4	0.54	0.23	$2786 + 392 \cdot \text{vol}(M)^{\frac{1}{2}}$
5	0.46	0.22	$7013 + 1119 \cdot \text{vol}(M)^{\frac{1}{2}}$
6	0.4	0.21	$13221 + 3366 \cdot \text{vol}(M)^{\frac{1}{2}}$
7	0.36	0.21	$25138 + 10644 \cdot \text{vol}(M)^{\frac{1}{2}}$
8	0.33	0.2	$50033 + 34890 \cdot \text{vol}(M)^{\frac{1}{2}}$
9	0.31	0.19	$63876 + 119533 \cdot \text{vol}(M)^{\frac{1}{2}}$
10	0.29	0.18	$139412 + 425554 \cdot \text{vol}(M)^{\frac{1}{2}}$

Pour une variété M de dimension 4 et de volume 1000, on obtient donc que 15 182 points suffisent à déterminer la dimension de M avec probabilité 90%, ce qui est une estimation plus de 1 000 fois moins élevée que ce qui est nécessaire pour retrouver le type d'homotopie de M grâce au théorème de [NSW08a].

Notre second résultat provient d'un modèle basé sur un raisonnement heuristique. Ce modèle se base sur l'idée selon laquelle la donnée importante à estimer n'est pas le nombre de points n du jeu de données mais le nombre de paires de points N à distance au plus ϵ_1 .

On peut estimer ce nombre N grâce à la formule

$$N \simeq \frac{n(n-1)}{2} \frac{\text{vol}(B(\epsilon_1))}{\text{vol}(M)} \quad (2)$$

où $\text{vol}(B(\epsilon_1))$ est le volume d'une boule euclidienne de dimension d et de rayon ϵ_1 . On raisonne ensuite comme si l'on avait tiré indépendamment N nombres dont la loi suit celle de la distance entre deux points à distance au plus ϵ_1 dans \mathbb{R}^d . Il est alors possible d'estimer la valeur de N nécessaire pour obtenir une probabilité de 90% et en déduire la valeur de n grâce à la formule (2).

Nous obtenons les valeurs suivantes

d	n
1	$5 \cdot \text{vol}(M)^{\frac{1}{2}}$
2	$12 \cdot \text{vol}(M)^{\frac{1}{2}}$
3	$22 \cdot \text{vol}(M)^{\frac{1}{2}}$
4	$50 \cdot \text{vol}(M)^{\frac{1}{2}}$
5	$128 \cdot \text{vol}(M)^{\frac{1}{2}}$
6	$355 \cdot \text{vol}(M)^{\frac{1}{2}}$
7	$964 \cdot \text{vol}(M)^{\frac{1}{2}}$
8	$2949 \cdot \text{vol}(M)^{\frac{1}{2}}$
9	$9458 \cdot \text{vol}(M)^{\frac{1}{2}}$
10	$33021 \cdot \text{vol}(M)^{\frac{1}{2}}$

Pour une variété M de dimension 4 et de volume 1000, on obtient donc que 1 581 points. L'estimation est donc encore réduite d'un facteur 10 par rapport à la borne obtenue dans le théorème précédent.

Ce modèle est simplifié par rapport à la situation réelle car on ne peut pas considérer les paires indépendamment les unes des autres et car les boules de rayons ϵ_1 dans M ne sont pas nécessairement euclidiennes. Cependant, les résultats obtenus sont confirmés par l'expérience : nous avons créé un programme informatique [Gri22] permettant de tirer des points sur différentes variétés puis de calculer la dimension de corrélation. Avec le nombre de points estimés par ce modèle heuristique, nous avons obtenus des taux de réussite de l'ordre de 90% sur chaque exemple, comme prévu.

PART I

The Smith conjecture in low regularity

Every first appearance of a new technical term will be written in **bold**.

We will note

- \mathbb{R}^n the Euclidean space of dimension n .
- S^n the sphere of dimension n : $\{x \in \mathbb{R}^{n+1} \mid \|x\|_2 = 1\}$
- \mathring{B}^n the open ball of dimension n : $\{x \in \mathbb{R}^n \mid \|x\|_2 < 1\}$
- \overline{B}^n (or D^n) the closed ball (or disk) of dimension n : $\{x \in \mathbb{R}^n \mid \|x\|_2 \leq 1\}$
- X^σ the fixed set of X under the self-map σ : $\{x \in X \mid \sigma(x) = x\}$
- $A \setminus B$ the difference of two sets A and B : $\{x \in A \mid x \notin B\}$

Let A be a subset of a topological space X , we will note

- \mathring{A} the interior of A
- \overline{A} the closure of A
- ∂A the boundary of A : $\overline{A} \setminus \mathring{A}$

In 1939, P. A. Smith proved [Smi39, 7.3 Theorem 4] that the fixed set M^σ of a finite-order and orientation-preserving homeomorphism σ of the 3-sphere S^3 was empty or a circle S^1 . He then asked [Eil49, Problem 36] if this circle could be knotted. This question is known as the Smith conjecture. More generally, the question is whether such a map is conjugate to an orthogonal map. Since the work on geometrization by Thurston and Perelman, we know that this is the case for smooth maps [BLP05]. On the other hand, Bing [Bin52] gave an example of a continuous orientation-reversing involution with a wildly embedded 2-sphere as fixed set. Therefore, this involution could not be conjugate to an orthogonal map. Montgomery and Zippin [MZ54] also modified Bing's example to obtain an orientation-preserving involution with a wild circle as a fixed set. Jani Onninen and Pekka Pankka showed in 2019 [OP19] that there also exists wild involutions in the Sobolev class $W^{1,p}$.

One can then wonder what happens for maps with more regularity but which are not differentiable. In [Ham08], D. H. Hamilton announced that quasi-conformal reflections are tame, but the proof seems to remain unpublished. In fact, even the Lipschitz case seems to be considered open. For example, as recently as 2013, Michael Freedman asked in [Fre13, Conjecture 3.21] if the Bing involution could be conjugate to a Lipschitz homeomorphism. Jani Onninen and Pekka Pankka reiterate this question in 2019 [OP19]. We give a partial answer to Freedman's question, proving that for $\varepsilon > 0$ small enough, such wild finite-order maps can not be $(1 + \varepsilon)$ -bilipschitz. More precisely, we will show the following theorem.

Theorem 1. *For $\varepsilon = \frac{1}{4000}$, any action of a finite cyclic group by $(1 + \varepsilon)$ -bilipschitz homeomorphisms on a closed 3-manifold is conjugate to a smooth action.*

Theorem 1 is proved by showing that the fixed set of such an action is always tamely embedded. We will show in Section 2.2.4 how Theorem 1 can be reduced to the tameness of this fixed set.

The first chapter of this part is a discussion about topological tameness. We present the different notions of tameness in topology and how they are linked together. We also give important examples that one has to keep in mind. Secondly, we expose results on the shrinkings of manifolds and in Section 1.4, we state and prove Proposition 1, a tameness criterion that will be at the core of the proof of Theorem 1.

Proposition 1. *Let Σ be a closed topological submanifold of a closed 3-manifold M . Suppose that its complement $M \setminus \Sigma$ is homeomorphic to the interior of a compact manifold X with boundary. If the inclusion $i : M \setminus \Sigma \rightarrow M$ extends to a continuous map from X to M*

$$\begin{array}{ccc} & X & \\ \uparrow & \cdots \searrow & \\ M \setminus \Sigma & \xrightarrow{i} & M \end{array}$$

then Σ is tamely embedded in M .

The second chapter is dedicated to the Smith conjecture and to the proof of Theorem 1. We introduce the ideas which led to this conjecture and we present the results known about it in literature. After stating some results about Lipschitz vector fields, we expose the proof of Theorem 1.

This proof is done by proving the following proposition

Proposition 2. *For $\varepsilon = \frac{1}{4000}$ and for every $(1 + \varepsilon)$ -bilipschitz action of a finite group G on a compact Riemannian manifold M , the fixed set M^G satisfies the conditions of Proposition 1.*

and by showing how Theorem 1 can be reduced to Proposition 1 and Proposition 2.

TOPOLOGICAL TAMENESS

1.1 Introduction

1.1.1 The Jordan-Brouwer separation theorem

A **simple closed curve** in the plane is an injective continuous map from the circle S^1 into \mathbb{R}^2 . The well-known Jordan-Brouwer separation theorem states that such a curve divides the plane into two components. Only one of these two components is bounded and is called the interior of γ . The other component is called the exterior.

Theorem 2 (Jordan-Brouwer separation theorem [Jor87]). *Let γ be a simple closed curve in \mathbb{R}^2 . Then $\mathbb{R}^2 \setminus \gamma$ has exactly two connected components.*

This theorem is not the strongest result that one can get. For example it does not describe the topology of the two components. However, we will see in section 1.1.2 that it generalizes well to higher dimensions.

An other result, the Jordan-Schoenflies theorem, completely describes the situation from topological point of view. It tells us that any simple closed curve has every topological property of the standard embedding of S^1 into \mathbb{R}^2 .

Theorem 3 (Jordan-Schoenflies theorem [Sch06]). *Let γ be a simple closed curve in \mathbb{R}^2 . Then there is a homeomorphism of \mathbb{R}^2 into itself that sends γ on the standard circle $\{x \in \mathbb{R}^2 \mid \|x\|_2 = 1\}$.*

This theorem is stronger than the Jordan-Brouwer separation theorem. For example, it gives us the following properties:

- (A) The interior of a simple closed curve is homeomorphic to an open disk.
- (B) The union of a simple closed curve with its interior is homeomorphic to a closed disk.




Figure 1.1 – Illustration of the Jordan-Schoenflies theorem

In other words, the point (B) says that the homeomorphism from the open disk to the interior of γ obtained in point (A) can be chosen so that it extends homeomorphically to γ .

1.1.2 Generalizations to higher dimensions

The Jordan-Brouwer separation theorem can be generalized to higher dimensions. Namely, one can show that any injective map from the sphere S^{n-1} into the sphere S^n cuts the latter into two pieces. This can be proved using **Alexander duality**, which is a result in homology theory linking the cohomology of a subset of the n -sphere S^n with the homology of its complement.

Theorem 4 (Alexander duality [Hat02][Corollary 3.45]). *Let X be a compact and locally contractible subset of S^n . There is an isomorphism*

$$\tilde{H}_k(S^n \setminus X) \simeq \tilde{H}^{n-k-1}(X)$$

for every $k \geq 0$ and where \tilde{H} stands for reduced (co)homology with coefficient in \mathbb{Z} .

In particular, this theorem asserts that the homology of the complement of X only depends on the topology of X , and not on the way it is embedded. Applying this theorem with X being homeomorphic to the $(n-1)$ -sphere S^{n-1} , we see that the complement of X has the homology of the complement of the standard $(n-1)$ -sphere. The complement of a subset of S^n homeomorphic to a $(n-1)$ -sphere is thus always the union of two

homology balls.

The Jordan-Schoenflies theorem, however, does not generalize to higher dimensions. Indeed, homology balls can have wild topologies. For example, the wild complement of the Alexander horned sphere (see Section 1.1.3) is not simply connected and its fundamental group is not even finitely generated.

However, under stronger conditions on the regularity of the embedding of S^{n-1} this theorem generalizes to every dimension. The generalized Schoenflies theorem states that this is the case if X is locally flat.

We say that an embedding of S^{n-1} in S^n is locally flat if it can be extended to an embedding of $S^{n-1} \times [-1, 1]$, which agrees with the first embedding on $S^{n-1} \times \{0\}$.

Theorem 5 (generalized Schoenflies theorem [Bro60]). *Let X be a subset of S^n homeomorphic to S^{n-1} . If X is locally flat, there is a homeomorphism of S^n into itself that sends X on the standard $(n - 1)$ -sphere.*

1.1.3 The Alexander horned sphere

We saw that any subset X of S^3 homeomorphic to S^2 must divide S^3 into two connected components. However, Alexander duality does not provide an ambient homeomorphism of S^3 sending X on the standard S^2 as in the Jordan-Schoenflies theorem. In particular, we cannot conclude that the complement of X in S^3 is the disjoint union of two open balls. We will present an important example with this behavior.




Figure 1.2 – The Alexander horned sphere

In [Ale24], James Waddell Alexander described a subset of S^3 homeomorphic to S^2 with a complement which is not simply connected. This object is called the **Alexander horned sphere**. To construct it, start with a sphere and grow two horns out of it. Divide

each of these horns into two other horns and interlace one horn of the first division with one horn of second division. Repeat this process with each new pair of horns over and over infinitely many times with smaller and smaller horns until they accumulate on a Cantor set.




Figure 1.3 – Two interlaced pairs of pants.

To understand better this object, we can decompose it into "pairs of pants". A pair of pants is a surface homeomorphic to a 2-sphere with 3 open disks removed or, equivalently, to a closed disk with two open disks removed.




Figure 1.4 – Decomposition of the Alexander horned sphere.

As shown on Figure 1.4, the Alexander horned sphere can be decomposed into a disk, uncountably many pairs of pants, and a limit Cantor set. This decomposition defines an explicit homeomorphism between the Alexander horned sphere and the standard 2-sphere S^2 .

Proposition 3. *The exterior of the Alexander horned sphere is not simply connected.*

Proof. Define a loop γ around one of the first horns, as in Figure 1.5. We will show that γ cannot be shrunk in the exterior of the Alexander horned sphere.

Following the decomposition of Figure 1.4, define a sequence of boxes containing each pair of pants. Note $B_{1,1}$ the box containing the first pair of pairs of pants, $B_{2,1}$ and $B_{2,2}$ the two boxes containing each of following smaller pair of pairs of pants, and continue this notation with boxes $B_{i,j}$ for $1 \leq i$ and $1 \leq j \leq 2^{i-1}$. We choose the boxes $B_{i,j}$ so that their diameters go to zero with i .

Note A the Alexander horned sphere and suppose that there exists a homotopy $h : S^1 \times [0, 1] \rightarrow S^3 \setminus A$ with $h_0 = \gamma$ and h_1 constant. Note $X = h(S^1 \times [0, 1])$.

It is clearly impossible to shrink γ in the complement of $A \cup_j B_{i,j}$ for any $1 \leq i$, so X intersects $\bigcup_j B_{i,j}$ for every $1 \leq i$. For every i , choose a index $J(i)$ such that X intersects $B_{i,J(i)}$. So X intersects every closed set $K_n = \overline{\bigcap_{i \geq n} B_{i,J(i)}}$ and, as X is closed and as $K_{n+1} \subset K_n$ for every n , X intersects $\bigcap_n K_n$. As A is also closed and intersects every $B_{i,j}$, A also intersects $\bigcap_n K_n$. As the diameter of K_n goes to zero with n , $\bigcap_n K_n$ is a point, so X must intersect A . So γ cannot be shrunk in $S^3 \setminus A$. \square




Figure 1.5 – Decomposition of the Alexander horned sphere

1.2 Wildness and tameness

The Alexander horned sphere is an example of what we call a wild object. In the following, we will often use the terms wild and tame to qualify some objects or maps. These terms may have different meaning in different contexts. The two main concepts related to these terms are the concepts of tame manifolds and of tame embeddings. The term wild will be use in the same contexts to refer to "non-tame" objects.

1.2.1 Wild and Tame embeddings

Let Σ be a compact subset of a m -manifold M . If Σ is a topological n -manifold for the topology induced by M , we say that Σ is a **topological submanifold** of M . We say that Σ is **tamely embedded** (or simply that Σ is tame) if there is a self-homeomorphism of M sending Σ to a polyhedron.



Figure 1.6 – A wild arc.

In \mathbb{R}^2 , every continuous arc (i.e. a continuous map from $[0, 1]$ to \mathbb{R}^2) is tame, but wild behaviors can appear in higher dimensions. For example, the arc illustrated on Figure 1.6 is a wild arc (see [FA48]). In this example, the arc is tame everywhere, except at its two endpoints. This means that every other point has a tamely embedded neighborhood in the arc.

A method to detect if a topological submanifold is wild is to study the topology of its complement. If this complement is not homeomorphic to the complement of a polyhedron, the topological submanifold is necessarily wild. For example, the complement of the arc illustrated on Figure 1.6 in S^3 is not simply connected, but the complement of a polyhedral arc in S^3 is necessarily a ball.



Figure 1.7 – A wild arc with a complement homeomorphic to a ball.

However, this condition is not sufficient. Indeed, there exist wild arcs in S^3 whose complement are homeomorphic to open balls. (see Section 1.2.3)

Wild embeddings of spheres can also be obtain by thickening the previous wild arcs, with the same discussion concerning their complements.

1.2.2 Wild and tame manifolds

We saw that a topological submanifold is wild if we can show that its complement has a sufficiently "bad" behavior, this lead to the notion of wild and tame manifolds. We say that manifold M is a **tame manifold** if it is homeomorphic to the interior of a compact manifold with boundary or, more generally, if it is homeomorphic to the complement of a closed subset of the boundary of a compact manifold. We also say that M is a **missing boundary manifold**.

The problem of knowing whether a manifold is tame or not, known as the missing boundary problem, is not always an easy task. Even if the answer is affirmative, finding the missing boundary explicitly can be challenging. In dimension three, a necessary condition for M to be tame is to have a finitely generated fundamental group. For example, one can show that the fundamental group of the exterior component of the complement of the Alexander horned sphere is not finitely generated, which shows that this component is wild manifold.

More generally, the following tameness criterion due to Tucker [Tuc74] states it is sufficient to consider the fundamental groups of complements of compact polyhedron.

Theorem 6 (Tucker). *Let M be a connected, P^2 -irreducible 3-manifold. Then M is tame if and only if for any compact polyhedron C in M , the fundamental group of each component of $M \setminus C$ is finitely generated.*

We say that a manifold M is P^2 -irreducible if every smooth sphere in M bounds a ball and if M contains no 2-sided real projective plane. In the theorem, we can replace "polyhedron" by "smooth compact submanifold", by "tame compact submanifold" or even by "tame arcs" (see [Mes77][§3]).

We now give some examples of wild manifolds.

Example 1. Consider an infinite concatenation of knots C in \mathbb{R}^3 as in Figure 1.8.



Figure 1.8 – An infinite concatenation of knots.

More precisely, let l be a smooth injective arc joining $(0, 0, 0)$ to $(1, 0, 0)$ in $[0, 1] \times \mathbb{R}^2$ and let $C = \bigcup_n l + (n, 0, 0)$. We choose l so that l and $l + (1, 0, 0)$ only intersect in $(1, 0, 0)$ and so that they meet smoothly at this point. We also ask l to be non-trivially knotted.

The space $\mathbb{R}^3 \setminus C$ is not a tame manifold. Indeed, we will show that its fundamental group is not finitely generated. Choose a base point x_0 and suppose by contradiction that there exists a finite number of generators $(\gamma_1, \dots, \gamma_n)$. As $\bigcup_i \gamma_i$ is bounded, there is a $N \in \mathbb{N}$ such that, for any loop γ based on x_0 , there is a homotopy sending γ inside $] - \infty, N] \times \mathbb{R}^2$.




Figure 1.9 – A loop γ far away from generators of $\pi_1 (]-\infty, N] \times \mathbb{R}^2)$.

By Seifert-Van Kampen theorem, the fundamental group of $\mathbb{R}^3 \setminus C$ is

$$\pi_1 \left((]-\infty, N] \times \mathbb{R}^2) \setminus C \right) \underset{\pi_1((\{N\} \times \mathbb{R}^2) \setminus C)}{*} \pi_1 \left(([N, +\infty[\times \mathbb{R}^2) \setminus C \right).$$

In the light of the foregoing, this means that the map from $\pi_1 \left((\{N\} \times \mathbb{R}^2) \setminus C \right)$ to $\pi_1 \left(([N, +\infty[\times \mathbb{R}^2) \setminus C \right)$ is surjective. However, this fact is known to be false: there exist loops in $([N, +\infty[\times \mathbb{R}^2) \setminus C$ (for example, the loop γ on Figure 1.9) that are not homotopic to a loop of $(\{N\} \times \mathbb{R}^2) \setminus C$ inside $([N, +\infty[\times \mathbb{R}^2) \setminus C$.

Remark 1. Note that, in the preceding example, C is smooth but not compact and \mathbb{R}^3 is a tame manifold. This shows the importance of the compactness of C in Theorem 6.

It is also important to remark that we cannot only consider polyhedrons in the interior of M . After proving a preliminary lemma, we illustrate this fact by an example.

Lemma 1. *Let X be a topological space and let $(B_i)_{i \in \mathbb{N}}$ be a sequence of subsets of X such that*

- $X = \bigcup_i B_i$
- each B_i is homeomorphic to an open 3-ball
- $\overline{B_i} \subset B_{i+1}$ for every i .

Then X is homeomorphic to a open 3-ball.

Proof. Let $B(0, i)$ be the open 3-ball $\{p \in \mathbb{R}^3 \mid \|p\|_2 < i\}$. Each B_i is homeomorphic to $B(0, i)$, but this fact cannot be used to define a global homeomorphism between X and \mathbb{R}^3 , because $B_{i+1} \setminus B_i$ is not necessarily homeomorphic to $S^2 \times]0, 1]$. For example, the inclusion $B_i \rightarrow B_{i+1}$ could be an Alexander horned sphere.

Let $f_i : B(0, 1) \rightarrow B_{i+1}$ be a homeomorphism between the standard open 3-ball and B_{i+1} . There is an $\varepsilon_i > 0$ such that $\overline{B_i} \subset f_i(B(0, 1 - \varepsilon_i))$. The sequence $(f_i(B(0, 1 - \varepsilon_i)))_{i \in \mathbb{N}}$ then verifies the three hypothesis of the proposition and each sphere $\partial(f_i(B(0, 1 - \varepsilon_i)))$ is flat. Then, by Schoenflies theorem, we can find for each i a homeomorphism $g_i : B(0, i) \rightarrow f_i(B(0, 1 - \varepsilon_i))$ which agrees with every g_j for $j < i$.

This defines a global homeomorphism between \mathbb{R}^3 and X . □

Example 2. *Consider two cylinders C_0 and C_1 , and a smooth embedding $f : C_0 \rightarrow C_1$ as shown on Figure 1.10.*




Figure 1.10 – Embedding $f : C_0 \rightarrow C_1$.

Let $h_i : D^2 \times [0, 1] \rightarrow C_i$ be a homeomorphism, note A_i and B_i the two disks $h_i(D^2 \times \{0\})$ and $h_i(D^2 \times \{1\})$. The embeddings h_i are chosen so that A_0 is smoothly embedded in the interior of A_1 and so that B_0 is smoothly embedded in the interior of B_1 .

The disjoint union of C_0 and C_1 quotiented by the map f is then homeomorphic to C_1 . Now, consider a sequence $(C_n)_{n \in \mathbb{N}}$ of cylinders and $f : C_n \rightarrow C_{n+1}$ as previously. Finally, let X be the disjoint union of every C_n quotiented by f . Here, "quotienting by f " we mean quotienting by the equivalence relation

$$x \sim y \Leftrightarrow \exists n \in \mathbb{N}, x = f^n(y) \text{ or } y = f^n(x),$$

$$X = \bigsqcup_i C_i / \sim.$$

We note π the quotient map associated to this equivalence relation.

The space X is a smooth manifold, but it is not a tame manifold. Indeed, the arc $\gamma = \{(0, 0)\} \times [0, 1]$ is smooth in C_0 , and $\pi(\gamma)$ is smooth in X . However, the fundamental group of the complement of $\pi(\gamma)$ is not finitely generated.




Figure 1.11 – The arc γ in C_0 .

Indeed, the interior of X can be written as

$$\mathring{X} = \bigcup_i \pi(h_i(D^2 \times [\frac{1}{i}, 1 - \frac{1}{i}])).$$

So, by Lemma 1, \mathring{X} is homeomorphic to \mathbb{R}^3 , and one can show that $\mathring{X} \setminus \pi(\gamma)$ is homeomorphic to $\mathbb{R}^3 \setminus C$ in the previous example. So, the complement of $\pi(\gamma)$ has a fundamental group which is not finitely generated. Finally, by Theorem 6, X is a wild manifold, even if its interior is tame.

1.2.3 The Fox-Artin sphere

We saw with Theorem 6 that the complement of a compact tame submanifold of tame manifold is also a tame manifold. We will now discuss the opposite question: if a compact topological submanifold of a tame manifold has a tame complement, is this submanifold tamely embedded ?

For example, consider the following question.

Question 1. *Let K be a compact subset of \mathbb{R}^3 . If*

- *\mathring{K} is homeomorphic to an open ball B^3*
- *∂K is homeomorphic to a 2-sphere S^2*

Is K homeomorphic to a closed ball D^3 ?

A difficulty involved in trying to answer this question is that the homeomorphism between B^3 and \mathring{K} does not necessarily extend to a homeomorphism from D^3 to K .

Remark that the Alexander horned sphere does not provide a counterexample to this question. Indeed, the union of the Alexander horned sphere with its "nice" complement is homeomorphic to a closed 3-ball, and its wild complement is not homeomorphic to an open 3-ball.

In 1948, Ralph Fox and Emil Artin described a wild sphere providing a negative answer the preceding question [FA48][Example 3.2]. They started with an arc created from an infinite number of interlacing loops getting smaller and smaller, converging to an end point.




Figure 1.12 – A half Fox-Artin arc.

This arc, which is half of what is known as a **Fox-Artin arc**, is tame everywhere except at one of its end points. By thickening this arc with a thickness approaching zero near its wild end point, we obtain a **Fox-Artin sphere**, which has the following remarkable properties :

- It is homeomorphic to a 2-sphere
- Its complement in S^3 is homeomorphic to the disjoint union of two open 3-balls
- It is wildly embedded : its union with its exterior does not yield a closed 3-ball

We provide short proofs of these facts

Lemma 2. *The solid Fox-Artin sphere (the Fox-Artin sphere together with its interior) is a closed 3-ball.*

Proof. The solid Fox-Artin sphere is the one-point compactification of an infinite connected sum of closed balls. It is thus also a closed ball. □




Figure 1.13 – The Fox-Artin sphere.

It is clear that the exterior of the Fox-Artin sphere is simply connected, as every closed arc inside it can be slid all over the ball. However, this is not sufficient to show that it is homeomorphic to a ball.

Lemma 3. *The exterior of the Fox-Artin sphere is an open 3-ball.*

Proof. We begin by showing that the complement of a half Fox-Artin arc is an open ball. As only one of the ends of the arc is wild, we can slide the other one to shrink the arc to a point. More precisely, the solid Fox-Artin sphere is a closed ball containing the half Fox-Artin arc, and this arc is tamely embedded in the solid Fox-Artin sphere, meaning that the homeomorphism between the solid Fox-Artin sphere and the standard closed ball brings the half Fox-Artin arc to a tame arc.




Figure 1.14 – Homeomorphism between the solid Fox-Artin sphere and the standard closed ball.

Every tame arc in a closed ball meeting its boundary only at one of its end points can be shrunk: there is always a continuous map from the closed ball to itself fixing its boundary, sending the tame arc to a point and being a homeomorphism out of the arc. This map can be pulled back in the solid Fox-Artin sphere to define a homeomorphism between the complement of the half Fox-Artin arc and the complement of a point in S^3 , which is an open 3-ball.

In fact, the exact same argument also work if we replace the half Fox-Artin arc by a solid Fox-Artin sphere, using a thicker solid Fox-Artin sphere containing the first one. \square

Lemma 4. *The Fox-Artin sphere is wildly embedded.*

Proof. The Fox-Artin sphere together with its exterior is not a closed ball. Indeed, the complement of any tame arc in a tame manifold has finitely-generated fundamental group. However, we can define a tame arc γ in the exterior of the Fox-Artin sphere that meet the sphere at its two endpoints and whose complement in the exterior of the sphere is not finitely generated.




Figure 1.15 – A tame arc with a wild complement.

A complete proof of this lemma can be found in [FA48]. \square

1.2.4 Some objects seem wild but are not

Wild arcs and wild spheres often contain infinitely repeating patterns, but it is not always easy to determine if such objects are really wild. For example, consider the arc created by the concatenation of infinitely many knots, getting smaller and smaller, as in Figure 1.16.



Figure 1.16 – Is it wild ?

This arc seems to be wild, but it is not easy to prove this fact since its complement is clearly simply connected. In fact, this arc is tame, and we can describe an explicit homeomorphism of \mathbb{R}^3 sending it to a linear segment.

First, remark that we can separate each knot with spheres as in Figure 1.17. Each sphere intersects the arc transversely at a unique point. These spheres divide the space in an infinite number of regions homeomorphic to $S^2 \times [0, 1]$ crossed by a smooth arc joining its to boundaries. The topology of the complement of this arc in these regions is in fact easily determined by what is known as the **lamp-cord trick**.




Figure 1.17 – Dividing \mathbb{R}^3 into a decreasing sequence of balls.

Proposition 4 (Lamp-cord trick). *Let X be the manifold $S^2 \times [0, 1]$ and let γ be a smooth arc joining the two boundaries of X . Then the pair (X, γ) is homeomorphic to the standard pair (X, γ') where γ' is a straight arc $\{p\} \times [0, 1]$ for a point $p \in S^2$.*

Proof. The lamp-cord trick is named after its very visual proof. If we embedded X in \mathbb{R}^3 and make the inner sphere very small, the arc γ resembles a cable joining the outer sphere to a light-bulb. This lamp-cord can then be unknotted inside the outer sphere to obtain a straight arc. This unknotting defines a self-homeomorphism of X sending γ to a straight arc. \square

The homeomorphism described by the lamp-cord trick can be chosen so that it fixes the boundaries of $S^2 \times [0, 1]$. For the arc of Figure 1.16, this homeomorphism can be applied in each region separated by the spheres simultaneously. This defines a homeomorphism of \mathbb{R}^3 sending the arc on a straight line.

This trick can however only be used for concatenations of knots that can be separated by spheres.




Figure 1.18 – Illustration of the lamp-cord trick.



Figure 1.19 – An arc on which the lamp-cord trick cannot be used.

1.2.5 Local flatness

A criterion to show that a topological submanifold is tamely embedded is to check if it has a tubular neighborhood. We say that a topological n -submanifold Σ of a m -manifold M has a **tubular neighborhood** if there is a vector bundle E of basis Σ and of fibers \mathbb{R}^{m-n} whose total space can be embedded as an open set in M and whose null section is Σ .

One can give a direct proof of this fact, but we prefer to sketch a quicker argument in dimension 3, using the uniqueness of smooth structures.

Proposition 5. *Let Σ be a closed topological submanifold of a 3-manifold M . If Σ has a tubular neighborhood E , then Σ is tamely embedded.*

Proof. Let A be the unit disk subbundle of E and let $B = \overline{M \setminus A}$. As A and B have

topological product structures on their boundaries, they are topological manifolds with boundaries. We choose any smooth structure on B and we choose a smooth structure on A that makes Σ smooth.

For these smooth structures, we can choose a diffeomorphism g from ∂B to ∂A which is isotopic to the identity map (see [Mun60, Theorem 6.3]). Gluing A and B along this diffeomorphism gives us a smooth manifold M' homeomorphic to M . To define this homeomorphism, remark that ∂A has a neighborhood in A homeomorphic to $\partial A \times [0, 1]$. Define a map h from M' homeomorphic to M by sending B on itself via the identity map and, in the collar $\partial A \times [0, 1]$, send (x, t) on $(g_t(x), t)$, where $(g_t)_t$ is an isotopy from g to the identity, and using the identity map elsewhere on A . The topological submanifold Σ is then a smooth submanifold of M' .

As M and M' are homeomorphic and as the smooth structure on a 3-manifold is unique, there is a diffeomorphism d between M' and M . The map $d \circ h$ is thus a homeomorphism of M that makes Σ smooth. This shows that the submanifold Σ is tame. \square

1.3 Shrinkings of manifolds

Before introducing our tameness criterion needed for the proof of Theorem 1, we will discuss the topic of shrinkings of manifolds.

Let X be a topological space and $(A_i)_{i \in I}$ be a collection of disjoint subsets of X . Consider the equivalence relation \sim on X by the following condition.

$$x \sim y \Leftrightarrow (x = y) \text{ or } (\exists i \in I, x \in A_i \text{ and } y \in A_i)$$

By "shrinking the A_i ", we mean considering the quotient space X / \sim .

It is sometimes convenient to define a space as such a shrinking. The principal questions that we will discuss here are the conditions under which a shrinking preserves the topology of the starting space, and under which the quotient map $\pi : X \rightarrow X / \sim$ is approximable by homeomorphisms.

An elementary example of such a behavior is the shrinking of a closed ball in Euclidean space \mathbb{R}^n . In dimension 1, consider the shrinking of the closed interval $A = [0, 1]$ in $X = \mathbb{R}$.

The resulting space X/\sim is still homeomorphic to \mathbb{R} , with the following relabeling.

$$\bar{x} \mapsto \begin{cases} x & \text{if } x < 0 \\ 0 & \text{if } x \in [0, 1] \\ x - 1 & \text{if } x > 1 \end{cases}$$

This formula is also the expression of the quotient map $\pi : X \rightarrow X/\sim$ for this new labeling of X/\sim .

$$\pi : X \rightarrow X/\sim$$

$$x \mapsto \begin{cases} x & \text{if } x < 0 \\ 0 & \text{if } x \in [0, 1] \\ x - 1 & \text{if } x > 1 \end{cases} = x + \frac{|x - 1| - |x| - 1}{2}$$

Which is approximable as close as desired by homeomorphisms π_ε .

$$\pi_\varepsilon : X \rightarrow X/\sim$$

$$x \mapsto x + \frac{\sqrt{(x - 1)^2 + \varepsilon} - \sqrt{x^2 + \varepsilon} - 1}{2}$$




Figure 1.20 – Graphs of π and π_ε .

This argument easily generalizes to shrinkings of closed balls in higher dimensional Euclidean spaces.

If A is homeomorphic to a closed ball and is tamely embedded in a manifold, the shrinking of A still yields the manifold X . However, this process fails when A is wildly embedded. For example, if A is the union of the Alexander horned sphere with the tame component of its complement, the shrinking of A in S^3 will yield a space in which the complement of a point can be non-simply connected, which is impossible in S^3 .

The process also becomes harder if there are (possibly uncountably) many subsets A_i to shrink.

1.3.1 Cell-like maps and Moore's Theorem

In dimension 2, a famous shrinking criterion is Moore's Theorem. It states that shrinking a collection $(A_i)_{i \in I}$ of contractible subsets of a closed surface X yields this same surface X and that the quotient map is approximable by homeomorphisms. This theorem is usually stated in term of cell-like maps. For this section, we mainly refer to [Dav86].

Let X be a topological space, we say that a subset of X is **cell-like** if it can be contracted inside any of its neighborhoods. We also say that a map is **cell-like** if the preimage of any point is a cell-like set. In particular, a contractible subset is cell-like.

A similar but more restricting notion is the cellularity. We say that a subset of a manifold is **cellular** if it has a basis of neighborhoods $(U_i)_i$ homeomorphic to balls. We also say that a subset of a manifold is **point-like** if its complement is homeomorphic to the complement of a point.

In a manifold, these notions are linked in the following way.

Proposition 6. *Let A be a compact subset of a manifold, the three following statements verify the following implications: $(1) \Leftrightarrow (2) \Rightarrow (3) \Leftrightarrow (4)$*

1. A is point-like
2. A is cellular
3. A is cell-like
4. A is contractible

Proof. The equivalences $(1) \Leftrightarrow (2)$ and $(3) \Leftrightarrow (4)$ are proved in [Edw80].

The implication $(2) \Rightarrow (3)$ is direct: let U be a neighborhood of A , and let V be neighborhood of A contained in U and homeomorphic to a ball. As V is contractible, it can be contracted to a point within U , so the same is true for A . \square

The union of the Alexander horned sphere with the tame component of its complement or the Fox-Artin arc (Figure 1.21) are examples of cell-like but non-cellular subsets. Indeed, they are contractible (they are homeomorphic to a closed ball and to $[0, 1]$) but not point-like (their complement are not simply connected).



Figure 1.21 – A contractible but non-cellular arc

Cell-like sets are important in the theory of decomposition of surfaces: the shrinking of a family of cell-like sets in a surface always yields the same surface, and the quotient map is approximable by homeomorphisms.

Theorem 7 (Moore [Dav86]). *Let $f : M \rightarrow X$ be a cell-like map defined on a closed surface M . Then X is also a surface, and f is approximable by homeomorphisms.*

This theorem can be generalized to non-compact manifolds by asking f to be a proper map.

Remark that this theorem is specific to dimension 2. For example, let γ be a wild arc in \mathbb{R}^3 whose exterior is not simply connected. The shrinking of this arc is a cell-like maps, but the resulting space is not homeomorphic to \mathbb{R}^3 as it has a point with a non-simply connected complement.

1.4 A tameness criterion

This section is dedicated to the proof of the following tameness criterion.

Proposition 1. *Let Σ be a closed topological submanifold of a closed 3-manifold M . Suppose that its complement $M \setminus \Sigma$ is homeomorphic to the interior of a compact manifold X with boundary. If the inclusion $i : M \setminus \Sigma \rightarrow M$ extends to a continuous map from X to M*

$$\begin{array}{ccc} & X & \\ \uparrow & \searrow & \\ M \setminus \Sigma & \xrightarrow{i} & M \end{array}$$

then Σ is tamely embedded in M .

With the notations and setting of Proposition 1, we denote by f the continuous map from ∂X to M defined by the extension of i . Remark that f is a surjective map from ∂X to Σ . The boundary ∂X has a neighborhood in X homeomorphic to $\partial X \times [0, 1]$ (see [Bro62]).

The subset Σ thus has a closed neighborhood U in M such that $U \setminus \Sigma$ is homeomorphic to $\partial X \times [0, 1[$ via a homeomorphism which extends to a continuous map π from $\partial X \times [0, 1]$ to U .

$$\begin{aligned} \pi : \partial X \times [0, 1[&\longrightarrow U \setminus \Sigma && \text{is a homeomorphism} \\ \pi : \partial X \times \{1\} &\longrightarrow \Sigma && \text{is the continuous map } f \end{aligned}$$

The neighborhood U of Σ is then homeomorphic to the quotient of $\partial X \times [0, 1]$ by π (that is, by gluing the points of $\partial X \times \{1\}$ having the same image by f).

$$U \simeq \partial X \times [0, 1] / \pi$$

The topological submanifold Σ is a finite disjoint union of connected manifolds of possibly different dimensions. We will study these connected components separately. As the components of dimension 0 and 3 are automatically tame, the only cases of interest are the components of dimension 1 and 2.

1.4.1 First case : Σ is a surface

In this subsection, we suppose that the topological submanifold Σ of M is a connected surface.

To show that Σ is tamely embedded, we want to show that it has a tubular neighborhood, as explained in Section 1.2.5. Morally, the map π is not a homeomorphism due to two obstructions. The first is the fact that the map f should be of degree two, as Σ should locally have two sides. The second is that the preimages of f are not necessarily discrete (for example, f can shrink a disk of ∂X to a point).




Figure 1.22 – $\partial X \times [0, 1]$ before and after the quotient by π .

We will deal with the first obstruction by defining a double cover $\tilde{\Sigma}$ of Σ together with a new map \tilde{f} from ∂X to $\tilde{\Sigma}$ giving rise to a new map $\tilde{\pi}$ defining a new quotient space \tilde{U} .




Figure 1.23 – $\partial X \times [0, 1]$ before and after the quotient by $\tilde{\pi}$.

We will deal with the second obstruction by approximating the map \tilde{f} by homeomorphisms.

As we do not give conditions on the orientability of Σ and M , we do not know if Σ cuts U into one or two connected components. However, a local version of the Jordan-Brouwer separation theorem determines this behavior at a local scale. This theorem is a direct consequence of [Lem18, Theorem 2].

Theorem 8 (Local Jordan-Brouwer separation theorem). *Let S be a closed topological $(n-1)$ -submanifold of a n -manifold M . Then every point x of S has a basis of neighborhoods $(V_k^x)_k$ in M such that $V_k^x \setminus S$ has exactly two connected components which both approach x arbitrarily closely.*

We begin by unfolding Σ to a surface $\tilde{\Sigma}$ with underlying set

$$\tilde{\Sigma} = \bigcup_{x \in \Sigma} \varprojlim \pi_0(\mathcal{U}_x \setminus \Sigma).$$

Here, the inverse limit is taken over the neighborhoods \mathcal{U}_x of x in M .

Theorem 8 tells us that each point of Σ has a basis of open neighborhoods $(V_k^x)_k$ in M such that $V_k^x \setminus \Sigma$ has exactly two connected components. As each component of $V_k^x \setminus \Sigma$ approaches x arbitrarily closely for every k , the inclusion map from $\pi_0(V_k^x \setminus \Sigma)$ to $\pi_0(V_{k-1}^x \setminus \Sigma)$ is a bijection. Then, for each point $x \in \Sigma$, $\varprojlim \pi_0(\mathcal{U}_x \setminus \Sigma)$ has two elements.

Let $y \in V_k^x \cap \Sigma$. Each element of $\varprojlim \pi_0(\mathcal{U}_y \setminus \Sigma)$ naturally maps to one of the components

of $V_k^x \setminus \Sigma$. The subset \tilde{V} defined by

$$\tilde{V} = \bigcup_{y \in V_k^x \cap \Sigma} \varprojlim \pi_0(\mathcal{U}_y \setminus \Sigma)$$

separates into two subsets, depending on the component of $V_k^x \setminus \Sigma$ to which its elements map. We endow $\tilde{\Sigma}$ with the topology for which each of these subsets are open neighborhoods, for every x and every k . This makes the projection map $p : \tilde{\Sigma} \rightarrow \Sigma$ a double covering space of Σ .

For a point $y \in \partial X$, the sequence $\pi(y, 1 - \frac{1}{n})$ ends up in only one component of $V_k^{f(y)}$, where π is the map defined in the beginning of Section 1.4. This defines a new continuous map $\tilde{f} : \partial X \mapsto \tilde{\Sigma}$. This allows us to define a new quotient \tilde{U} of $\partial X \times [0, 1]$ by gluing the points of $\partial X \times \{1\}$ having the same image by \tilde{f} instead of f . We denote by $\tilde{\pi}$ the quotient map from $\partial X \times [0, 1]$ to \tilde{U} .

$$\tilde{U} = \partial X \times [0, 1] / \tilde{\pi}$$

The surface $\tilde{\Sigma}$ will be identified with the subset $\tilde{\pi}(\partial X \times \{1\})$. Note that the quotient of \tilde{U} by the covering map $p : \tilde{\Sigma} \rightarrow \Sigma$ is homeomorphic to U .

The map \tilde{f} is not a homeomorphism, but is approximable by homeomorphisms. To prove this, we will show that \tilde{f} is cell-like (see Section 1.3.1).

Lemma 5. *The map \tilde{f} is cell-like.*

To show lemma 5, we begin by proving Lemma 6 and Lemma 7.

Lemma 6. *For any point $x \in \tilde{\Sigma}$, the fiber $\tilde{f}^{-1}(\{x\})$ is connected.*

Proof. From Theorem 8, every point of $\tilde{\Sigma}$ has a basis of closed neighborhoods $(C_k)_k$ in \tilde{U} such that $C_k \setminus \tilde{\Sigma}$ is connected.

Suppose that the fiber $\tilde{\pi}^{-1}(\{x\}) \subset \partial X \times \{1\}$ can be written as a disjoint union of two non-empty compact sets A_1 and A_2 . Choosing a distance d on $\partial X \times [0, 1]$, the closed set Y defined by

$$Y = \{y \in \partial X \times [0, 1] \mid d(y, A_1) = d(y, A_2)\}$$

is disjoint from A_1 and A_2 . The sets $\tilde{\pi}^{-1}(C_k)$ are a basis of closed neighborhoods of $\tilde{\pi}^{-1}(\{x\})$, and the sets $\tilde{\pi}^{-1}(C_k \setminus \tilde{\Sigma})$ are homeomorphic to the sets $C_k \setminus \tilde{\Sigma}$ and are thus connected. The

sets $\tilde{\pi}^{-1}(C_k)$ thus always intersect Y and the intersections $\tilde{\pi}^{-1}(C_k) \cap Y$ are therefore closed and non-empty. Finally, the intersection $\bigcap_k \tilde{\pi}^{-1}(C_k) \cap Y$ is non-empty and contained in $\bigcap_k \tilde{\pi}^{-1}(C_k) = \tilde{\pi}^{-1}(\{x\})$ and in Y , which is impossible since Y is disjoint from A_1 and A_2 . The fiber $\tilde{\pi}^{-1}(\{x\}) = \tilde{f}^{-1}(\{x\})$ is then connected. \square

Lemma 7. *Let $g : A \rightarrow B$ be a continuous surjective map with connected fibers between compact Hausdorff spaces. Let $B' \subset B$ and $A' = g^{-1}(B')$. If B' is connected, then A' is also connected. And if B' is non-separating in B , then A' is also non-separating in A .*

Proof. We consider that B' is connected. Suppose that A' is the disjoint union of two non-empty subsets A_1 and A_2 closed in A' . This means that $\overline{A_1}$, the closure of A_1 in A , does not intersect A_2 and that $\overline{A_2}$, the closure of A_2 in A , does not intersect A_1 . As A and B are Hausdorff and compact, $g(\overline{A_1})$ and $g(\overline{A_2})$ are closed in B . As B' is connected, the sets $g(\overline{A_1}) \cap B'$ and $g(\overline{A_2}) \cap B'$ cannot be disjoint. So, let x be an element of $g(\overline{A_1}) \cap g(\overline{A_2}) \cap B'$. The fiber $g^{-1}(x)$ is connected and then cannot intersect both A_1 and A_2 . We can consider without loss of generality that it is contained in A_1 . However, as x is in $g(\overline{A_2})$, there is an element y of $\overline{A_2}$ such that $g(y) = x$. But $g^{-1}(x)$ is contained in A_1 and we saw that A_1 and $\overline{A_2}$ were disjoint. So this situation is impossible, which means that A' must be connected.

We now consider that B' is a non-separating. This means that $B \setminus B'$ is connected. Using the first part of this lemma, we obtain that $A \setminus A' = g^{-1}(B \setminus B')$ is connected. So A' is non-separating. \square

Now, we can show that \tilde{f} is cell-like.

Proof of Lemma 5. Let x be a point of $\tilde{\Sigma}$. Let $(U_i)_{i \in \mathbb{N}}$ be a decreasing sequence of open disks in $\tilde{\Sigma}$ whose intersection is $\{x\}$. We want to show that the sets $\tilde{f}^{-1}(U_i)$ are also homeomorphic to disks. The set $\tilde{f}^{-1}(\{x\})$ will then be the intersection of a decreasing family of disks, which implies that the map \tilde{f} is cell-like.

First, we know that $\tilde{f}^{-1}(U_i)$ is a connected surface thanks to Lemma 7. This surface has only one end. Indeed, U_i is an increasing union of non-separating compacts set, and so is $\tilde{f}^{-1}(U_i)$ by Lemma 7.

Consider the subset $\pi(f^{-1}(p(U_i)) \times [0, 1])$ of M , where $p : \tilde{\Sigma} \rightarrow \Sigma$ is the cover map. This neighborhood is a manifold since its boundary $\pi(f^{-1}(p(U_i)) \times \{0\})$ has a collared neighborhood $\pi(f^{-1}(p(U_i)) \times [0, \varepsilon])$. Endow this manifold with a smooth structure so that the map π is smooth on this collar. Suppose that $\tilde{f}^{-1}(U_i)$ is not a disk, it is then

homeomorphic to a non-simply connected compact surface with a point removed. We can then find two smooth simple curves γ_1 and γ_2 in $\tilde{f}^{-1}(U_i)$ with a non-zero modulo 2 intersection number. Let us define the smooth curve $\gamma' = \pi(\gamma_1 \times \{\frac{\varepsilon}{3}\})$ and a smooth surface S' in the following way. First, we define a topological surface S by two pieces S_1 and S_2 . The first piece S_1 is $\pi(\gamma_2 \times [0, 1])$. The curve $\pi(\gamma_2 \times \{1\})$ is contained in U_i and can be contracted into a point, as U_i is a disk. This homotopy defines the second piece S_2 . This surface S defined by the two pieces S_1 and S_2 can be approximated by a smooth surface S' without being modified on $\pi(\gamma_2 \times [0, \frac{2\varepsilon}{3}])$. Finally, γ' and S' have an intersection number of 1, but this is not possible since γ' can also be homotoped to U_i and then contracted to a point in the same way we constructed S . So $\tilde{f}^{-1}(U_i)$ is a disk. \square

We can now prove Proposition 1.

Proof of Proposition 1 if Σ is a surface. Moore's theorem (Theorem 7) states that cell-like maps between surfaces are approximable by homeomorphisms. The surface $\tilde{\Sigma}$ is then homeomorphic to ∂X and the map \tilde{f} is thus approximable by homeomorphisms f_n . Choose any homeomorphism h from ∂X to $\tilde{\Sigma}$ and define the following maps.

$$\begin{aligned} h_n : \partial X &\rightarrow \partial X \\ x &\mapsto f_n^{-1} \circ h(x) \end{aligned}$$

$$\begin{aligned} \varphi_n : \partial X &\rightarrow \tilde{U} \\ x &\mapsto \tilde{\pi}(x, 1 - \frac{1}{n}) \end{aligned}$$

The map φ_n sends points of ∂X close to $\tilde{\Sigma}$, and the map h_n is a reparametrization of ∂X . As the limits of φ_n and of f_n are both \tilde{f} , the limit of the sequence $\varphi_n \circ h_n = \varphi_n \circ f_n^{-1} \circ h$ is the map h , as shown in Lemma 8 bellow.

From here, one can apply a Bing's approximation theorem [Bin59, Theorem 11.1] to obtain the tameness of Σ , but we will provide a more direct proof.

For any positive real number $t \in [n, n+1]$, we can interpolate the maps φ_n and φ_{n+1} and the maps h_n and h_{n+1} to obtain maps φ_t and h_t with the same properties. To define φ_t , we just replace n by t in the definition, and h_t is defined using local arcwise connectivity

of $\text{Homeo}(\partial X)$ (see Lemma 10 below). We can then define the following homeomorphism.

$$\begin{aligned} \Psi : \partial X \times [0, 1] &\rightarrow \tilde{U} \\ (x, t) &\mapsto \begin{cases} \varphi_{\frac{1}{1-t}} \circ h_{\frac{1}{1-t}} & \text{if } t < 1 \\ h(x) & \text{if } t = 1 \end{cases} \end{aligned}$$

For Ψ to be continuous at $t = 1$, the paths between the maps h_n and h_{n+1} must however be chosen carefully so that their diameters go to zero with n . This is done thanks to Lemma 10 below.

Finally, taking the quotient of \tilde{U} by the covering map $p : \tilde{\Sigma} \rightarrow \Sigma$ yields the total space of a vector bundle over Σ . As this quotient is homeomorphic to U , this proves that Σ has a tubular neighborhood. By proposition 5, Σ is then tamely embedded. \square

The final arguments of this proof are justified by the following results. First, we make use of Lemma 8 to show the convergence of $\varphi_n \circ h_n$. Indeed, we have $\varphi_n \circ h_n = \varphi_n \circ f_n^{-1} \circ h$ and φ_n converges to f .

Lemma 8. *Let f be a map between two compact metric spaces X and Y . Suppose that f is approximable by homeomorphisms, meaning that there is a sequence f_n of homeomorphisms converging uniformly to f . Then the sequence $f \circ f_n^{-1}$ converges uniformly to the identity function of Y .*

Proof. Let $\varepsilon > 0$ and $y \in Y$. There is a $n \in \mathbb{N}$ such that $d(f(x), f_n(x)) < \varepsilon$ for every x of X . Let $z = f_n^{-1}(y)$. We have

$$d(f(z), f_n(z)) < \varepsilon.$$

We then have

$$d(f \circ f_n^{-1}(y), y) < \varepsilon.$$

So $f \circ f_n^{-1}$ converges uniformly to the identity function. \square

Remark 2. *The sequence $f_n^{-1} \circ f$ does not necessarily converge to the identity function of X . Indeed, if f is not injective, $f_n^{-1} \circ f$ cannot converge to an injective map. For example, consider the following function f*

$$\begin{aligned} f : [0, 2] &\rightarrow [0, 1] \\ x &\mapsto \begin{cases} 0 & \text{if } x < 1 \\ x - 1 & \text{if } x \geq 1 \end{cases} \end{aligned}$$

which is approximable by the homeomorphisms f_n .

$$f_n : [0, 2] \rightarrow [0, 1]$$

$$x \mapsto \begin{cases} x/n & \text{if } x < 1 \\ \frac{n-1}{n}(x-2) + 1 & \text{if } x \geq 1 \end{cases}$$

Then the sequence $f_n^{-1} \circ f$ does not converge to the identity function of $[0, 2]$. Indeed to




Figure 1.24 – Graphs of f and f_n .

converges to the following map g .

$$g : [0, 2] \rightarrow [0, 2]$$

$$x \mapsto \begin{cases} 0 & \text{if } x \leq 1 \\ x & \text{if } x > 1 \end{cases}$$

Secondly, we prove that σ can be approached by appropriate surfaces. To obtain our final product structure, we need to fill the gaps between these surfaces. The flow φ already gives us the supports of the desired intermediate surfaces, what remains to be done is to connect the homeomorphisms h_n . We will show that the surface homeomorphism groups are locally path connected and that we have some control on the length of these paths.

The first Lemma is known as the Alexander trick [Ale23].

Lemma 9 (Alexander trick [Ale23]). *Let $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a homeomorphism such that $d(x, g(x))$ is bounded. There is a path γ in $\text{Homeo}(\mathbb{R}^n)$ from the identity to g . This path is continuous for the topology of the uniform convergence and we have the equality*

$$\max_{x \in \mathbb{R}^n} d(x, \gamma_t(x)) = t \max_{x \in \mathbb{R}^n} d(x, g(x)).$$

Proof. Consider the map $\gamma_t(x) = t g(\frac{x}{t})$. The equality stated in the lemma is clear from this formula and shows that

$$\lim_{t \rightarrow 0} d(x, \gamma_t(x)) = 0$$

for all x .

One can think of this proof as a "step back" argument : when looking to g from further and further away, the transformation looks more and more like the identity. \square

In particular, we can apply this lemma to any homeomorphism of \mathbb{R}^n that fixes the exterior of a compact set. In our situation, we want a path γ in $\text{Homeo}(\varphi_n(\partial X))$ from $\varphi_n \circ h_n \circ \varphi_n^{-1}$ to $\varphi_n \circ h_{n+1} \circ \varphi_n^{-1}$.

Lemma 10. *For $n \in \mathbb{N}$ large enough, there is a path γ from in $\text{Homeo}(\varphi_n(\partial X))$ from the identity to $g_n = \varphi_n \circ h_{n+1} \circ \varphi_n^{-1}$. This path is continuous for the topology of the uniform convergence and we have the inequality*

$$\max_{x \in \varphi_n(\partial X)} d(x, \gamma_t(x)) \leq K \max_{x \in \varphi_n(\partial X)} d(x, g_n(x)).$$

for some constant K independent of n .

Proof. Let $(A_i)_i$ for $0 \leq i \leq N$ be closed disks covering Σ and $(B_i)_i$ be closed disks such that $A_i \subset \overset{\circ}{B}_i$ and let $\delta = \min_i d(A_i, \partial X \setminus B_i)$. Choose diffeomorphisms ψ_i from $S^1 \times [-1, 1[$ to $B_i \setminus A_i$ and choose $L > 0$ such that the $\psi_i|_{[-\frac{1}{2}, \frac{1}{2}]}$ are L -bilipschitz. Noting $\epsilon_n = \max_{x \in \varphi_n(\partial X)} d(x, g_n(x))$, choose n_0 big enough so that $\epsilon_n < \frac{\delta}{100}$, $\epsilon_n < \frac{\delta}{100L^2}$, $\epsilon_n < \frac{1}{2L}$ and $d(\varphi_n(\partial X), \Sigma) < \frac{\delta}{100}$ for any $n \geq n_0$. From now on, we will consider that n is greater than n_0 .

We begin by showing that the map g_n is the composition of a finite number of homeomorphisms b_i that fix the exterior of some ball. Let $\tilde{B}_i = \varphi_n \circ f^{-1}(B_i)$ and $\tilde{A}_i = \varphi_n \circ f^{-1}(A_i)$. From Lemma 5, these subsets of $\varphi_n(\partial X)$ are disks and we have $d(\tilde{A}_i, \varphi_n(\partial X) \setminus \tilde{B}_i) > 98 \epsilon_n$.

We will write g_n as the composition $b_n \circ \dots \circ b_1$ where each b_i is defined as follow :

- b_i fixes $(\varphi_n(\partial X) \setminus \tilde{B}_i) \cup \psi_i(S^1 \times [-1, -L\epsilon_n])$
- b_i fixes the disks \tilde{A}_j for $j < i$
- b_i equals $g_n \circ b_1^{-1} \circ \dots \circ b_{i-1}^{-1}$ on $\tilde{A}_i \cup \psi_i(S^1 \times [L\epsilon_n, 1])$

These conditions are compatible and the only part where a map b_i still needs to be defined is on $\psi_i(S^1 \times [-L\epsilon_n, L\epsilon_n])$. Let $x \in \psi_i(S^1 \times \{L\epsilon_n\})$, the distance $d(x, \psi_i(S^1 \times \{-L\epsilon_n\}))$ is at least ϵ_n and at most $L^2\epsilon_n$. Choose m points x_j in S^1 for $j \in \mathbb{Z}_m$ at distance at least $2L^3\epsilon_n$ and at most $4L^3\epsilon_n$ from their two neighbors. The cells \tilde{C}_j between $\{x_j\} \times [-L\epsilon_n, L\epsilon_n]$ and $\{x_{j+1}\} \times [-L\epsilon_n, L\epsilon_n]$ in $S^1 \times [-L\epsilon_n, L\epsilon_n]$ produce cells $C_j = \psi_i(\tilde{C}_j)$ of a diameter at most $5L^4\epsilon_n$ in $\varphi_n(\partial X)$. We can then find paths from $g_n \circ b_1^{-1} \circ \dots \circ b_{i-1}^{-1}(\psi_i(x_j, L\epsilon_n))$ to $\psi_i(x_j, -L\epsilon_n)$ of length at most $2L^2\epsilon_n$. These paths produce new cells C'_j of diameter at




Figure 1.25 – Illustration of Lemma 10

most $6L^4\epsilon_n$. When $C_j = C'_j$, define b_i on the cells C_j by the identity, and define b_i on the other cells by any homeomorphism from C_j to C'_j .

For each of these homeomorphisms b_i , Lemma 9 tells us that there is a path b_i^t from the identity to b_i . The composition $\gamma_t = b_n^t \circ \dots \circ b_2^t \circ b_1^t \circ f$ is then a path from f to g verifying the desired inequality for $K = 10L^4N$. \square

1.4.2 Second case : Σ is a circle

In this subsection, we suppose that the subset Σ of M is a circle S^1 . The approach will be very similar to what we did for the first case. We will show that ∂X is the total space of a circle bundle over S^1 (that is, a torus or a Klein bottle) and that the map π from $\partial X \times [0, 1]$ to U can be reparametrized to a map whose restriction to $\partial X \times \{1\}$ is the standard projection of this bundle onto its basis.

In this case, there are no cell-like map from ∂X to some surface, but the general method still works. Most of the lemmas used for the first case will be adapted by hand, and previously used arguments will only be sketched.

Lemma 11. *For any point $x \in \Sigma$, the fiber $f^{-1}(x)$ is connected and non-separating.*

Proof. The proof of Lemma 6 can be reused as it is. The important points being that the connected neighborhoods $(C_k)_k$ still exist and that the sets $C_k \setminus \Sigma$ are still connected, as removing a 1-dimensional topological submanifold from a 3-dimensional connected manifold cannot disconnect it.

Lemma 7 still shows that the fibers are also non-separating in ∂X . \square

Lemma 12. *Let I be an open interval of Σ . The set $f^{-1}(I)$ is homeomorphic to an open cylinder $S^1 \times]0, 1[$.*

Proof. We know that $f^{-1}(I)$ is a connected surface thanks to Lemma 7 and Lemma 11. This surface has two ends. Indeed, I is an increasing union of compact sets I_n separating I into two connected components C_n and C'_n with $C_{n+1} \subset C_n$ and $C'_{n+1} \subset C'_n$. So, by Lemma 7 and Lemma 11, $f^{-1}(I)$ is also the increasing union of the compact sets $f^{-1}(I_n)$ separating $f^{-1}(I)$ into two connected components $f^{-1}(C_n)$ and $f^{-1}(C'_n)$ with $f^{-1}(C_{n+1}) \subset f^{-1}(C_n)$ and $f^{-1}(C'_{n+1}) \subset f^{-1}(C'_n)$.

Suppose that $f^{-1}(I)$ is not an open cylinder, it is then homeomorphic to a non-simply connected compact surface with two points removed. We can then find two smooth curves γ_1 and γ_2 in $f^{-1}(I)$ with a non-zero modulo 2 intersection number. The argument used in Lemma 5 then work in the same way. \square

Lemma 13. *The surface ∂X is homeomorphic to the total space of a circle bundle over Σ .*

Proof. Write Σ as the union of two open intervals I_1 and I_2 . These two intervals intersect in two other intervals I_3 and I_4 . Let S_3 and S_4 be two non-null-homotopic circles in the cylinders $f^{-1}(I_3)$ and $f^{-1}(I_4)$. The circles S_3 and S_4 are also non-null-homotopic in $f^{-1}(I_1)$, so the compact set between S_3 and S_4 in $f^{-1}(I_1)$ is then homeomorphic to a closed cylinder $S^1 \times [0, 1]$. The same is true for $f^{-1}(I_2)$, so ∂X is the union of two closed cylinder glued along their boundaries. So ∂X is homeomorphic to the total space of a circle bundle over Σ . \square

We can now conclude the proof of Proposition 1.

Proof of Proposition 1 if Σ is a circle. Let h be a map from ∂X to Σ so that the triple $(\partial X, \Sigma, h)$ is a circle bundle. For $n \in \mathbb{Z}$, let f_n be a smooth $\frac{1}{n}$ -approximation of f . By Sard's theorem, there is a homeomorphism j between \mathbb{R}/\mathbb{Z} and Σ such that each $q \in \mathbb{Q}/\mathbb{Z} \subset j^{-1}(\Sigma)$ is a regular value for every f_n . Each preimage $f_n^{-1}\left(j\left(\frac{i}{2^n}\right)\right)$ for $0 \leq i < 2^n$ then contains a non-null-homotopic smooth circle $S_{i,n}$. We then define a homeomorphism h_n from ∂X to itself that sends each circle $h^{-1}\left(j\left(\frac{i}{2^n}\right)\right)$ on $S_{i,n}$. And we also define the following map.

$$\begin{aligned} \varphi_n : \partial X &\rightarrow M \\ x &\mapsto \pi\left(x, 1 - \frac{1}{n}\right) \end{aligned}$$

On a circle $h^{-1}\left(j\left(\frac{i}{2^n}\right)\right)$, the limit of the sequence $(\varphi_n \circ h_n)_{n \in \mathbb{N}}$ is $j\left(\frac{i}{2^n}\right)$. So this sequence converges to the map h .

For any positive real number t , we can interpolate the maps φ_n and h_n to obtain maps φ_t and h_t as in section 1.4.1. We can then define the following map.

$$\begin{aligned} \Psi : \partial X \times [0, 1] &\rightarrow M \\ (x, t) &\mapsto \begin{cases} \varphi_t \circ h_t(x) & \text{if } t < 1 \\ h(x) & \text{if } t = 1 \end{cases} \end{aligned}$$

By identifying the points having the same image by Ψ , we obtain the total space of a disk bundle over Σ that maps homeomorphically to a neighborhood of Σ and whose core is mapped on Σ . By proposition 5, Σ is tamely embedded. \square

Now that we have introduced every notion of tameness useful for our problem and that we have proved our tameness criterion, we can start discussing the Smith conjecture.

THE SMITH CONJECTURE

2.1 Introduction

Let X be a topological space. A **self-homeomorphism** of X is a homeomorphism $\sigma : X \rightarrow X$ from X onto itself. We say that two self-homeomorphisms σ and σ' of X are **conjugate** if there is there is a self-homeomorphism f of X such that

$$\sigma' = f^{-1} \circ \sigma \circ f.$$

A self-homeomorphism σ of X is **of finite order** (or **periodic**) if there is $n \in \mathbb{N}_{>0}$ such that f^n (the composition of f with itself n times) is the identity. If f is of finite order, the smaller number n with this property is called the **order** of f .

The understanding and the classification of such finite order self-homeomorphisms is an important topic in topology. If X is the euclidean space \mathbb{R}^n and if we only consider linear transformations, this classification is well known. For example, every finite order linear automorphism of \mathbb{R}^3 is conjugate either to the identity, to a reflection, to a rotation of a rational number of turn around some axis, or to the antipodal map $[x \mapsto -x]$. We can then define a large variety of finite order self-homeomorphism of \mathbb{R}^3 by conjugating these automorphisms by any self-homeomorphism of \mathbb{R}^3 . As the linear automorphisms seem to be the only examples that arise naturally, the following question then follows

Question 2. *Is there a finite order self-homeomorphism of \mathbb{R}^3 which is not conjugate to a linear automorphism ?*

Although it appears to be difficult to find such a map, it turns out that the answer to this question is positive. This answer may however change if one consider other spaces than \mathbb{R}^3 and consider only maps with some regularity.

Our study will mostly focus on compact 3-manifolds, and the fundamental example to keep in mind will be the 3-sphere S^3 . The non-compactness of the euclidean space \mathbb{R}^n

should not eclipse the previous discussion. Indeed, the euclidean spaces and spheres behave in the same way in many aspects.

First, as S^{n-1} is the subset $\{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$, transformations of \mathbb{R}^n that fix S^{n-1} yield transformations of S^{n-1} . The only linear transformations of \mathbb{R}^n mapping S^{n-1} onto itself are the isometries of \mathbb{R}^n and are well known. The group of every such transformation, the orthogonal group O_n , then acts on S^{n-1} .

Secondly, the n -sphere S^n is homeomorphic to the one-point compactification of \mathbb{R}^n , that is \mathbb{R}^n with an additional point for which a family of neighborhoods are the complement of compact subsets of \mathbb{R}^n . In other words, \mathbb{R}^n is homeomorphic to S^n with a point removed. Any self-homeomorphism of \mathbb{R}^n is proper and can then be extended to self-homeomorphism of S^n by fixing its additional point. On the contrary, a self-homeomorphism of S^n with a fixed point can yield a self-homeomorphism of \mathbb{R}^n by removing one its fixed point. Question 2 can then be reworded as follows

Question 3. *Is there a finite order self-homeomorphism of S^3 which is not conjugate to an element of O_4 ?*

The only real difference between Question 2 and Question 3 is the fact that O_4 contains a free involution (the restriction of the antipodal map $[x \mapsto -x]$ in \mathbb{R}^4) which does not exist in \mathbb{R}^3 .

2.1.1 Smith theory

One of the earliest results in the study of finite-order homeomorphisms of S^3 is the determination of the topology of the fixed set by P. A. Smith. He showed ([Smi39, 7.3 Theorem 4]) that the fixed set of a finite-order homeomorphism of S^3 is homeomorphic to a lower dimensional sphere (that is, S^3 itself, a 2-sphere, a circle, a pair of points or no fixed points at all). This result can be generalized to any 3-manifold as follows.

Theorem 9. *Let $\sigma : M \rightarrow M$ be a finite order homeomorphism of a topological 3-manifold M . The fixed set M^σ is a disjoint union of open pieces $(\Sigma_i)_i$ where Σ_i is a topological manifold.*

A proof of this result can be found, for example, in [Par20, Theorem 4.5]). As in John Pardon's paper, Smith theory is usually stated for prime orders, but Theorem 9 is still valid for any finite order. Indeed, suppose that we have Theorem 9 for prime orders, and let σ be a homeomorphism of a topological 3-manifold of order n , and let $p_1 \cdot \dots \cdot p_k$ be

a decomposition of n into prime factors. The map $\sigma^{\frac{n}{p_1}}$ is of order p_1 , so the result of Theorem 9 applies to the fixed set $M^{\sigma^{\frac{n}{p_1}}}$. The map $\sigma^{\frac{n}{p_1 \cdot p_2}}$ is of order p_2 on $M^{\sigma^{\frac{n}{p_1}}}$ and we have the inclusion $M^{\sigma^{\frac{n}{p_1 \cdot p_2}}} \subset M^{\sigma^{\frac{n}{p_1}}}$, and as Theorem 9 is also true for dimensions 2 and 1, the result still applies. We can continue this argument until it applies to σ .

2.1.2 The Bing Involution

Smith theory describes the intrinsic topology of the fixed set M^σ , but it does not describe how the latter is embedded in M . For example, Smith theory does not answer the following questions:

- If $M^\sigma \simeq S^1$, can it be a non-trivial knot in M ?
- Can M^σ be wildly embedded ?

If one only requires continuity on the map σ , the answer is affirmative for those two questions. We will describe an example of an involution of S^3 whose fixed set is a wildly embedded 2-sphere.

Consider the Alexander horned sphere presented in Section 1.1.3. We recall that it is a 2-sphere wildly embedded in S^3 . Note W the wild component of its complement, accordingly to Figure 2.1. The horned sphere is then its boundary ∂W and its union with its exterior component is the closure \overline{W} , called the Alexander solid horned sphere.





Figure 2.1 – the Alexander horned sphere

Bing proved in 1952 [Bin52] that the double of the Alexander solid horned sphere is homeomorphic to S^3 . This double is defined as the disjoint union of two Alexander solid horned spheres \overline{W} glued along their boundary ∂W by identifying each point of ∂W in the first copy with the corresponding point in the second copy: $\overline{W} \bigsqcup_{\partial W} \overline{W}$.

Exchanging those two copies of \overline{W} yields an involution i of S^3 called the Bing involution. The fixed set of this involution is the 2-sphere ∂W , which is wildly embedded as its

complement is the disjoint union of two copies of W , which are not simply connected.

$$\text{The Bing involution: } \overline{W} \sqcup \overline{W} \simeq S^3$$


2.2 The Lipschitz Smith Conjecture

2.2.1 Early results

In 1939, P. A. Smith proved [Smi39, 7.3 Theorem 4] that the fixed set M^σ of a finite-order and orientation-preserving homeomorphism σ of the 3-sphere S^3 was empty or a circle S^1 . He then asked [Eil49, Problem 36] if this circle could be knotted. This question is known as the Smith conjecture.

Conjecture 1 (Smith). *The fixed set of a finite order map from S^3 into itself cannot be a nontrivial knot.*

Question 3 is a generalization of this question, and its answer may be different if we only consider maps of a given regularity. Since the work on geometrization by Thurston and Perelman, we know that the Smith conjecture is true for smooth maps [BLP05]. On the other hand, if we only ask σ to be continuous, we saw in Section 2.1.2 that the Bing involution was a counterexample. Montgomery and Zippin [MZ54] also modified Bing's example to obtain an orientation-preserving involution with a wild circle as a fixed set.

One can then wonder what happens for maps with more regularity but which are not differentiable. For example, Jani Onninen and Pekka Pankka showed in 2019 [OP19] that there also exists wild involutions in the Sobolev class $W^{1,p}$.

Our study will focus on the Lipschitz continuity. The answer to Question 3 for Lipschitz map seems to be considered positive, but the ideas to prove this fact only existed in unpublished work until the present thesis. In [Ham08], D. H. Hamilton announced that quasi-conformal reflections are tame, but the proof seems to remain unpublished. In 2013, Michael Freedman asked in [Fre13, Conjecture 3.21] if the Bing involution could be conjugate to a Lipschitz homeomorphism. Jani Onninen and Pekka Pankka reiterate this question in 2019 [OP19].

Conjecture 2 (Freedman). *Every bilipschitz action of a finite group on a closed 3-manifold is conjugate to a smooth action.*

In this thesis, we give a partial answer to Freedman's question, proving that for $\varepsilon > 0$ small enough, wild finite-order maps can not be $(1 + \varepsilon)$ -bilipschitz. More precisely, we will show the following theorem.

Theorem 1. *For $\varepsilon = \frac{1}{4000}$, any action of a finite cyclic group by $(1 + \varepsilon)$ -bilipschitz homeomorphisms on a closed 3-manifold is conjugate to a smooth action.*

This theorem started with an idea of Pekka Pankka and Juan Souto. They obtained the following result in unpublished work, we will sketch their argument.

Theorem 10 (Pekka Pankka and Juan Souto). *Let σ be an involution of $M = \mathbb{R}^3$ with a fixed set M^σ homeomorphic to \mathbb{R}^2 . If the complement of the fixed set is not a tame manifold, then there exists an ε such that σ is not conjugate to a $(1 + \varepsilon)$ -bilipschitz involution.*

Sketch of the proof. We want to show that, if σ is conjugate to a $(1 + \frac{1}{n})$ -bilipschitz involution σ_n of \mathbb{R}^3 for n arbitrarily large, then the map

$$\varphi_n : x \mapsto x + \frac{1}{4}(\sigma_n(x) - x)$$

would take every point of \mathbb{R}^3 closer to M^{σ_n} for some n . It would follow from the Tucker Theorem 6 that the complement of M^{σ_n} is a tame manifold. Indeed, if such a map exists, it is possible to find a compact submanifold with boundary N arbitrarily large in the complement of M^{σ_n} such that $\lim_k \varphi_n^k(N) = M \setminus M^{\sigma_n}$. The fundamental group of the complement of a compact subset in $M \setminus M^{\sigma_n}$ would then be isomorphic to the fundamental group of the complement of a compact subset in N , which is finitely generated.

Suppose by contradiction that for every $n \in \mathbb{N}$ we could find a $(1 + \frac{1}{n})$ -bilipschitz involution φ_n of \mathbb{R}^3 fixing 0 and conjugate to σ such that

$$d(\varphi_n(x_n), M^{\sigma_n}) \geq 0.99 d(x_n, M^{\sigma_n})$$

for some x_n . With rotations and homotheties, we can suppose that $x_n = (1, 0, 0) =: x_\infty$ for every n . By the Arzelà-Ascoli theorem, the sequence σ_n converges to some map σ_∞ (with a corresponding map φ_∞). This map verifies

$$d(\varphi_\infty(x_\infty), M^{\sigma_\infty}) \geq 0.99 d(x_\infty, M^{\sigma_\infty}).$$

But it is also an isometry as it is 1-bilipschitz, which is a contradiction. Indeed, for an isometry of \mathbb{R}^3 , we have $d(\varphi_\infty(x_\infty), M^{\sigma_\infty}) = 0.5 d(x_\infty, M^{\sigma_\infty})$. So, there exist a number n

so that, for every $(1 + \frac{1}{n})$ -bilipschitz involution φ of \mathbb{R}^3 fixing 0 and conjugate to σ , we have

$$d(\varphi_\infty(x_\infty), M^{\sigma_\infty}) < 0.99 d(x_\infty, M^{\sigma_\infty}).$$

for every x in \mathbb{R}^3 .

So there exists a n such that the map φ_n takes every point of \mathbb{R}^3 closer to M^{σ_n} .

□

Theorem 1 is a more general version of Theorem 10: it gives an explicit value for ε and it applies to a larger set of spaces and homeomorphisms.

Michael Freedman and Michael Starbird recently published another result. They proved that the Bing involution could not be conjugate to a Lipschitz map.

Theorem 11 (Freedman, Starbird [FS22]). *Let σ be a map conjugate to the Bing involution and let $\delta(\varepsilon)$ be its modulus of continuity at scale ε . Then we have the following inequality*

$$\delta(\varepsilon)^{-1} \geq 1.167\sqrt{1/\varepsilon}$$

for a sequence of ε 's converging to zero. In particular, σ cannot be Lipschitz.

Here, the modulus of continuity $\delta(\varepsilon)$ is defined as the highest value so that $|x - y| \leq \delta(\varepsilon)$ implies $|\sigma(x) - \sigma(y)| \leq \varepsilon$.

This theorem is stronger than Theorem 1 because it has no bound on the Lipschitz constant, but it only applies to the Bing involution.

2.2.2 Some results on Lipschitz vector fields

Let \mathbf{v} be a vector field on a compact manifold M . If \mathbf{v} is Lipschitz continuous, the Picard-Lindelöf theorem tells us that one can solve the autonomous ordinary differential equation

$$\gamma'(t) = \mathbf{v}(\gamma(t)) , \gamma(0) = x_0 \tag{2.1}$$

which is a C^1 curve following the vector field \mathbf{v} and starting at a point x_0 . One can then define the flow $\varphi : M \times \mathbb{R} \rightarrow M$ of \mathbf{v} by letting $\varphi(x_0, t)$ be the solution of (2.1) at time t with the initial condition x_0 .

The results about such a flow are often stated for C^1 vector fields. In fact, the Lipschitz continuity is sufficient to obtain the continuity of φ . For example, we state the following

lemma from [CB08, Lemma 2], which gives the Lipschitz continuous dependence on initial conditions.

Lemma 14 (Continuous dependence on initial conditions [CB08]). *Let f be a Lipschitz vector field with constant K defined on an open subset of a Banach space. Let σ_x and σ_y be solutions to f for initial conditions x and y with interval $I \ni 0$ contained in their common domains. Then*

$$\|\sigma_x(t) - \sigma_y(t)\| \leq \|x - y\| e^{K|t|}$$

for all $t \in I$.

An other important result is the **flow-box Theorem**. If \mathbf{v} is a C^1 vector field and if x is a point of \mathbb{R}^d such that $\mathbf{v}(x) \neq 0$, then there is a neighborhood U of x and a neighborhood of 0 such that there exist a diffeomorphism from U to V sending \mathbf{v} to the standard vector field $(1, 0, 0, \dots, 0)$. In particular, The flow of \mathbf{v} locally acts by diffeomorphisms.

If \mathbf{v} is only Lipschitz continuous, we can only hope to obtain a conjugacy to a standard vector field via a bilipschitz homeomorphism. However, a bilipschitz homeomorphism cannot transfer vector fields, they can however transfer the action of the flow. We say that two vector fields \mathbf{v}_1 and \mathbf{v}_2 defined on two open sets U_1 and U_2 are **topologically conjugate** near $x_1 \in U_1$ and $x_2 \in U_2$ if there exist two neighborhoods W_1 and W_2 of x_1 and x_2 and a homeomorphism h from W_1 to W_2 with $h(x_1) = x_2$ and such that

$$h(\varphi_1(x, t)) = \varphi_2(h(x), t)$$

for every $x \in W_1$ and $|t|$ sufficiently small, where φ_1 and φ_2 are the local flows of \mathbf{v}_1 and \mathbf{v}_2 . The following theorem is then proved in [CB08, Theorem 4].

Theorem 12 (Lipschitz Flow box Theorem[CB08]). *Let X be a Banach space with open subset U . Let $f : U \rightarrow X$ be a Lipschitz vector field. Let $z \in X$ be nonzero and let $g : X \rightarrow X$ be the constant vector field $g(x) = z$. Then for any point $x_1 \in U$ with $f(x_1) \neq 0$, f and g are locally topologically conjugate near x_1 and $x_2 := 0$.*

The homeomorphism which gives the conjugacy is bilipschitz.

2.2.3 The Smith Conjecture for C^1 maps

Theorem 1 also applies to C^1 homeomorphisms, without any bound on its derivatives. Indeed, we will show that every C^1 homeomorphism is conjugate to a $(1 + \varepsilon)$ -bilipschitz homeomorphism for every $\varepsilon > 0$.

Let ρ be the riemannian metric on M and suppose that, for any $g \in G$, the map $x \mapsto gx$ is C^1 . Any such map lifts to a continuous map from TM to TM . We can define the pullback $g^*\rho$ of ρ by any such g :

$$g^*\rho(u, v) = \rho(gu, gv)$$

This metric $g^*\rho$ is only continuous.

We can define the sum of these pullbacks:

$$\tilde{\rho} = \sum_{g \in G} g^*\rho$$

This metric verifies the property $\tilde{\rho}(hu, hv) = \tilde{\rho}(u, v)$ for every $h \in G$.

$$\begin{aligned} \tilde{\rho}(hu, hv) &= \sum_{g \in G} g^*\rho(hu, hv) \\ &= \sum_{g \in G} \rho(ghu, ghv) \\ &= \sum_{g' \in G} \rho(g'u, g'v) \quad \text{by setting } g' = gh \\ &= \tilde{\rho}(u, v) \end{aligned}$$

So G acts by isometries on M for the metric $\tilde{\rho}$. This metric is only continuous but we can approximate it by a smooth metric ρ_ε such that

$$\frac{1}{1+\varepsilon} \tilde{\rho}(u, v) \leq \rho_\varepsilon(u, v) \leq (1+\varepsilon) \tilde{\rho}(u, v).$$

So we have

$$\left(\frac{1}{1+\varepsilon} \right)^2 \rho_\varepsilon(u, v) \leq \rho_\varepsilon(gu, gv) \leq (1+\varepsilon)^2 \rho_\varepsilon(u, v).$$

This means that G acts by $(1+\varepsilon)^2$ -lipschitz homeomorphisms on M for the metric ρ_ε . So, by Theorem 1, the action of G is smoothable.

2.2.4 Reducing Theorem 1 to two propositions

Theorem 1 will be proved by studying the tameness of the fixed set of the action. If the fixed sets of the elements of such an action are tamely embedded, it is known that this

action is smoothable (see Theorem 13). Theorem 1 can be reduced to two propositions.

The first proposition, proved in Section 1.4, is a tameness criterion allowing us to obtain, in some specific cases, the tameness of an embedding from the topology of its complement.

Proposition 1. *Let Σ be a closed topological submanifold of a closed 3-manifold M . Suppose that its complement $M \setminus \Sigma$ is homeomorphic to the interior of a compact manifold X with boundary. If the inclusion $i : M \setminus \Sigma \rightarrow M$ extends to a continuous map from X to M*

$$\begin{array}{ccc} & X & \\ \uparrow & \searrow & \\ M \setminus \Sigma & \xrightarrow{i} & M \end{array}$$

then Σ is tamely embedded in M .

We recall that a topological submanifold is a subset which is a topological manifold for the induced topology. For example, the Alexander horned sphere is a topological submanifold of \mathbb{R}^3 . Note that the extension of the map i just needs to be continuous, without any injectivity hypothesis. This proposition will be proved by showing that the resulting map from ∂X to Σ is approximable by coverings.

In the last part, we show how this criterion can be applied to finite order $(1 + \varepsilon)$ -bilipschitz homeomorphisms by proving Proposition 2.

Proposition 2. *For $\varepsilon = \frac{1}{4000}$ and for every $(1 + \varepsilon)$ -bilipschitz action of a finite group G on a compact Riemannian manifold M , the fixed set M^G satisfies the conditions of Proposition 1.*

This proposition is proved by defining a Lipschitz vector field on M and by showing that the flow of this vector field converges to the fixed set. This allows us to define a product structure on a neighborhood of the fixed set which extends continuously to the latter. The Lipschitz continuity of the action is crucial to define this vector field, and the bound on ε is used to show the convergence of its flow.

Note that Proposition 2 works in any dimension with any finite group. The conditions of Theorem 1 are imposed by Theorem 13 and Proposition 1.

Remark 3. *As Proposition 2 only uses local arguments, the compactness hypothesis in its statement could be removed with some work. However, we cannot get rid of compactness needed in Theorem 1 as it is needed for Theorem 13 stated in the next part.*

The fixed set of a smooth action on a 3-manifold is a smooth submanifold. In particular, this fixed set is tamely embedded. It turns out that the tameness of the fixed set for every power of a finite order self-homeomorphism of a 3-manifold is a sufficient condition for this map to be smoothable.

Theorem 13. *[KL88, Corollary 2.3] A topological action of a finite cyclic group G on a closed 3-manifold M is smoothable if and only if, for every subgroup H of G , the fixed set M^H is tame.*

Note that the tameness of the global fixed set M^G is not sufficient. For example, consider the disjoint union of three 3-spheres $S^3 \sqcup S^3 \sqcup S^3$. Let r be a circular permutation exchanging these three copies, and let s be a wild involution on each 3-sphere such that r and s commute. The group generated by r and s is cyclic of order six but its action has an empty (and thus tame) global fixed set, and this action is not smoothable.




Figure 2.2 – A wild action on a compact manifold with an empty global fixed set.

Theorem 1 can be reduced to Proposition 1 and Proposition 2 stated earlier.

Proof of Theorem 1. Let G be a finite cyclic group acting by $(1 + \varepsilon)$ -bilipschitz homeomorphisms on a closed 3-manifold. Each subgroup H of G is cyclic and acts by $(1 + \varepsilon)$ -bilipschitz homeomorphisms. By Proposition 2 and Proposition 1, the fixed sets of these subgroups are tame. Theorem 13 then applies, implying that the action of G is smoothable. \square

Section 1.4 is dedicated to the proof of Proposition 1, and Section 2.3 to Proposition 2.

2.3 Proof of Proposition 2

In this section, we show that we can use the tameness criterion developed in the previous section on finite order $(1 + \varepsilon)$ -bilipschitz mappings. More precisely, we show the following.

Proposition 2. *For $\varepsilon = \frac{1}{4000}$ and for every $(1 + \varepsilon)$ -bilipschitz action of a finite group G on a compact Riemannian manifold M , the fixed set M^G satisfies the conditions of Proposition 1.*

Remark that this proposition works for any dimension and any finite group.

To obtain Proposition 2, we will define a continuous flow φ near the fixed set M^G such that

$$\lim_{t \rightarrow \infty} \varphi_t(x) \in M^G$$

for any point x close enough to M^G , and such that this convergence is uniform in x . We will show that this flow allows us to define a compactification of the complement of M^G which extends continuously (but not homeomorphically) to M^G .

This flow φ will be constructed as the flow of a Lipschitz vector field \mathbf{v} . By Lemma 14, Lipschitz continuity on the parameters of an ODE is enough to obtain a continuous dependency on the initial conditions and to produce a continuous flow.

2.3.1 A vector field near the fixed set

Let $\varepsilon = \frac{1}{4000}$ and let G and M be as in Proposition 2. If we were given a point x near the fixed set M^G , then a natural place where we could look for a fixed point would be near the center of mass $B(x)$ of the orbit of x . We will build a vector field \mathbf{v} pointing towards this center of mass. In what follows, we will freely use some classic notions and facts of Riemannian geometry, such as the convexity radius and the center of mass. We refer to [Ber03] to learn more about these topics.

Let r be the convexity radius of M . The center of mass of a finite set S of points in a ball of radius r is the unique minimum of the function

$$z \mapsto \sum_{s \in S} d(z, s)^2$$

Equivalently, it is the unique zero of the vector field

$$z \mapsto \sum_{s \in S} \exp_z^{-1}(s).$$

Note that, for any point x at a distance $\frac{r}{1+\varepsilon}$ from a fixed point, the orbit G_x of x is included in a ball of radius r around the latter. This means that G_x has a well-defined center of mass for any such x .

Let us proceed with some definitions.

— If $d(x, M^G) < \frac{r}{1+\varepsilon}$, let $B(x)$ be the center of mass of the orbit of x .

As explained in Section 2.3.3, B has a Lipschitz dependency on x . This means that $\exp_x^{-1}(B(x))$ is Lipschitz in x .

— Let \mathbf{v} be a Lipschitz vector field on M with $\mathbf{v}(x) = \exp_x^{-1}(B(x))$ wherever $d(x, M^G) < \frac{r}{1+\varepsilon}$.

— Let φ be the flow of \mathbf{v} .

By Lemma 14, the flow φ is well-defined and acts by homeomorphisms.

To prove Proposition 2, we will make use of the following lemma.

Lemma 15. *There are positive constants $\tau > 0$ and $k < 1$ and an open neighborhood V of M^G such that we have the inequality*

$$\|\mathbf{v}(\varphi_\tau(x))\| \leq k \|\mathbf{v}(x)\|$$

for all x in V .

We will also need the following result, asserting that, as long as x is close to M^G , then the point $B(x)$ is almost fixed by G .

Lemma 16. *There is a constant $R > 0$ depending only on M and an open neighborhood V of M^G such that, for $g_0 \in G$, we have*

$$d(B(x), g_0 B(x)) \leq R d(x, B(x)).$$

for all x in V .

Now, we show how Lemma 15 and Lemma 16 can be used to prove Proposition 2.

Proof of Proposition 2. As $\|\mathbf{v}(x)\|$ converges to 0 as x approaches M^G , there is a neighborhood U of M^G such that every flow line of \mathbf{v} starting in it will stay in V and then converge. Such limits are fixed points: the vector field \mathbf{v} vanishes at the limits of the flow lines and, from the expression of \mathbf{v} , these limits are thus points fixed by the map B . Using Lemma 16, we see that a point fixed by B is also fixed by the action of G .

To show that the complement $M \setminus M^G$ of the fixed set is homeomorphic to the interior of a compact manifold with boundary, we will define a codimension 1 topological submanifold Z intersecting every flow line only once. This will show that the end of $M \setminus M^G$ is homeomorphic to $Z \times [0, 1[$, proving that $M \setminus M^G$ is homeomorphic to the interior of a compact manifold with boundary.

We begin by defining a map l that measures the length of the flow line from a point x to the fixed set :

$$l(x) = \int_0^\infty \|\mathbf{v}(\varphi_t(x))\| dt$$

We claim that this quantity is finite and depends continuously on x . Indeed, from Lemma 15, when x is sufficiently close to the fixed set, we have the two following inequalities :

$$\begin{aligned} \|\mathbf{v}(\varphi_{t+\tau}(x))\| &\leq k \|\mathbf{v}(\varphi_t(x))\| \\ \|\mathbf{v}(\varphi_t(x))\| &\leq 1 \end{aligned}$$

From which we obtain :

$$\|\mathbf{v}(\varphi_t(x))\| \leq \|\mathbf{v}(x)\| k^{\frac{t}{\tau}-1} \leq k^{\frac{t}{\tau}-1} \quad (\star)$$

We see that the map $t \mapsto \|\mathbf{v}(\varphi_t(x))\|$ is bounded by an integrable map independent of x . Thanks to the dominated convergence theorem, this uniform integrable bound shows that $l(x)$ is finite and that the map l inherits the continuity of its integrand.

Choose $b > 0$ so that every point at a distance at most b of M^G is in U and define the subset Z of M as follows.

$$Z = l^{-1}(b)$$

When t increases, the length $l(\varphi_t(x))$ decreases, the set Z thus intersects each flow line

at most once. On the other hand, every flow line must intersect Z at some point, because it cannot converge to a fixed point both in positive and negative times.

This set is a closed topological submanifold. Indeed, as the map l is continuous and as M^G is compact, the set Z is also compact. Moreover, from the Lipschitz flow-box Theorem (Theorem 12), we see that every point of Z has a neighborhood homeomorphic to \mathbb{R}^{n-1} .

The map $\varphi : Z \times [0, 1[\rightarrow M : (x, t) \mapsto \varphi_{\frac{t}{1-t}}(x)$ is a homeomorphism onto its image which describes the topology of the end of $M \setminus M^G$ and which can be compactified by adding a boundary $Z \times \{1\}$. This is done by taking the limit of the flow.

$$\begin{aligned} Z \times \{1\} &\longrightarrow M^G \\ (x, 1) &\longmapsto \lim_{t \rightarrow \infty} \varphi_t(x) \end{aligned}$$

The uniform bound (\star) also shows that the limit $\lim_{t \rightarrow \infty} \varphi_t$ is uniform, making this map continuous.

This proves that the inclusion of $M \setminus M^G$ in M extends to a continuous map from X to M . □

Our goal is now to show Lemma 15 and Lemma 16.

2.3.2 Proof of Lemma 15 and Lemma 16 in a flat setting

To make the computations easier, we will begin by working in a flat geometry. Namely, for a fixed point $p \in M^G$, we will consider that p has a sufficiently large neighborhood isometric to an open subset of \mathbb{R}^n . As we will only work locally, we will be able to reduce to this case. We will discuss this question in section 2.3.3.

We will show the following.

Lemma 17. *If M is flat around a point p , there are positive constants $\tau > 0$ and $k' < 1$ and an open neighborhood V_p of p such that we have the inequality*

$$\|\mathbf{v}(\varphi_\tau(x))\| \leq k' \|\mathbf{v}(x)\|$$

for all x in V_p .

We suppose that, for x sufficiently close to p , there is a flat ball V_p centered at p and of radius $2d(x, p)$. For some coordinate system on this ball, the map B and the vector field \mathbf{v}

have simple expressions:

$$\begin{aligned} B(x) &= \frac{1}{|G|} \sum_{g \in G} gx \\ \mathbf{v}(x) &= B(x) - x \end{aligned}$$

The maps B has a $(1 + \varepsilon)$ -Lipschitz dependency on x and the vector field \mathbf{v} has a $(2 + \varepsilon)$ -Lipschitz dependency on x . Some classical computations about centers of mass lead to the following equation, for x and y in \mathbb{R}^n .

$$\frac{1}{|G|} \sum_{g \in G} d(y, gx)^2 = d(y, B(x))^2 + \frac{1}{|G|} \sum_{g \in G} d(B(x), gx)^2 \quad (*)$$

We begin by showing the flat version of Lemma 16.

Lemma 18. *If M is flat around a point p , there is an open neighborhood V_p of p such that, for $g_0 \in G$, we have*

$$d(B(x), g_0 B(x)) \leq R d(x, B(x)).$$

for all x in V , where $R = \frac{1}{40}$.

Proof. From the Lipschitz continuity of the action of g_0 , we have

$$\frac{1}{|G|} \sum_{g \in G} d(g_0 B(x), gx)^2 \leq (1 + \varepsilon)^2 \frac{1}{|G|} \sum_{g \in G} d(B(x), gx)^2$$

Evaluating $(*)$ at $y = g_0 B(x)$, we obtain

$$d(g_0 B(x), B(x))^2 = \frac{1}{|G|} \sum_{g \in G} d(g_0 B(x), gx)^2 - \frac{1}{|G|} \sum_{g \in G} d(B(x), gx)^2$$

Together with the previous inequality, this leads to

$$d(g_0 B(x), B(x)) \leq \sqrt{2\varepsilon + \varepsilon^2} \max_{g \in G} d(B(x), gx)$$

We would like to obtain a inequality depending only on ε and on the distance $d(x, B(x))$.

To do this, note that we have the inequality :

$$\begin{aligned}
 \max_{g \in G} d(B(x), gx) &\leq (1 + \varepsilon) \max_{g \in G} d(gB(x), x) \\
 &\leq (1 + \varepsilon) \max_{g \in G} \left(d(gB(x), B(x)) + d(B(x), x) \right) \\
 &\leq (1 + \varepsilon) \left(\sqrt{2\varepsilon + \varepsilon^2} \max_{g \in G} d(B(x), gx) + d(x, B(x)) \right)
 \end{aligned}$$

So

$$\max_{g \in G} d(B(x), gx) \leq \frac{1 + \varepsilon}{1 - (1 + \varepsilon)\sqrt{2\varepsilon + \varepsilon^2}} d(x, B(x))$$

And finally

$$d(g_0 B(x), B(x)) \leq \frac{(1 + \varepsilon)\sqrt{2\varepsilon + \varepsilon^2}}{1 - (1 + \varepsilon)\sqrt{2\varepsilon + \varepsilon^2}} d(x, B(x))$$

Taking $\varepsilon = \frac{1}{4000}$, the quantity $\frac{(1 + \varepsilon)\sqrt{2\varepsilon + \varepsilon^2}}{1 - (1 + \varepsilon)\sqrt{2\varepsilon + \varepsilon^2}}$ is smaller than $\frac{1}{40}$. □

We can now prove Lemma 17.

Proof of Lemma 17. Let x be a point of V_p , $\tau = \frac{1}{5}$, $k' = \frac{9999}{10000}$ and $\delta = d(x, B(x))$.

Step 1 : We prove that $d(\varphi_t(x), x) \leq \frac{\delta}{3}$ for every $t \leq \tau$.

If $d(\varphi_t(x), x) > \frac{\delta}{3}$ for some $t \leq \tau$, let

$$t_0 = \min \left\{ t \leq \tau \mid d(\varphi_t(x), x) \geq \frac{\delta}{3} \right\}.$$

For $t \leq t_0$, we have $d(\varphi_t(x), x) \leq \frac{\delta}{3}$, so $d(B(\varphi_t(x)), B(x)) \leq \frac{\delta}{3}(1 + \varepsilon)$ as B is $(1 + \varepsilon)$ -Lipschitz. As $\mathbf{v}(\varphi_t(x)) = B(\varphi_t(x)) - \varphi_t(x)$, $\varphi_t(x)$ is contained in the convex hull of $\{x\} \cup \mathcal{B}(B(x), \frac{\delta}{3}(1 + \varepsilon))$ for every $t \leq t_0$ (see Figure 2.3, where \mathbf{v} points inward on the boundary of the circular sector).

The distance between $\varphi_t(x)$ and $B(\varphi_t(x))$ is then smaller than the diameter of this convex hull, which is $\delta + \frac{\delta}{3}(1 + \varepsilon) = \frac{\delta(4 + \varepsilon)}{3}$. So $\|\mathbf{v}(\varphi_t(x))\| \leq \frac{\delta(4 + \varepsilon)}{3}$, then $d(x, \varphi_{t_0}(x)) \leq t_0 \frac{\delta(4 + \varepsilon)}{3} \leq \frac{1}{5} \frac{\delta(4 + \varepsilon)}{3} < \frac{\delta}{3}$. This is in contradiction with the definition of t_0 , so $d(\varphi_t(x), x)$ cannot exceed $\frac{\delta}{3}$ for $t \leq \tau$.




Figure 2.3 – Setting of Step 1

Step 2 : We prove that $d(\varphi_\tau(x), B(x)) < \frac{19}{20} \delta$.

As $d(\varphi_t(x), x) \leq \frac{\delta}{3}$ for every $t \leq \tau$, we obtain $d(B(\varphi_t(x)), B(x)) \leq \frac{\delta}{3}(1 + \varepsilon)$ and then $\|\mathbf{v}(\varphi_t(x))\| \geq \delta - \frac{\delta}{3} - \frac{\delta}{3}(1 + \varepsilon) = \frac{\delta}{3}(1 - \varepsilon)$.

The distance between $\varphi_\tau(x)$ and $B(x)$ will be the greatest if $\mathbf{v}(\varphi_t(x))$ makes an angle $\alpha = \arcsin(\frac{1 + \varepsilon}{3})$ with the vector $B(x) - x$ for every $t \leq \tau$ (see Figure 2.3). At time τ , if $\mathbf{v}(\varphi_t(x))$ made an angle α with $B(x) - x$ for every $t \leq \tau$, $\varphi_\tau(x)$ would be at distance at least $\tau \frac{\delta}{3}(1 - \varepsilon)$ from x , and at distance at most $\sqrt{(\tau \frac{\delta}{3}(1 - \varepsilon))^2 + \delta^2 - 2\tau \frac{\delta^2}{3}(1 - \varepsilon) \cos(\alpha)} \leq \frac{19}{20} \delta$ from $B(x)$.

Step 3 : We prove that $d(\varphi_\tau(x), B(\varphi_\tau(x))) \leq \frac{9999}{10000} \delta$.

Let $Q = \frac{d(\varphi_\tau(x), B(x))}{\delta}$. By Step 2, we know that $Q < \frac{19}{20}$. The points in the orbit of $\varphi_\tau(x)$ verify the following inequality, for every $g \in G$.

$$d(g\varphi_\tau(x), gB(x)) \leq (1 + \varepsilon) d(\varphi_\tau(x), B(x))$$

Which, by Lemma 18 and Step 2, leads to :

$$d(g\varphi_\tau(x), B(x)) \leq (1 + \varepsilon + \frac{R}{Q}) d(\varphi_\tau(x), B(x)) \quad (1)$$

These points also verify :

$$d(g\varphi_\tau(x), gB(\varphi_\tau(x))) \geq \frac{1}{1 + \varepsilon} d(\varphi_\tau(x), B(\varphi_\tau(x))).$$

Which, also by Lemma 18, leads to :

$$d\left(g\varphi_\tau(x), B(\varphi_\tau(x))\right) \geq \left(\frac{1}{1+\varepsilon} - R\right) d\left(\varphi_\tau(x), B(\varphi_\tau(x))\right). \quad (2)$$




Figure 2.4 – Limit case of Step 3

As $B(\varphi_\tau(x))$ is the center of mass of the orbit $G\varphi_\tau(x)$, it must be contained in the convex hull of this orbit. As the possible positions of the points of this orbit are restricted by inequalities (1) and (2), $B(\varphi_\tau(x))$ must be in the convex hull of the region defined by these two inequalities. As shown in Figure 2.4 which presents the limit case, this condition is satisfied when the triangle shown on the figure is obtuse at $B(\varphi_\tau(x))$. This condition is given by the inequality :

$$\begin{aligned} \left(1 + \varepsilon + \frac{R}{Q}\right)^2 d(\varphi_\tau(x), B(x))^2 &\geq \left(\frac{1}{1+\varepsilon} - R\right)^2 d(\varphi_\tau(x), B(\varphi_\tau(x)))^2 \\ &\quad + d(B(\varphi_\tau(x)), B(x))^2. \end{aligned}$$

Note that this equation is still satisfied when the spheres defined by the two inequalities do not intersect.

Taking $\varepsilon = \frac{1}{4000}$, $R = \frac{1}{40}$ and $Q < \frac{19}{20}$, the positions of $B(\varphi_\tau(x))$ satisfying this inequality are in the interior of an ellipse contained in the ball of center $\varphi_\tau(x)$ and of radius $\frac{9999}{10000} \delta$. So we necessarily have $d(\varphi_\tau(x), B(\varphi_\tau(x))) \leq \frac{9999}{10000} \delta$.

So $\|\mathbf{v}(\varphi_\tau(x))\| \leq \frac{9999}{10000} \delta = k' \|\mathbf{v}(x)\|$ for $k' = \frac{9999}{10000}$. □

2.3.3 Reduction to the flat case

In section 2.3.2, we worked locally on a flat neighborhood of a fixed point. In general, the manifold M cannot be flattened in a neighborhood of M^G , but as every manifold is locally almost flat, Lemma 17 will allow us to produce the inequalities needed for Lemma 15 and Lemma 18 will allow us to prove Lemma 16.

Proof of Lemma 15. Let p be a fixed point of G . On a small neighborhood of p , the manifold M is almost flat. The goal of this proof is to compare the flow φ obtained with the metric of M and the flow φ^E obtained by the flat metric of $T_p M$. Every object obtained in the flat metric of $T_p M$ will be noted with the exponent E .

Let U and U' be the balls of center p and of radii δ and $(1 + \varepsilon)\delta$ for a δ small enough so that U' is contained in V .

First, notice that $B(x)$ has a Lipschitz dependency on x and on the Riemannian metric ρ . This fact can be proved using a Lipschitz version of the implicit function theorem on the map

$$\begin{aligned} \Phi : U \times U' \times \mathcal{M} &\longrightarrow \mathbb{R}^n \\ (x, y, \rho) &\longmapsto \rho\left(\sum_{g \in G} \exp_y^{-1}(gx), E_i\right) \end{aligned}$$

where \mathcal{M} is the space of all metrics on U' for the uniform operator norm (according to the starting metric of M) and $(E_i)_i$ is a basis of sections of TM . The map Φ is then K_1 -bilipschitz in x and y for some constant K_1 . As the norm of $\sum_{g \in G} \exp_y^{-1}(gx)$ is smaller than $K_2\delta$ for some constant K_2 , the map Φ is $K_3\delta$ -Lipschitz in ρ for some constant K_3 . The map B is then K_1^2 -Lipschitz in x and $K_1K_3\delta$ -Lipschitz in ρ .

The Riemannian metric ρ on U' is always at a distance $K_4\delta$ from a Euclidean metric (namely, the metric of $T_p M$ induced by the exponential map), for some constant K_4 . The ratio $\frac{d^E(x, y)}{d(x, y)}$ is then between $1 - K_4\delta$ and $1 + K_4\delta$ for any x and y in U' .

With the above, the map B_E is at a distance at most $K_1K_3K_4\delta^2$ from B . The distance between \mathbf{v} and \mathbf{v}^E is then itself bounded by $K_5\delta^2$ for some constant K_5 . And the point $\varphi_\tau(x)$ is at a distance at most $K_6\delta^2$ from $\varphi_\tau^E(x)$ for some K_6 .

Lemma 17 gives us

$$d^E\left(\varphi_\tau^E(x), B^E(\varphi_\tau^E(x))\right) \leq k'd^E(x, B^E(x))$$

from which we obtain from what precedes

$$d\left(\varphi_\tau(x), B(\varphi_\tau(x))\right) \leq k'\delta + K\delta^2 = (k' + K\delta)\delta$$

for some constant K depending only on M . Consequently, for any manifold M , it is always possible to work as locally as we want (i.e. to choose δ small enough) so that we have

$$d\left(\varphi_\tau(x), B(\varphi_\tau(x))\right) \leq \frac{k' + 1}{2}\delta$$

which can be rewritten as follows.

$$\|\mathbf{v}(\varphi_\tau(x))\| \leq \frac{k' + 1}{2}\|\mathbf{v}(x)\|$$

This proves Lemma 15. □

We can also prove Lemma 16 using the preceding computations.

Proof of Lemma 16. From lemma 18, we obtain

$$d^E(B^E(x), g_0 B^E(x)) \leq \frac{1}{40} d^E(x, B^E(x))$$

for every x in V . With the preceding method, we obtain

$$d(B(x), g_0 B(x)) \leq R d(x, B(x))$$

for some constant $R > 0$ depending only on M . □

PART II

Manifold learning

This chapter is a reproduction of the article "*Effective estimation of the dimension of a manifold from random samples*" written with Juan Souto.

EFFECTIVE ESTIMATION OF THE DIMENSION OF A MANIFOLD FROM RANDOM SAMPLES

3.1 Introduction

The *manifold hypothesis* asserts that naturally occurring data sets $X \subset \mathbb{R}^s$ behave as if they had been sampled from a low-dimensional submanifold M . There is then a large literature on *manifold learning*, that is on understanding how properties of the underlying manifold M can be "learned" from the data set X —see for example [BNS06; FMN16; NSW08b; NM10; LV07; Bre+18b]—, but it is maybe fair to say that estimating the dimension $\dim(M)$ of the manifold, that is figuring out the *intrinsic dimension* of the data set, seems to be of particular interest. This is maybe so because the dimension is one of the most basic quantities associated to a manifold, but maybe also because of the importance of understanding the applicability to the data set X of different dimension reduction schemes. In any case, there are very numerous estimators for the intrinsic dimension of a data set. We refer to the surveys [Cam03; CS16; Bre+18b] for a brief discussion of many of those estimators and to the references therein for details—see also Chapter 3 in the monograph [LV07]. Anyways, all these dimension estimators are based on the idea that manifolds look locally like their tangent spaces, that is like euclidean space. For example, non-linear PCA aims at finding the dimension of tangent spaces. ANOVA aims at computing the average angle between vectors in a tangent space, exploiting the fact that for euclidean space itself one can read the dimension from the expected value. The estimator (3.1) below, as well as closely related algorithms due to Takens [Tak83; Tak85], Theiler [The90] or Grassberger-Procaccia [GP04], exploits the fact that the *correlation dimension* of the manifold is just its dimension. There are other such estimators where the correlation dimension is replaced by the *Box counting dimension* or the *Kolmogorov*

capacity. This list of estimators is far from being all inclusive.

All those estimators work very well—in the sense that they yield with high probability the desired result—as long as we apply them with common sense and under suitable conditions, that is if

- 1) we work at a scale at which the manifold M resembles euclidean space, and
- 2) the data set X is large enough and the points therein have been obtained by uniform sampling.

However, amazingle little is known when it comes to quantifying 1) or 2). Indeed, there are classical estimates due to Grassberger [Gra86], Procaccia [Pro88], and Eckmann-Ruelle [ER92] giving absolute lower bounds, in terms of the dimension, for the number of measurements needed for the results to be reliable: Grassberger argues that there are absolute lower bounds at all, Procaccia argues that if the intrinsic dimension is \dim then one might need a data set of at least 10^{\dim} measurements, while Eckmann-Ruelle suggest a lower bound of the form $C \cdot 10^{\frac{1}{2} \dim}$ for some large but undetermined C . One can summarize these results as asserting that, when the dimension grows, the minimal cardinality of a data set allowing to compute the dimension grows exponentially—this is indeed already the case (see [Wei14]) when one just wants to distinguish between spheres of consecutive dimensions.

The absolute lower bounds for the data size that we just mention do not help however with the following less philosophical question: *I suspect, or hypothesize, that my data set is sampled out of a manifold with this or that properties. To what extent can I trust the intrinsic dimension estimation given by this or that estimator?* The only result we know along those lines is due to Niyogi-Smale-Weinberger [NSW08b], at least as long as one wants to allow for variable curvature manifolds. In [NSW08b] the authors give namely an algorithm to compute the homology of a closed submanifold $M \subset \mathbb{R}^s$ out of a set X sampled from M , and they estimated how large does the data size have to be so that their algorithm has success in at least 90% of the cases—observe that knowing the homology we also know the dimension of the manifold. A problem is that the needed data size is astronomical. For example, the estimate in [NSW08b] for the number of points needed to be sampled to compute with 90% probability of success the homology of the 4-dimensional Clifford torus $\mathbb{T}^4 = S^1 \times S^1 \times S^1 \times S^1$ is 24.967.788 points, that is about 25 million points.

One might contend that the Niyosi-Smale-Weinberger algorithm computes something much more sophisticated than the dimension, that they are basically learning the whole manifold, and that being able to do that is an overkill if what one wants to do is to

estimate its dimension. We agree with that point of view. We will consider the estimator

$$\dim_{\text{Corr}(\epsilon_1, \epsilon_2)}(X) = \text{Round} \left(\frac{\log |PX(\epsilon_1)| - \log |PX(\epsilon_2)|}{\log(\epsilon_1) - \log(\epsilon_2)} \right) \quad (3.1)$$

where

$$PX(\epsilon) = \{\{x, y\} \subset X \text{ with } 0 < |x - y| \leq \epsilon\}, \quad (3.2)$$

and our goal will be to estimate how large does $X \subset M$ need to be for the estimator (3.1) to have a 90% success rate. It is definitively much smaller. For example, applying (3.1) to randomly sampled data sets in \mathbb{T}^4 we have a 90% rate of success, as long as we sample 18.262 points, that is more than 1300 times less than before.

Evidently, the value of $\dim_{\text{Corr}(\epsilon_1, \epsilon_2)}(X)$ does not only depend on X but also on the chosen scales. In some sense the goal of this paper is to decide how to choose ϵ_1 and ϵ_2 in such a way that one does not need to sample too many points to be 90% sure that (3.1) returns the correct value. This is the kind of results that we will prove:

Theorem 14. *For $d = 1, \dots, 10$ let ϵ_1 and ϵ_2 be scales as in the table below. Also, given a closed d -dimensional manifold $M \subset \mathbb{R}^s$ with reach $\tau(M) \geq 1$ let n be also as in the following table:*

d	ϵ_1	ϵ_2	n
1	1.5	0.19	$9 + 21 \cdot \text{vol}(M)^{\frac{1}{2}}$
2	0.78	0.2	$94 + 58 \cdot \text{vol}(M)^{\frac{1}{2}}$
3	0.63	0.23	$635 + 146 \cdot \text{vol}(M)^{\frac{1}{2}}$
4	0.54	0.23	$2786 + 392 \cdot \text{vol}(M)^{\frac{1}{2}}$
5	0.46	0.22	$7013 + 1119 \cdot \text{vol}(M)^{\frac{1}{2}}$
6	0.4	0.21	$13221 + 3366 \cdot \text{vol}(M)^{\frac{1}{2}}$
7	0.36	0.21	$25138 + 10644 \cdot \text{vol}(M)^{\frac{1}{2}}$
8	0.33	0.2	$50033 + 34890 \cdot \text{vol}(M)^{\frac{1}{2}}$
9	0.31	0.19	$63876 + 119533 \cdot \text{vol}(M)^{\frac{1}{2}}$
10	0.29	0.18	$139412 + 425554 \cdot \text{vol}(M)^{\frac{1}{2}}$

Then, if we sample independently and according to the riemannian volume form a subset $X \subset M$ consisting of at least n points, then we have

$$\dim_{\text{Corr}(\epsilon_1, \epsilon_2)}(X) = d$$

with at least 90% probability.

Here the *reach* $\tau(M)$ (see Definition 1), or in the terminology of [NSW08b] the *condition number*, is taken as a measure for the local regularity of the submanifold M . It is evident that in order to have specific bounds for the needed data size, we do need to have some a priori control on the local geometry: otherwise we could have a 1-dimensional submanifold so interwoven that it looks as a d -dimensional manifold for some $d \geq 2$. Taking the reach as a measure to quantify to which extent does a submanifold $M \subset \mathbb{R}^s$ resemble euclidean space seems to be actually pretty common in the field of manifold learning [Aam+19; NSW08b; FMN16; BLW19].

After discussing briefly the correlation dimension and the estimator (3.1) in Section 3.2 we discuss some aspects of the geometry of the reach in Section 3.3. We will mostly care about the volume of the thick diagonal

$$DM(\epsilon) = \{(x, y) \in M \times M \text{ with } |x - y| \leq \epsilon\}$$

For example, in Theorem 16 we use the Bishop-Gromov theorem and the properties of CAT(1)-spaces to give upper and lower bounds for the ratio $\frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))}$ for sufficiently small $\epsilon_1 > \epsilon_2$ positive:

Theorem 15. *Suppose that $M \subset \mathbb{R}^s$ is a d -dimensional ($d \geq 1$) closed submanifold with reach $\tau(M) \geq 1$. Then we have*

$$\frac{\frac{\epsilon_1}{2}}{\arcsin(\frac{\epsilon_2}{2})} \left(\frac{\sin(\epsilon_1)}{\sin(2 \arcsin \frac{\epsilon_2}{2})} \right)^{d-1} \leq \frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))} \leq \frac{\int_0^{\sqrt{2} \cdot 2 \cdot \arcsin(\frac{\epsilon_1}{2})} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot \epsilon_2} \sinh^{d-1}(t) dt}$$

for any two $1 > \epsilon_1 > \epsilon_2 > 0$.

In Section 3.4 we come then to core of the present pamphlet. The reason why we care about the volume of the thick diagonal is that for $X \subset M$ we have

$$\frac{\log |PX(\epsilon_1)| - \log |PX(\epsilon_2)|}{\log(\epsilon_1) - \log(\epsilon_2)} \sim \frac{\log(\text{vol}(DM(\epsilon_1))) - \log(\text{vol}(DM(\epsilon_2)))}{\log(\epsilon_1) - \log(\epsilon_2)}$$

with large probability, at least if X has been obtained by independently sampling a large number of points according to the Riemannian measure on M . Basically the goal of Section 3.4, or maybe even the goal of this paper, is to find $\epsilon_1, \epsilon_2 > 0$ so that a relatively small set X is such that with high probability the left side lies in the interval $(d - \frac{1}{2}, d + \frac{1}{2})$. We

prove Theorem 14 above in Section 3.4. The scales in the table in the theorem are obtained numerically—the process is also described in Section 3.4, and a computer implementation can be found in [Gri22].

Heuristical bounds

It is more or less evident that the bounds given by Theorem 14 are far from being sharp. In Section 3.5 we add a heuristical discussion of what could be, in practice, more realistic bounds. The starting point is that the number of sampled points should not be what determines how reliable is the obtained result, but rather the number $|PX(\epsilon_1)|$ of pairs at our larger scale $\epsilon_1 > \epsilon_2 > 0$. Arguing as if

1. at our scales all balls in M were euclidean, and
2. the distances between the two points in pairs as in (3.2) were independent,

we get that, with N as in Table 3.1, it would suffice to have a data set $X \subset M$ with at least N pairs as in (3.2) for (3.1) to give the correct answer in about 90% (resp. 70%) of the cases.

d	ϵ_1	ϵ_2	N for 90%	N for 70%
1	1.5	0.19	30	10
2	0.78	0.2	122	40
3	0.63	0.23	249	111
4	0.54	0.23	516	238
5	0.46	0.22	878	360
6	0.4	0.21	1329	554
7	0.36	0.21	1719	698
8	0.33	0.2	2481	1070
9	0.31	0.19	3900	1604
10	0.29	0.18	5849	2414

Table 3.1 – Heuristic bound N for how large should the cardinality of $PX(\epsilon_1)$ at least be to have 90% (resp. 70%) rate of success when using (3.1).

Note that the assumptions (1) and (2) are not that outlandish, at least if ϵ_1 is small and if the number of pairs is small when compared to the cardinality of X . In any case, for what it is worth, numerical simulations (see Table 3.5) seem to support the values given in Table 3.1. In particular, our numerical simulations also indicate that if $|PX(\epsilon_1)|$ is less or equal than the value in the right column in Table 3.1, then we should count with about a 30% failure rate when we use (3.1).

Remark 4. *As a bigger value of $|PX(\varepsilon_1)|$ results in better performance, it is a natural idea to try to maximize it. This can be done in two ways. First, we can try to obtain more data to increase the number n , but as this is not always possible in concrete situations, the second solution is the increase the scale ε_1 . The downside of this second method is that considering bigger balls increases the effect of the curvature. One must carefully balance those two effects.*

Thinking on the reliability of (3.1) in terms of the number of pairs has the huge advantage that one does not need to have any a priori knowledge of the volume of the manifold. It is however not clear how different are the heuristic bound from Table 3.1 and the formal bound given in Theorem 14. To be able to compare both bounds we note that, always under the assumption that all ε_1 balls in our d -manifold M are euclidean, then if we sample a set $X \subset M$ with n points then we expect to have

$$|PX(\varepsilon_1)| = \frac{n(n-1)}{2} \cdot \frac{\text{vol}(B^{\mathbb{R}^d}(\varepsilon_1))}{\text{vol}(M)}$$

pairs of points. Using this relation we get for example that, using the heuristic bound, it would suffice to sample 1958 points from the torus \mathbb{T}^4 for (3.1) to be 90% of the time correct. See Table 3.4 for more on the comparison between the heuristic bound and the bound in Theorem 14 for the number of points that suffice to have 90% success rate when using (3.1). This is, once again, supported by numerical experiments—see Table 3.6.

Note now that thinking of the applicability of (3.1) in terms of the cardinality of $PX(\varepsilon_1)$ leads to an implementation of (3.1) which could be applied to data sets X sampled from a manifold M whose reach we ignore. The basic idea is the following:

If we want to check if M has dimension d then we take ε_1 minimal so that $|PX(\varepsilon_1)|$ is as large as Table 3.1 asks for in dimension d , then we choose ε_2 so that the ratio $\frac{\varepsilon_1}{\varepsilon_2}$ is as in Table 3.1), and then compute $\dim_{\text{Corr}(\varepsilon_1, \varepsilon_2)}(X)$.

The preceeding discussion implies that for sufficiently rich synthetic data sets this algorithm has a reliability of about 90%. For the sake of completeness we decided to test it also on data sets each consisting in 200 grayscale pictures of the 3D-model *Suzanne* (see Figure 3.1) obtained by randomly choosing 1, 2 and 3 Euler angles—the pictures are 64 by 64 pixels large and can thus be represented as points in \mathbb{R}^{4096} . See Table 3.7 for the obtained results.




Figure 3.1 – 3 sample pictures from the data sets.

Comparison with other estimators

In Section 3.6 we compare briefly the performance of (3.1) with that of a few other dimension estimators: reading the correlation dimension via a log-log plot, ANOVA, and local PCA. The basic observation is that, at least if the volume of M is greater or equal to that of the d -dimensional Clifford torus and if we take sets of cardinality close to the heuristic bound proposed either in Table 3.1 or in Table 3.4, then (3.1) seems to perform better than ANOVA and local PCA, at least with the specific implementations we proposed. It is however harder to compare any of these estimators with the incarnation of (3.1) using a log-log plot: it is namely unclear how to meaningfully quantify what does it mean for the non-constant slope of a curve to look constant. It thus only seems to make sense to test the log-log technique numerically, but it is also unclear how to formally do this: while one might well print diagrams for 100 randomly chosen random subsets of a given manifold, it is not clear how can one formally decide what any given diagram is suggesting to us. Still, the limited experiments we conducted lead to the conclusion that the log-log plot method is pretty reliable for relatively small data sets.

Acknowledgements

The second author would like to thank Kaie Kubjas for getting him interested in this topic.

3.2 The estimator

In this section we recall what is the correlation dimension and how it can be approximated to obtain dimension estimators, but first we introduce some notation that we will

use through out the paper. First, distances in Euclidean space will be denoted either by $d_{\mathbb{R}^n}(x, y)$ or by $|x - y|$. If x and y are points in a Riemannian manifold M then $d_M(x, y)$ is their distance with respect to the Riemannian metric. The ball in M of radius R centered at some point $x \in M$ is denoted by $B^M(x, R)$, although sometimes, when there is no risk of confusion, we drop the subscript writing thus simply $B(x, R)$. Again, if there is no risk of confusion with the dimension, balls in the ambient euclidean space \mathbb{R}^s will be denoted by blackboard bold, that is $\mathbb{B}(x, R) = B^{\mathbb{R}^s}(x, R)$. Finally, the intersection of balls in the ambient space \mathbb{R}^s with the manifold M will be denoted by

$$\mathbb{B}^M(x, R) = M \cap \mathbb{B}(x, R)$$

With this notation in place, we turn our attention to the correlation dimension and the estimator (3.1).

As we mentioned earlier, there are plenty of dimension estimators (see [Cam03; CS16] and [LV07, Chapter 3]). In the language of [CS16], (3.1) is a *fractal-based estimator*. Fractal-based because what one is aiming at, is computing a dimension which makes sense for fractal objects—in this case the correlation dimension. Recall that the *correlation integral* at scale ϵ of a Borel measure μ on \mathbb{R}^s is the integral

$$C(\mu, \epsilon) = \int \mu(\mathbb{B}(x, \epsilon)) d\mu(x)$$

The *upper correlation dimension* and *lower correlation dimension* are then defined as

$$D^+(\mu) = \limsup_{\epsilon \rightarrow 0} \frac{\log(C(\mu, \epsilon))}{\log(\epsilon)} \text{ and } D^-(\mu) = \liminf_{\epsilon \rightarrow 0} \frac{\log(C(\mu, \epsilon))}{\log(\epsilon)}.$$

When both of them agree, then one refers to

$$D(\mu) = D^+(\mu) = D^-(\mu)$$

as the *correlation dimension* of μ . There are numerous situations of interest in dynamics [Sim98] in which the correlation dimension of a measure exists. It is also trivial that it exists if μ is a smooth measure whose support is a submanifold of euclidean space. Since this is the case in which we will find ourselves, we state this fact as a lemma:

Lemma 19. *If μ is a smooth finite measure of full support of a d -dimensional submanifold $M \subset \mathbb{R}^s$, then the correlation dimension exists and we have $D(\mu) = d$. \square*

In [GP04] Grassberger and Procaccia note that if $\epsilon > 0$ is small enough and if we have a set $X = (x_1, \dots, x_n)$ consisting of n points (n very large) sampled independently with respect to the measure μ then $D(\mu)$ is approximated by

$$\dim_{GP}(X, \epsilon) = \frac{\log \left(\frac{|DX(\epsilon)|}{n(n-1)} \right)}{\log(\epsilon)} \quad (3.3)$$

where

$$DX(\epsilon) = \{(x_i, x_j) \in X \times X \text{ with } i \neq j \text{ and } |x_i - x_j| \leq \epsilon\} \quad (3.4)$$

is the set of (ordered) pairs of points in X within ϵ of each other. Grassberger and Procaccia propose (3.3) as an estimator for the intrinsic dimension of the data set X —variants are discussed by Takens [Tak83; Tak85] and Theiler [The90]. We will however focus here on the version from [LV07], or more precisely the quantity

$$\dim_{\text{Corr}(\epsilon_1, \epsilon_2)}(X) = \text{Round} \left(\frac{\log |DX(\epsilon_1)| - \log |DX(\epsilon_2)|}{\log \epsilon_1 - \log \epsilon_2} \right) \quad (3.5)$$

Note that this expression is nothing other than (3.1) above.

A problem when using (3.5) as an estimator, or for that matter when we use (3.3), is that we have to pick up the appropriate scales. In practical implementations, this is often by-passed by taking many possible scales $\epsilon_1 < \epsilon_2 < \dots < \epsilon_k$, computing $|DX(\epsilon_i)|$ for each one of those scales, plotting the result in a log-log-plot and choosing what looks to us as the slope at some region where the slope looks constant. In practice, at least when tested on synthetic data, the log-log plot method is surprisingly effective, but it is very unclear how one can get out of that a formal statistical test, or how can one evaluate how much confidence can one have on the obtained number. Theorem 14 from the introduction gives such bounds, at least if the manifold we are working with is sufficiently regular, in the sense of having *reach* $\tau(M) \geq 1$. We discuss a few aspects of the geometry of the reach in the next section.

3.3 Some Geometry

In this section we recall a few facts about the geometry of reach-1 submanifolds M of euclidean space. Combining these facts with standard arguments from Riemannian geometry we give upper and lower bounds for the volume of the ϵ fat diagonal in $M \times M$.

3.3.1 Reach-1 manifolds

We start recalling the definition of the *reach* of a closed subset of euclidean space:

Definition 1. *The reach $\tau(S)$ of a closed subset $S \subset \mathbb{R}^s$ is the supremum of those $T \geq 0$ with the property that for every $x \in \mathbb{R}^s$ with $d_{\mathbb{R}^s}(x, S) \leq T$ there is a unique point in S closest to x .*

Since the reach was introduced by Federer in [Fed59], it has proved to be a useful notion. Indeed, it follows from the very definition that positive reach sets, that is sets S with $\tau(S) > 0$, have some neighborhood $\mathcal{N}(S)$ on which the closest point projection $\pi : \mathcal{N}(S) \rightarrow S$ is well-defined—the existence of such a projection is enough to show that sets of positive reach share many regularity properties with convex sets.

Here we will be working from the very beginning with very regular objects, namely closed smooth submanifolds $M \subset \mathbb{R}^s$ of euclidean space. In this setting the reach, or rather a lower bound for the reach, helps to quantify how distorted is the inner geometry of M with respect to that of the ambient euclidean space. We summarize what we will need in the following proposition:

Proposition 7. *Let $M \subset \mathbb{R}^s$ be a closed submanifold with reach $\tau(M) \geq 1$. Then we have:*

1. *$d_M(x, y) \leq 2 \arcsin\left(\frac{|x-y|}{2}\right)$ for any two $x, y \in M$ with $|x - y| < 2$.*
2. *The set $\mathbb{B}^M(x, r) = M \cap \mathbb{B}(x, r)$ is geodesically convex for any $x \in \mathbb{R}^s$ and any $r < 1$.*
3. *We have $\angle(\gamma'(0), \gamma'(\ell)) \leq d_M(\gamma(0), \gamma(\ell))$ for every geodesic $\gamma : [0, \ell] \rightarrow M$ parametrized by arc length.*
4. *The manifold M has sectional curvature pinched by $-2 \leq \kappa_M \leq 1$.*

See, in that order Lemma 3, Corollary 1 and Lemma 5 in [BLW19] for the first three claims of Proposition 7. See then Proposition A.1 (iii) in [Aam+19] for the final claim. In any case, we refer to [Aam+19] and [BLW19], and to the references therein for general facts about submanifolds $M \subset \mathbb{R}^s$ of positive reach.

3.3.2 Injectivity radius

Armed with Proposition 7 we can now derive a lower bound for the injectivity radius of those submanifolds $M \subset \mathbb{R}^s$ with reach $\tau(M) \geq 1$. Recall that the *injectivity radius* of a geodesically complete Riemannian manifold M at a point $x \in M$ is defined as

$$\text{inj}(M, x) = \sup\{t > 0 \mid \exp_x : T_x M \rightarrow M \text{ is injective on } B^{T_x M}(0, t)\}$$

where $B^{T_x M}(0, t) = \{v \in T_x M \text{ with } |v| \leq r\}$ and where \exp_x is the Riemannian exponential map. The *injectivity radius* of M itself is then defined to be

$$\text{inj}(M) = \inf_{x \in M} \text{inj}(M, x).$$

Recall also that the *systole* $\text{syst}(M)$ of a closed manifold M is the length of the shortest non-trivial closed geodesic. The importance of the systole now is that, together with an upper bound $\kappa_M \leq \kappa$ for the sectional curvature, the systole yields very a usable bound [Berger] for the injectivity radius:

$$\text{inj}(M) \geq \min \left\{ \frac{\pi}{\sqrt{\kappa}}, \frac{1}{2} \text{syst}(M) \right\}.$$

Altogether, see [Berger] for basic facts and definitions from Riemannian geometry.

Suppose now that $M \subset \mathbb{R}^s$ is a closed submanifold of euclidean space with $\tau(M) \geq 1$, and note that the assumption that M is closed implies that it is metrically complete and thus geodesically complete by the Hopf-Rinow theorem. On the other hand if $\gamma : S^1 \rightarrow M \subset \mathbb{R}^s$ is any smooth curve and if t is such that $\langle \gamma'(0), \gamma(t) \rangle$ is maximal, then $\langle \gamma'(0), \gamma'(t) \rangle = 0$. We thus get from (3) in Proposition 7 that every closed geodesic reaches at least distance $\frac{\pi}{2}$, and hence that

$$\text{syst}(M) \geq \pi.$$

Now, this fact together with the upper bound for the sectional curvature from (4) in Proposition 7 and with the lower bound for the injectivity radius yields the following:

Corollary 1. *If $M \subset \mathbb{R}^s$ is a closed submanifold of euclidean space with reach $\tau(M) \geq 1$, then we have $\text{inj}(M) \geq \frac{\pi}{2}$.* \square

It should be noted that if $\dim(M) = 1$ and $\tau(M) \geq 1$ then $\text{syst}(M) \geq 2\pi$ and hence $\text{inj}(M) \geq \pi$. The standard circle shows that this is optimal. We do not know by how much can one improve the bound given in Corollary 1 in other dimensions.

3.3.3 Volumes of balls

Although we will be mostly interested in the volume of sets $\mathbb{B}^M(x, r) = M \cap \mathbb{B}(x, r)$ we start by considering the volume of actual metric balls in the manifold. First note that having bounds on the curvature, we also get bounds on the volumes of balls, at least as long

as the radius remains below the injectivity radius. More precisely, if M is a Riemannian manifold of dimension $d \geq 2$ and with curvature pinched in $[-2, 1]$ then we have

$$\text{vol}(B^{S^d}(r)) \leq \text{vol}(B^M(x, r)) \leq \text{vol}(B^{\frac{1}{\sqrt{2}}\mathbb{H}^d}(r)) \quad (3.6)$$

for all $r < \text{inj}_M(x)$ (see [Berger] and [Berger]). Here, the sphere S^d and the scaled hyperbolic space $\frac{1}{\sqrt{2}}\mathbb{H}^d$ are respectively the simply connected complete d -dimensional manifolds of constant curvature 1 and -2 . Having bounds for the volumes of balls we also have bounds for the ratios between volumes of balls of different radius. We can get however somewhat better results:

Proposition 8. *Suppose that M is a d -dimensional ($d \geq 2$) Riemannian manifold with sectional curvature pinched in $[-2, 1]$. We then have*

$$\frac{R}{r} \cdot \left(\frac{\sin(R)}{\sin(r)} \right)^{d-1} \leq \frac{\text{vol}(B^M(R, x))}{\text{vol}(B^M(r, x))} \leq \frac{\int_0^{\sqrt{2} \cdot R} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot r} \sinh^{d-1}(t) dt}$$

for any two $0 < r < R < \min\{\pi, \text{inj}(M)\}$ and any $x \in M$.

Proof. We begin with the upper bound. We get from the curvature bound $\kappa \geq -2$ and the Bishop-Gromov comparison theorem [Berger] that

$$\frac{\text{vol}(B^M(x, R))}{\text{vol}(B^M(x, r))} \leq \frac{\text{vol}(B^{\frac{1}{\sqrt{2}}\mathbb{H}^d}(R))}{\text{vol}(B^{\frac{1}{\sqrt{2}}\mathbb{H}^d}(r))} = \frac{\int_0^{\sqrt{2} \cdot R} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot r} \sinh^{d-1}(t) dt}$$

for all $x \in M$, as we wanted.

Let us now deal with the lower bound. Well, the fact that M has curvature pinched from above by 1 implies that M is locally a CAT(1)-space—see [BH13; CEE75] for facts about CAT(κ)-spaces and comparison geometry. Recall now that we are working at a scale smaller than π and the injectivity radius. In particular, the CAT(1) property implies that geodesic triangles we encounter are thinner in our manifold than in S^d . This implies in particular that, for $0 < r < R < \min\{\pi, \text{inj}(M)\}$, the radial projection

$$\text{proj} : S^M(x, R) \rightarrow S^M(x, r), \quad \text{proj}(y) = \exp_x(r \cdot R^{-1} \cdot \exp_x^{-1}(y))$$

contracts distances more (expands distances less) than the corresponding map in the

sphere, meaning that proj is $\frac{\sin(r)}{\sin(R)}$ -Lipschitz. We deduce thus that

$$\text{Area}(S^M(x, r)) \leq \left(\frac{\sin(r)}{\sin(R)} \right)^{d-1} \cdot \text{Area}(S^M(x, R))$$

where $S^M(x, r) = \partial B^M(x, r)$ is the distance- r -sphere in M and $\text{Area}(\cdot)$ stands for the $(d-1)$ -dimensional volume. Anyways, if we set $T = \frac{r}{R}$ then we get from the co-area formula that

$$\begin{aligned} \text{vol}(B^M(x, r)) &= \int_0^r \text{Area}(S^M(x, t)) dt \\ &\stackrel{s=\frac{1}{T}t}{=} \int_0^R \text{Area}(S^M(x, T \cdot s)) \cdot T ds \\ &\leq \frac{r}{R} \cdot \int_0^R \left(\frac{\sin(Ts)}{\sin(s)} \right)^{d-1} \text{Area}(S^M(x, s)) ds \end{aligned}$$

The function

$$(0, \pi) \rightarrow \mathbb{R}, \quad s \rightarrow \frac{\sin(Ts)}{\sin(s)}$$

is monotonically increasing (because $T = \frac{r}{R} \in (0, 1)$). This means that

$$\begin{aligned} \text{vol}(B^M(x, r)) &\leq \frac{r}{R} \cdot \left(\frac{\sin(r)}{\sin(R)} \right)^{d-1} \cdot \int_0^R \text{Area}(S^M(x, s)) ds \\ &= \frac{r}{R} \cdot \left(\frac{\sin(r)}{\sin(R)} \right)^{d-1} \cdot \text{vol}(B^M(x, R)) \end{aligned}$$

And we are done. □

We come now to the result we really care about:

Corollary 2. *Suppose that $M \subset \mathbb{R}^s$ is a closed d -dimensional ($d \geq 1$) submanifold with reach $\tau(M) \geq 1$. We then have*

$$\text{vol}(B^{S^d}(r)) \leq \text{vol}(\mathbb{B}^M(x, r)) \leq \text{vol} \left(B^{\frac{1}{\sqrt{2}}\mathbb{H}^d} \left(2 \arcsin \left(\frac{r}{2} \right) \right) \right)$$

and

$$\frac{\frac{R}{2}}{\arcsin(\frac{r}{2})} \left(\frac{\sin(R)}{\sin(2 \arcsin(\frac{r}{2}))} \right)^{d-1} \leq \frac{\text{vol}(\mathbb{B}^M(x, R))}{\text{vol}(\mathbb{B}^M(x, r))} \leq \frac{\int_0^{\sqrt{2} \cdot 2 \cdot \arcsin(\frac{R}{2})} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot r} \sinh^{d-1}(t) dt}$$

for any two $0 < r < R < 1$ and any $x \in M$.

Proof. Let's assume for the time being that $d \geq 2$. From part (1) in Proposition 7 we get for any $t < 2$ that

$$B^M(x, t) \subset \mathbb{B}^M(x, t) \subset B^M\left(x, 2 \arcsin\left(\frac{t}{2}\right)\right) \quad (3.7)$$

The first claim then follows directly from (3.6).

On the other hand, if we combine (3.7) with the upper bound in Proposition 8 we get that

$$\begin{aligned} \frac{\text{vol}(\mathbb{B}^M(x, R))}{\text{vol}(\mathbb{B}^M(x, r))} &\leq \frac{\text{vol}\left(B^M\left(x, 2 \arcsin\left(\frac{R}{2}\right)\right)\right)}{\text{vol}(B^M(r, x))} \\ &\leq \frac{\int_0^{\sqrt{2} \cdot 2 \arcsin(\frac{R}{2})} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot r} \sinh^{d-1}(t) dt}, \end{aligned}$$

and we are done with the upper bound of the second claim. The lower bound is obtained analogously and we leave the details to the reader.

So far we have been focusing on the case of dimension $d \geq 2$. In dimension $d = 1$ we actually get from (3.7) that

$$R - r \leq \text{vol}(\mathbb{B}^M(x, R)) - \text{vol}(\mathbb{B}^M(x, r)) \leq 2 \arcsin\left(\frac{R - r}{2}\right)$$

This implies directly that

$$\frac{R}{r} \leq \frac{\text{vol}(\mathbb{B}^M(x, R))}{\text{vol}(\mathbb{B}^M(x, r))} \leq 1 + \frac{2 \arcsin\left(\frac{R-r}{2}\right)}{r}$$

for all $0 < r < R < 1$. The so-obtained bound for $d = 1$ is slightly better than the one we had claimed. We also note that this bound still works for $0 < r < R < 2$. \square

3.3.4 Volume of thick diagonal

Our goal here is to estimate how the volume $\text{vol}(DM(\epsilon))$ of the ϵ -thick diagonal

$$DM(\epsilon) = \{(x, y) \in M \times M \text{ with } |x - y| \leq \epsilon\}$$

of a reach-1 submanifold $M \subset \mathbb{R}^s$ varies when we replace ϵ by something else. Here the volume is computed as a subset of the Riemannian manifold $M \times M$, but can be expressed as an integral over M as follows:

$$\text{vol}(DM(\epsilon)) = \int_M \text{vol}(\mathbb{B}^M(x, \epsilon)) \, dx \quad (3.8)$$

In other words, $\text{vol}(DM(\epsilon))$ is nothing other than the correlation integral at scale ϵ of the Riemannian measure of M when considered as a measure on the ambient euclidean space \mathbb{R}^s . We stress that the thick diagonal is defined in terms of the ambient distance in euclidean space, not in terms of the intrinsic distance of M . Anyways, now we prove the following:

Theorem 16. *Suppose that $M \subset \mathbb{R}^s$ is a d -dimensional ($d \geq 1$) closed submanifold with reach $\tau(M) \geq 1$. Then we have*

$$\frac{\frac{\epsilon_1}{2}}{\arcsin(\frac{\epsilon_2}{2})} \left(\frac{\sin(\epsilon_1)}{\sin(2 \arcsin \frac{\epsilon_2}{2})} \right)^{d-1} \leq \frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))} \leq \frac{\int_0^{\sqrt{2} \cdot 2 \cdot \arcsin(\frac{\epsilon_1}{2})} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot \epsilon_2} \sinh^{d-1}(t) dt}$$

for any two $0 < r < R < 1$.

Proof. From the expression (3.8) we get that

$$\min_{x \in M} \frac{\text{vol}(\mathbb{B}^M(\epsilon_1, x))}{\text{vol}(\mathbb{B}^M(\epsilon_2, x))} \leq \frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))} \leq \max_{x \in M} \frac{\text{vol}(\mathbb{B}^M(\epsilon_1, x))}{\text{vol}(\mathbb{B}^M(\epsilon_2, x))}$$

Now the claim follows from Corollary 2. □

3.3.5 The gap

As we mentioned in the introduction, if we sample more and more points from a manifold M and we apply the algorithm, then what we are doing is computing the quantity

$$\frac{\log(\text{vol}(DM(\epsilon_1))) - \log(\text{vol}(DM(\epsilon_2)))}{\log(\epsilon_1) - \log(\epsilon_2)}.$$

Now, armed with Theorem 16 we could analyze what happens when one of the scales ϵ_1 and ϵ_2 tends to 0, or when the gap between them tends to 0, or when the dimension grows. All of this would be nice and well, but what we actually care about is to find scales that

on the one hand keep our estimator reliable while working with as few points as possible. Implementing numerically a procedure described in Section 3.4.4 below, we find convenient scales for the dimensions we are mostly interested in. Note that the condition $R < 1$ can be replaced by $R < 2$ for dimension 1, as explained in the proof of Corollary 2.

Corollary 3. *Let $M \subset \mathbb{R}^s$ be a submanifold of dimension $d = 1, 2, \dots, 10$ and with reach $\tau(M) \geq 1$, and let ϵ_1, ϵ_2 and gap_d be as in the table below. Then, for every x in M , we have*

$$d - \frac{1}{2} + \text{gap}_d \leq \frac{\log \left(\frac{\text{vol}(\mathbb{B}^M(\epsilon_1, x))}{\text{vol}(\mathbb{B}^M(\epsilon_2, x))} \right)}{\log \left(\frac{\epsilon_1}{\epsilon_2} \right)} \leq d + \frac{1}{2} - \text{gap}_d.$$

d	ϵ_1	ϵ_2	gap_d
1	1.5	0.19	0.463241
2	0.78	0.2	0.387573
3	0.63	0.23	0.307476
4	0.54	0.23	0.249891
5	0.46	0.22	0.223958
6	0.4	0.21	0.208521
7	0.36	0.21	0.178814
8	0.33	0.2	0.166892
9	0.31	0.19	0.155560
10	0.29	0.18	0.152528

3.4 Sampling the thick diagonal

In this section we will be still assuming that $M \subset \mathbb{R}^s$ is a closed d -dimensional submanifold ($d \geq 1$) with reach $\tau(M) \geq 1$. Basically our goal is to bound the number of points that we have to sample from M to get a decent result when we use (3.1).

3.4.1 Some probability

Suppose that we have a symmetric, say bounded, function

$$f : M \times M \rightarrow \mathbb{R}$$

We are interested in the sequence of random variables

$$X_n^f : M^{\mathbb{N}} \rightarrow \mathbb{R}, \quad (x_1, x_2, \dots) \mapsto \sum_{i,j \leq n, i \neq j} f(x_i, x_j) \quad (3.9)$$

when n tends to ∞ . Here we have endowed M with the probability measure

$$\text{Prob} = \frac{1}{\text{vol}(M)} \text{vol}$$

proportional to the Riemannian measure. Accordingly, M^k and $M^{\mathbb{N}}$ are all endowed with the corresponding product measure, again denoted by Prob .

Being the sum of random variables, the expectation and variance of X_n^f are easy to get. Here they are:

$$\begin{aligned} E(X_n^f) &= n(n-1) \cdot E(f) \\ \text{Var}(X_n^f) &= 2 \cdot n(n-1) \cdot \text{Var}(f) + 4 \cdot n(n-1)(n-2) \cdot \text{cov}(f) \end{aligned}$$

where $\text{cov}(f)$ is the co-variance of $(x_1, \dots, x_n) \mapsto f(x_1, x_2)$ and $(x_1, \dots, x_n) \mapsto f(x_1, x_3)$, or in a formula

$$\text{cov}(f) = \int_{M \times M \times M} f(x, y) \cdot f(x, z) \, d\text{Prob}(x, z, y) - E(f)^2.$$

Besides the expectation and the variance, what we will need to estimate is the quantity $\frac{\text{Var}}{E^2}$ for the random variables X_n^f . Well, this is what we get if we just use our expressions for the expectation and the variance:

$$\frac{\text{Var}(X_n^f)}{E(X_n^f)^2} = \frac{2}{n(n-1)} \cdot \frac{\text{Var}(f)}{E(f)^2} + \frac{4(n-2)}{n(n-1)} \cdot \frac{\text{cov}(f)}{E(f)^2} \quad (3.10)$$

The reason why we will care about this last quantity is the following surely standard consequence of the Bienaymé-Chebyshev inequality:

Lemma 20. *For any integrable random variable X in a probability space X we have*

$$\text{Prob} \left(\left| \log \left(\frac{X}{E(X)} \right) \right| > \delta \right) \leq \frac{1}{(1 - e^{-\delta})^2} \cdot \frac{\text{Var}(X)}{E(X)^2}$$

Proof. Well, let us compute

$$\begin{aligned} \text{Prob} \left(\left| \log \left(\frac{X}{E(X)} \right) \right| > \delta \right) &= \text{Prob} \left(\frac{X}{E(X)} \notin [e^{-\delta}, e^{\delta}] \right) \\ &\leq \text{Prob} \left(\left| \frac{X}{E(X)} - 1 \right| > 1 - e^{-\delta} \right) \\ &\leq \text{Prob} (|X - E(X)| > (1 - e^{-\delta}) \cdot E(X)) \end{aligned}$$

Setting

$$k = \frac{(1 - e^{-\delta}) \cdot E(X)}{\sqrt{\text{Var}(X)}}$$

in the standard Bienaymé-Chebysheff inequality

$$\text{Prob} (|X - E(X)| \geq k \cdot \sqrt{\text{Var}(X)}) \leq k^{-2}$$

we get

$$\text{Prob} \left(\left| \log \left(\frac{X}{E(X)} \right) \right| > \delta \right) \leq \frac{\text{Var}(X)}{(1 - e^{-\delta})^2 \cdot E(X)^2}$$

as we had claimed. □

3.4.2 The function we care about

We are going to be interested in all of this in the case that $f = f_\epsilon$ is the characteristic function of $DM(\epsilon)$, that is

$$f_\epsilon(x, y) = \begin{cases} 1 & \text{if } |x - y| \leq \epsilon \\ 0 & \text{otherwise} \end{cases} \quad (3.11)$$

This function satisfies that

$$E(f_\epsilon) = \frac{\text{vol}(DM(\epsilon))}{\text{vol}(M)^2}$$

and hence we get that

$$E(X_n^{f_\epsilon}) = n(n-1) \cdot \frac{\text{vol}(DM(\epsilon))}{\text{vol}(M)^2}$$

for all $n \geq 2$. Besides knowing the expectation, to be able to use Lemma 20 when we need to know, or at least estimate, is the quantity $\frac{\text{Var}}{E^2}$. To apply (3.10) we need first to be able to estimate the variance and covariance of f_ϵ . Well, since f_ϵ only takes the values 0 and 1,

the variance is easily calculated:

$$\text{Var}(f_\epsilon) = E(f_\epsilon) - E(f_\epsilon)^2 = \frac{\text{vol}(DM(\epsilon))}{\text{vol}(M)^2} - \left(\frac{\text{vol}(DM(\epsilon))}{\text{vol}(M)^2} \right)^2$$

When it comes to the covariance we have

$$\begin{aligned} \text{cov}(f_\epsilon) &= \int \left(\frac{\text{vol}(\mathbb{B}^M(x, \epsilon))}{\text{vol}(M)} \right)^2 d\text{Prob}(x) - \frac{\text{vol}(DM(\epsilon))^2}{\text{vol}(M)^4} \\ &= \int \left(\frac{\text{vol}(\mathbb{B}^M(x, \epsilon))}{\text{vol}(M)} \right)^2 d\text{Prob}(x) - \left(\int_M \frac{\text{vol}(\mathbb{B}^M(x, \epsilon))}{\text{vol}(M)} d\text{Prob}(x) \right)^2 \\ &= \text{Var} \left(x \mapsto \frac{\text{vol}(\mathbb{B}^M(x, \epsilon))}{\text{vol}(M)} \right) \end{aligned}$$

This means that when the volume of $\mathbb{B}^M(x, \epsilon) = M \cap \mathbb{B}(x, \epsilon)$ is constant then the covariance vanishes. This is for example the case for $M = S^d \subset \mathbb{R}^{d+1}$ or for the Clifford torus $M = \mathbb{T}^d \subset \mathbb{R}^{2d}$. However, in general we do not get anything better than the bound coming from Popoviciu's inequality, that is

$$\text{cov}(f_\epsilon) \leq \frac{(V_{\max}^M(\epsilon) - V_{\min}^M(\epsilon))^2}{4 \cdot \text{vol}(M)^2}$$

where we have set

$$V_{\max}^M(\epsilon) = \max_{x \in M} \text{vol}(\mathbb{B}^M(x, \epsilon)) \text{ and } V_{\min}^M(\epsilon) = \min_{x \in M} \text{vol}(\mathbb{B}^M(x, \epsilon))$$

Now, using (3.10), the bound for $\text{cov}(f_\epsilon)$, as well as the bound $\text{vol}(DM(\epsilon)) \geq \text{vol}(M) \cdot V_{\min}(\epsilon)$ we get that

$$\frac{\text{Var}(X_n^{f_\epsilon})}{E(X_n^{f_\epsilon})^2} \leq \frac{2}{(n-1)^2} \cdot \frac{\text{vol}(M)}{V_{\min}^M(\epsilon)} + \frac{1}{n-1} \cdot \left(\frac{V_{\max}^M(\epsilon)}{V_{\min}^M(\epsilon)} - 1 \right)^2$$

To get bounds that only depend on the dimension and on $\epsilon < 2$ recall that from Corollary 2 we get that

$$\begin{aligned} V_{\min}^M(\epsilon) &\geq \text{vol}(B^{S^d}(\epsilon)) \stackrel{\text{def}}{=} \mathcal{V}(\epsilon) \\ \frac{V_{\max}^M(\epsilon)}{V_{\min}^M(\epsilon)} &\leq \frac{\text{vol} \left(B^{\frac{1}{\sqrt{2}} \mathbb{H}^d} \left(2 \arcsin \left(\frac{\epsilon}{2} \right) \right) \right)}{\text{vol}(B^{S^d}(\epsilon))} \stackrel{\text{def}}{=} \mathcal{R}(\epsilon) \end{aligned} \tag{3.12}$$

Using these bounds we get

$$\frac{\text{Var}(X_n^{f_\epsilon})}{E(X_n^{f_\epsilon})^2} \leq \frac{2}{(n-1)^2} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon)} + \frac{1}{n-1} \cdot (\mathcal{R}(\epsilon) - 1)^2$$

We record what we have so far:

Lemma 21. *Let $M \subset \mathbb{R}^s$ be a closed submanifold with dimension $\dim(M) = d$ and reach $\tau(M) \geq 1$, and for some $\epsilon < 2$ and $n \in \mathbb{N}$ consider f_ϵ and $X_n^{f_\epsilon}$ as in (3.11) and (3.9). Then we have*

$$\begin{aligned} E(X_n^{f_\epsilon}) &= n(n-1) \cdot \frac{\text{vol}(DM(\epsilon))}{\text{vol}(M)^2} \\ \frac{\text{Var}(X_n^{f_\epsilon})}{E(X_n^{f_\epsilon})^2} &\leq \frac{2}{(n-1)^2} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon)} + \frac{1}{n-1} \cdot (\mathcal{R}(\epsilon) - 1)^2 \end{aligned}$$

where $\mathcal{V}(\epsilon)$ and $\mathcal{R}(\epsilon)$ are as in (3.12). □

Before moving any further let us give explicit formulas for $\mathcal{V}(\epsilon)$ and $\mathcal{R}(\epsilon)$:

$$\begin{aligned} \mathcal{V}(\epsilon) &= \text{vol}(S^{d-1}) \cdot \int_0^\epsilon \sin^{d-1}(t) dt \\ \mathcal{R}(\epsilon) &= \frac{2^{-\frac{d}{2}} \int_0^{2\sqrt{2} \arcsin \frac{\epsilon}{2}} \sinh(t)^{d-1} dt}{\int_0^\epsilon \sin^{d-1}(t) dt} \end{aligned} \tag{3.13}$$

Note also that the two summands in the bound for $\frac{\text{Var}(X_n^{f_\epsilon})}{E(X_n^{f_\epsilon})^2}$ in Lemma 21 are rather different. Assume for example that d is fixed. Then the weight of the second factor decreases when ϵ decreases. On the other hand the value of the first one explodes. Recall also that the second factor can be ignored if cov vanishes, that is if all balls $\mathbb{B}^M(x, \epsilon)$ have the same volume.

3.4.3 Some technical results

Recall that to estimate the dimension of M via (3.1), or equivalently via (3.5), what we do is to take, for two scales $\epsilon_1 > \epsilon_2$ random values of $X_n^{\epsilon_1} = X_n^{f_{\epsilon_1}}$ and $X_n^{\epsilon_2} = X_n^{f_{\epsilon_2}}$, compute

$$\frac{\log \frac{X_n^{\epsilon_1}}{n(n-1)} - \log \frac{X_n^{\epsilon_2}}{n(n-1)}}{\log \epsilon_1 - \log \epsilon_2}$$

and hope the that obtained value has something to do with $\dim(M)$. Well, what we get from Lemma 20 is an estimate of the probability that this value is far from the expectation. Indeed, we have

$$\begin{aligned}
 & \text{Prob} \left(\left| \frac{\log X_n^{\epsilon_1} - \log X_n^{\epsilon_2}}{\log \epsilon_1 - \log \epsilon_2} - \frac{\log \text{vol}(DM(\epsilon_1)) - \log \text{vol}(DM(\epsilon_2))}{\log \epsilon_1 - \log \epsilon_2} \right| > \rho \right) \\
 &= \text{Prob} \left(\left| \frac{\log X_n^{\epsilon_1} - \log X_n^{\epsilon_2}}{\log \epsilon_1 - \log \epsilon_2} - \frac{\log E(X_n^{\epsilon_1}) - \log E(X_n^{\epsilon_2})}{\log \epsilon_1 - \log \epsilon_2} \right| > \rho \right) \\
 &= \text{Prob} \left(\left| \log \left(\frac{X_n^{\epsilon_1}}{E(X_n^{\epsilon_1})} \right) - \log \left(\frac{X_n^{\epsilon_2}}{E(X_n^{\epsilon_2})} \right) \right| > \log \left(\left(\frac{\epsilon_1}{\epsilon_2} \right)^\rho \right) \right) \\
 &\leq \sum_{i=1,2} \text{Prob} \left(\left| \log \left(\frac{X_n^{\epsilon_i}}{E(X_n^{\epsilon_i})} \right) \right| > \log \left(\left(\frac{\epsilon_1}{\epsilon_2} \right)^{\frac{\rho}{2}} \right) \right) \\
 &\leq \frac{1}{\left(1 - \left(\frac{\epsilon_2}{\epsilon_1} \right)^{\frac{\rho}{2}} \right)^2} \sum_{i=1,2} \frac{\text{Var}(X_n^{\epsilon_i})}{E(X_n^{\epsilon_i})^2}.
 \end{aligned}$$

Plugging in the statement of Lemma 21 we get:

Theorem 17. *Let $M \subset \mathbb{R}^s$ be a closed submanifold with reach $\tau(M) \geq 1$, pick two scales $0 < \epsilon_2 < \epsilon_1$. Also, for $n \geq 2$ set*

$$\rho = \sum_{i=1,2} \left(\frac{2}{(n-1)^2} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon_i)} + \frac{1}{n-1} \cdot (\mathcal{R}(\epsilon_i) - 1)^2 \right)$$

Then we have

$$\text{Prob} \left(\left| \frac{\log X_n^{\epsilon_1} - \log X_n^{\epsilon_2}}{\log \epsilon_1 - \log \epsilon_2} - \frac{\log \left(\frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))} \right)}{\log \epsilon_1 - \log \epsilon_2} \right| > \Delta \right) \leq \frac{\rho}{\left(1 - \left(\frac{\epsilon_2}{\epsilon_1} \right)^{\frac{\Delta}{2}} \right)^2}$$

for any $\Delta > 0$. □

Let us get a slightly more user friendly version:

Corollary 4. *With the same assumptions and notation as in Theorem 17 suppose that for some positive α_i 's with $\alpha_1 + \alpha_2 = 1$ and for some $\rho > 0$ we have*

$$n \geq 1 + \frac{1}{\alpha_i \cdot \rho} \cdot (\mathcal{R}(\epsilon_i) - 1)^2 + \sqrt{\frac{2}{\alpha_i \cdot \rho} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon_i)}}$$

for $i = 1, 2$. Then we also have

$$P \left(\left| \frac{\log X_n^{\epsilon_1} - \log X_n^{\epsilon_2}}{\log \epsilon_1 - \log \epsilon_2} - \frac{\log \left(\frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))} \right)}{\log \epsilon_1 - \log \epsilon_2} \right| > \Delta \right) \leq \rho \cdot \left(1 - \left(\frac{\epsilon_2}{\epsilon_1} \right)^{\frac{\Delta}{2}} \right)^{-2}$$

Proof. In terms of Theorem 17 what we have to do is to guarantee for $i = 1, 2$ that

$$\alpha_i \rho \geq \frac{2}{(n-1)^2} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon_i)} + \frac{1}{n-1} \cdot (\mathcal{R}(\epsilon_i) - 1)^2$$

It thus suffices to ensure that $n-1$ is larger than the solution X of the equation

$$a \stackrel{\text{def}}{=} \alpha_i \rho = \frac{1}{X^2} \cdot \frac{2 \cdot \text{vol}(M)}{\mathcal{V}(\epsilon_i)} + \frac{1}{X} \cdot (\mathcal{R}(\epsilon_i) - 1)^2 \stackrel{\text{def}}{=} \frac{1}{X^2} c + \frac{1}{X} b$$

This is now a quadratic equation with positive solution

$$X = \frac{b}{a} + \sqrt{\frac{c}{a}}$$

The claim follows. □

Again, if M is such that all balls $\mathbb{B}^M(x, \epsilon)$ have constant volume, then one can replace the first displayed equation in Corollary 4 by

$$n \geq 1 + \sqrt{\frac{2}{\alpha_i \cdot \rho} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon_i)}}.$$

3.4.4 Searching decent scales

Given the dimension d and the volume $\text{Vol}(M)$, what are the optimal scales to run (3.1) so that we have $\geq 90\%$ success probability? Let's see how we could find, if not the optimal scales, at least decent ones. First, for $1 > \epsilon_1 > \epsilon_2 > 0$ consider the quantity

$$\Delta = \Delta_{\epsilon_1, \epsilon_2} = \max\{\Delta_1, \Delta_2\}$$

where

$$\begin{aligned}\Delta_1 &\leq \frac{1}{2} - \frac{\log \left(\frac{\int_0^{\sqrt{2} \cdot 2 \cdot \arcsin(\frac{\epsilon_1}{2})} \sinh^{d-1}(t) dt}{\int_0^{\sqrt{2} \cdot \epsilon_2} \sinh^{d-1}(t) dt} \right)}{\log \frac{\epsilon_1}{\epsilon_2}} + d \\ \Delta_2 &\leq \frac{1}{2} + \frac{\log \left(\frac{\epsilon_1}{2 \cdot \arcsin(\frac{\epsilon_2}{2})} \cdot \left(\frac{\sin(\epsilon_1)}{\sin(2 \cdot \arcsin(\frac{\epsilon_2}{2}))} \right)^{d-1} \right)}{\log \frac{\epsilon_1}{\epsilon_2}} - d\end{aligned}$$

From Lemma 21 and Theorem 16 we get for all n that

$$\left| \frac{\log \frac{E(X_n^{\epsilon_1})}{E(X_n^{\epsilon_2})}}{\log \epsilon_1 - \log \epsilon_2} - d \right| = \left| \frac{\log \frac{\text{vol}(DM(\epsilon_1))}{\text{vol}(DM(\epsilon_2))}}{\log \epsilon_1 - \log \epsilon_2} - d \right| \leq \frac{1}{2} - \Delta$$

Note that Theorem 16 has the condition $R < 1$, but for dimension 1, it be replaced by the condition $R < 2$, as explained in the proof of Corollary 2.

It follows that, as long as $\Delta > 0$, if we take a very large number of points n then we get that it is very likely that the (3.1) returns the value d . Now, how many points we do actually need if we want to guarantee a 90% rate of success? Well, with notation as in Theorem 17 we start by setting

$$\rho = \rho_{\epsilon_1, \epsilon_2} = \frac{1}{10} \cdot \left(1 - \left(\frac{\epsilon_2}{\epsilon_1} \right)^{\frac{\Delta}{2}} \right)^2 \quad (3.14)$$

Now, once we have ρ we get from Corollary 4 that if we take $\alpha \in (0, 1)$, set $\alpha_1 = \alpha$ and $\alpha_2 = 1 - \alpha$, and if we take at least

$$\begin{aligned}n(\epsilon_1, \epsilon_2, \alpha, \text{vol}(M)) &= \max_{i=1,2} \left(1 + \frac{1}{\alpha_i \cdot \rho} \cdot (\mathcal{R}(\epsilon_i) - 1)^2 + \sqrt{\frac{2}{\alpha_i \cdot \rho} \cdot \frac{\text{vol}(M)}{\mathcal{V}(\epsilon_i)}} \right) \\ &\leq 1 + \max_{i=1,2} \left(\frac{1}{\alpha_i \cdot \rho} \cdot (\mathcal{R}(\epsilon_i) - 1)^2 \right) + \\ &\quad + \left(\max_{i=1,2} \sqrt{\frac{2}{\alpha_i \cdot \rho \cdot \mathcal{V}(\epsilon_i)^2}} \right) \cdot \text{vol}(M)^{\frac{1}{2}}\end{aligned}$$

points, then

$$\text{Prob} \left(\left| \frac{\log \frac{X_n^{\epsilon_1}}{X_n^{\epsilon_2}}}{\log \epsilon_1 - \log \epsilon_2} - d \right| < \frac{1}{2} \right) > 90\%. \quad (3.15)$$

d	ϵ_1	ϵ_2	α_1
1	1.5	0.19	0.15
2	0.78	0.2	0.11
3	0.63	0.23	0.09
4	0.54	0.23	0.06
5	0.46	0.22	0.04
6	0.4	0.21	0.03
7	0.36	0.21	0.03
8	0.33	0.2	0.02
9	0.31	0.19	0.02
10	0.29	0.18	0.01

Table 3.2 – Decent scales for $\text{vol} = \text{vol}(\mathbb{T}^d)$ in dimension $d = 1, 2, \dots, 10$. It is evident that the values in Table 3.2 can change if instead of using $\text{vol}(\mathbb{T}^d)$ as an input one chooses any other value. However, for whatever it is worth, if instead one chooses $10 \cdot \text{vol}(\mathbb{T}^d)$ or even $100 \cdot \text{vol}(\mathbb{T}^d)$ then there are not much changes: in small dimensions (that is, up to dimension 3) the scales increase a bit, but for dimensions at least 4 nothing changes.

If we are interested in manifolds with $\text{Vol}(M) \leq V$ then (3.15) holds as long as we take at least

$$\min_{\substack{1 > \epsilon_1 > \epsilon_2 > 0 \\ \text{with } \Delta_{\epsilon_1, \epsilon_2} > 0}} \min_{\alpha \in (0,1)} n(\epsilon_1, \epsilon_2, \alpha) \quad (3.16)$$

points and we use (3.1) with constants $1 > \epsilon_1 > \epsilon_2 > 0$. Now, to find decent scales we can now minimize (3.16). A program which numerically approximates that is available at [Gri22].

In fact, running also the program in each $d = 1, 2, \dots, 10$ for the volume of the corresponding d -dimensional torus we get that $(\epsilon_1, \epsilon_2, \alpha)$ as in Table 3.2 give smallish values for (3.16). If we plug these constants in the formula for $n(\epsilon_1, \epsilon_2, \alpha, \text{vol}(M))$ that we gave above we recover the statement of Theorem 14 stated in the introduction:

Theorem 14. *For $d = 1, \dots, 10$ let ϵ_1 and ϵ_2 be scales as in the table below. Also, given a closed d -dimensional manifold $M \subset \mathbb{R}^s$ with reach $\tau(M) \geq 1$ let n be also as in the following table:*

d	ϵ_1	ϵ_2	n
1	1.5	0.19	$9 + 21 \cdot \text{vol}(M)^{\frac{1}{2}}$
2	0.78	0.2	$94 + 58 \cdot \text{vol}(M)^{\frac{1}{2}}$
3	0.63	0.23	$635 + 146 \cdot \text{vol}(M)^{\frac{1}{2}}$
4	0.54	0.23	$2786 + 392 \cdot \text{vol}(M)^{\frac{1}{2}}$
5	0.46	0.22	$7013 + 1119 \cdot \text{vol}(M)^{\frac{1}{2}}$
6	0.4	0.21	$13221 + 3366 \cdot \text{vol}(M)^{\frac{1}{2}}$
7	0.36	0.21	$25138 + 10644 \cdot \text{vol}(M)^{\frac{1}{2}}$
8	0.33	0.2	$50033 + 34890 \cdot \text{vol}(M)^{\frac{1}{2}}$
9	0.31	0.19	$63876 + 119533 \cdot \text{vol}(M)^{\frac{1}{2}}$
10	0.29	0.18	$139412 + 425554 \cdot \text{vol}(M)^{\frac{1}{2}}$

Then, if we sample independently and according to the riemannian volume form a subset $X \subset M$ consisting of at least n points, then we have

$$\dim_{\text{Corr}(\epsilon_1, \epsilon_2)}(X) = d$$

with at least 90% probability.

In the next section we discuss some (much smaller) heuristic bounds, discuss some numerical experiments, and compare with the performance of other estimators.

3.5 Heuristics

Theorem 17 gives us a bound for the number of points needed in a data set to be able to get from (3.1) at least 90% of the time its dimension. In concrete examples, we expect that this confidence level can be achieved with significantly less points. We will discuss this difference between theory and practice, suggesting a simpler heuristic model supported by computational examples.

Heuristic bound

We begin by discussing a heuristic model representing an ideal situation without curvature. More concretely we will be running (3.1) at some scales ϵ_1 and ϵ_2 at which we

can ignore curvature effects. In effect, we will act as if all balls in M of radius at most ϵ_1 were euclidean and totally geodesic.

The first observation is that the statistic (3.1) is computed from information extracted from the distances $|x - y|$ for those (unordered) pairs

$$\{x, y\} \in PX(\epsilon) = \{\{x, y\} \subset X \text{ with } x \neq y \text{ with } |x - y| \leq \epsilon\}$$

rather than from the points themselves. This means that the performance of the algorithm should depend on $|PX(\epsilon_1)|$ and on the dimension d , instead of directly on the total number n of points.

We think of the distance $|x - y|$ for $\{x, y\} \in PX(\epsilon_1)$ as a random variable, and from now on, we will take the point of view that we have $N = |PX(\epsilon_1)|$ random variables $(X_i)_{1 \leq i \leq N}$ given by taking the distance between pairs of points at distance at most ϵ_1 and sampled uniformly on M . We then consider the variables Y_i equal to 1 if X_i is smaller than ϵ_2 and 0 otherwise. With this notation in place, the estimator (3.1) becomes

$$\dim_{\text{Corr}(\epsilon_1, \epsilon_2)(X)} = \text{Round} \left(\frac{\log(\frac{1}{N} \sum_{i=1}^N Y_i)}{\log(\epsilon_2) - \log(\epsilon_1)} \right)$$

Note that since we are assuming that all the balls are euclidean, the mean value of Y_i is then $E_d = (\epsilon_2/\epsilon_1)^d$ and as the variables Y_i only take the values 0 and 1, their variance is $\sigma_d^2 = E_d - E_d^2$.

In reality, the variables X_i , and thus the variables Y_i have no reason to be independent. Still, most of them are when the volume is large when compared to the size of the data set. So, from now on, we will put ourselves in the ideal situation that the N variables Y_i are independent. Independence implies that the distribution of the sample mean $Z = \frac{1}{N} \sum_i Y_i$ is binomial of parameters N and E_d , which can be approximated by a normal distribution of mean value E_d and of variance $\frac{1}{N} \sigma_d^2$ using the central limit theorem. It is then known that the probability of Z to be in the interval $[E_d - 1.64 \cdot \sigma_d/\sqrt{N}, E_d + 1.64 \cdot \sigma_d/\sqrt{N}]$ is about 90%. If we set $\text{gap}_d = \min(E_{d-0.5} - E_d, E_d - E_{d+0.5})$, we want to find N so that $1.64 \cdot \sigma_d/\sqrt{N} = \text{gap}_d$. This number gives a number of pairs sufficient to obtain the right dimension with a confidence of 90%.

Suppose for example that the manifold M has dimension 4. For $\epsilon_1 = 0.54$ and $\epsilon_2 = 0.23$, the scales coming from Theorem 14, we can then compute the values of gap_4 and of σ_4

and deduce the required value of N .

$$\begin{aligned}\text{gap}_4 &= \min(E_{3.5} - E_4, E_4 - E_{4.5}) \simeq 0.01143 \\ \sigma_4 &= \sqrt{E_4 - E_4^2} \simeq 0.1784 \\ N &= (1.64 \cdot \sigma_4 / \text{gap}_4)^2 \simeq 655 \text{ pairs}\end{aligned}$$

This is an approximation using the central limit theorem (that is, replacing the binomial distribution by the normal distribution), but more precise computation can be done working directly with the binomial distribution. Doing this, we can take N down to 516 (see [Gri22]). This reasoning can be applied to obtain the required number of pairs for each dimension—the results are summarized in Table 3.3.

d	ϵ_1	ϵ_2	N for 90%	N for 70%
1	1.5	0.19	30	10
2	0.78	0.2	122	40
3	0.63	0.23	249	111
4	0.54	0.23	516	238
5	0.46	0.22	878	360
6	0.4	0.21	1329	554
7	0.36	0.21	1719	698
8	0.33	0.2	2481	1070
9	0.31	0.19	3900	1604
10	0.29	0.18	5849	2414

Table 3.3 – Heuristic bounds for the size of $PX(\epsilon_1)$ needed to have 90% and 70% rate of success when applying (3.1) to data sets sampled from a reach 1 manifold.

We compare next these bounds with those in Theorem 14 and then discuss a few numerical experiments.

Comparison between theoretical and heuristic bounds

A problem when comparing the heuristic bounds in Table 3.3 and those in Theorem 14 is that the former ones are given in terms of the cardinality of $PX(\epsilon_1)$ while the latter ones are given in terms of the cardinality of X .

To connect these two quantities recall that we are acting as if all ϵ_1 -balls in M were euclidean. Now, we get for example from Lemma 21 that the number $|PX(\epsilon_1)|$ of unordered pairs of points at distance at most ϵ_1 is approximately given by the formula

$$|PX(\epsilon_1)| \simeq \frac{n(n-1)}{2} \frac{\text{vol}(B^{\mathbb{R}^d}(\epsilon_1))}{\text{vol}(M)}. \quad (3.17)$$

Acting as if (3.17) were to give a perfect relation between the number of points and that of pairs, we get that to get the estimated 516 pairs in the case that $M = \mathbb{T}^4$ is the 4-dimensional torus we need 1958 data points. In comparison, Theorem 14 gives an upper bound of 18262 points.

Arguing like this we can convert the heuristic bounds in Table 3.3 to bounds for the needed cardinality of a data set in terms of the dimension and the volume of the underlying manifold—see Table 3.4.

d	heuristic n	n from Theorem 14
1	$5 \cdot \text{vol}(M)^{\frac{1}{2}}$	$9 + 21 \cdot \text{vol}(M)^{\frac{1}{2}}$
2	$12 \cdot \text{vol}(M)^{\frac{1}{2}}$	$94 + 58 \cdot \text{vol}(M)^{\frac{1}{2}}$
3	$22 \cdot \text{vol}(M)^{\frac{1}{2}}$	$635 + 146 \cdot \text{vol}(M)^{\frac{1}{2}}$
4	$50 \cdot \text{vol}(M)^{\frac{1}{2}}$	$2786 + 392 \cdot \text{vol}(M)^{\frac{1}{2}}$
5	$128 \cdot \text{vol}(M)^{\frac{1}{2}}$	$7013 + 1119 \cdot \text{vol}(M)^{\frac{1}{2}}$
6	$355 \cdot \text{vol}(M)^{\frac{1}{2}}$	$13221 + 3366 \cdot \text{vol}(M)^{\frac{1}{2}}$
7	$964 \cdot \text{vol}(M)^{\frac{1}{2}}$	$25138 + 10644 \cdot \text{vol}(M)^{\frac{1}{2}}$
8	$2949 \cdot \text{vol}(M)^{\frac{1}{2}}$	$50033 + 34890 \cdot \text{vol}(M)^{\frac{1}{2}}$
9	$9458 \cdot \text{vol}(M)^{\frac{1}{2}}$	$63876 + 119533 \cdot \text{vol}(M)^{\frac{1}{2}}$
10	$33021 \cdot \text{vol}(M)^{\frac{1}{2}}$	$139412 + 425554 \cdot \text{vol}(M)^{\frac{1}{2}}$

Table 3.4 – Comparison between the heuristic bound and the bound in Theorem 14 for the number of points that suffice to have 90% success rate when applying (3.1) to data sets sampled from a reach 1 manifold.

One should keep in mind that entries the middle column in Table 3.4 are only meaningful for $\text{vol}(M)$ large. Still, there is a very clear difference between both bounds, the heuristic bound and that from Theorem 14. This difference is at least in part due to the curvature of the submanifold M , but things are not helped by either all the nested inequalities leading to Theorem 14 or the fact that the Bienaymé-Chebyshev inequality is not very precise.

Numerical evidence

Experimental examples seem to confirm the heuristical bounds presented on Table 3.3. By numerically sampling points on different manifolds, we obtain results close to what was expected. Here is the procedure we followed:

1. Choose a manifold M of known dimension d and consider the scales and number of pairs N given by Table 3.3.
2. Uniformly and independently sample points in M until we obtain N pairs at distance ε_1 . The sampling is done by repeating a program specific to the desired manifold that samples a single point randomly and uniformly on it.
3. Estimate the dimension using estimator (3.1) with scales ε_1 and ε_2 .
4. Repeat steps 2 and 3 one hundred times and count the number of success.

We ran this experiment a variety of manifolds of reach 1 and always obtained a rate of success in a range of $\pm 6\%$ of the target rate. This difference between the actual rate and the target would be totally normal even in an ideal situation. Indeed, repeating an experiment 100 time with a probability of success of 90% gives a result with a standard deviation of 3 successes. Every experimental result then falls into the usual range of two times the standard deviation.

The manifolds were chosen to observe different situations:

- **Worms:** 0-level set of a randomly produced function on \mathbb{R}^2 . The precise algorithm used for generating these manifolds and sampling from them is available in [Gri22], as well as the algorithms for the other manifolds.
- **Rotation torus:** Rotate around the z -axis the circle in the xy -plane of radius 1 and center $(2, 0, 0)$. This surface has reach 1 and there is a mix of positive and negative curvature.
- **Clifford torus:** The product $\mathbb{T}^d = S^1 \times \dots \times S^1 \subset \mathbb{R}^{2d}$ of d -circles of radius 1. The Clifford torus is curved in Euclidean space but is flat as a Riemannian manifold.
- **Flat torus:** This is an ideal situation. We consider namely the abstract manifold $\mathbb{R}^d / 2\pi \cdot \mathbb{Z}^d$ with its inner distance—it is not embedded in some larger Euclidean space.
- **Swiss roll:** The Swiss roll is one of the standard objects on which manifold learning algorithms seem to be tested, but it also adds a manifold with boundary to our list. For the sampling, we used the function `make_swiss_roll` from the library *scikit-learn*.

- **Schwarz P surface:** This is the triply periodic surface in \mathbb{R}^3 with equation $\cos(x) + \cos(y) + \cos(z) = 0$ —it approximates one of Schwartz's triply periodic minimal surfaces and thus shows features of negative curvature, both intrinsic and extrinsic. To be able to deal with a finite volume surface we consider it as a submanifold of the 3-dimension flat torus.
- **Spheres:** This is the standard sphere $S^d = \{x \in \mathbb{R}^{d+1} \text{ with } \|x\| = 1\}$. Spheres have reach 1 and are positively curved.
- **Gaussian distribution:** The standard Gaussian distribution in \mathbb{R}^d . The reason why we test this in particular is to include a non-uniform distribution in our list.

The way we sample points depends on the concrete manifold under consideration, but we stress that we are sampling each point independently. More precisely, we are not aiming at getting point at some uniform distance of each other.




Figure 3.2 – Examples of samplings from manifolds. 500 points on a "Worms" manifold (Left) and 2 000 points on a piece of the Schwarz surface (Right).

These results tend to confirm the conclusions of the heuristic model. We can test this model a bit more by testing the estimator with the number of points given by Table 3.4. This is what we did in Table 3.6, and we still obtain results close to what was expected.

Reach free estimator

A problem with all the results we have been discussing so far is that in practice one has little clue what the reach of the underlying manifold could be. However, as we already

Manifold	d	90% target	70% target
Worms	1	88%	66%
Rotation torus	2	92%	70%
Clifford torus	2	89%	69%
Flat torus	2	88%	66%
Swiss roll	2	93%	69%
Schwarz surface	2	88%	66 %
3-sphere	3	92%	76%
4-sphere	4	89%	75%
Product of two rotation tori	4	92%	70%
Clifford torus	4	93%	72%
Flat torus	4	90%	74%
Product of two Schwarz surfaces	4	92%	72%
Gaussian distribution in \mathbb{R}^4	4	90%	76%
5-sphere	5	93%	74%

Table 3.5 – Experimental rates of success for different manifolds by considering the amount of pairs of points suggested by Table 3.3.

Manifold	d	rate of success
Clifford torus	2	91%
3-sphere	3	91%
Flat torus	4	91%
Product of two tori	4	94%

Table 3.6 – Experimental rates of success for different manifolds by sampling the number of points given by the "heuristic" column in Table 3.4.

mentioned in the introduction, one can actually derive from Table 3.3 an estimator which does not need any a priori bound on the reach.

Assumption: We have a data set $X \subset \mathbb{R}^s$ of which we think that it has been sampled from some mysterious submanifold $M \subset \mathbb{R}^s$. We trust however that our data set is good enough and we want to test if M could plausibly have dimension d .

Test: For the chosen d , let ϵ_1, ϵ_2 and N be as in Table 3.3 (say from the 90% column). Now take $R > 0$ to be minimal with $|PX(R)| \geq N$ and set $r = \frac{\epsilon_2}{\epsilon_1} R$. Now check if $\dim_{\text{Corr}(R,r)}(X) = d$.

The heuristic discussion above, as well as the numerical experiments, suggest that as long as our data set is rich enough so that $R \leq 0.54 \cdot \tau(M)$ then we should get $\dim_{\text{Corr}(R,r)}(X) = d$ with a 90% probability.

Wondering what would happen if we run this algorithm on more "real" data sets, we chose 3 different data sets of respective dimension 1, 2 and 3 (see [Gri22]). Each data sets consist in 200 grayscale pictures of the 3D-model *Suzanne* with random rotations (see Figure 3.1). We respectively randomized 1, 2 and 3 Euler angles to obtain the desired dimensions. The pictures are 64 by 64 pixels large and can thus be represented as points in \mathbb{R}^{4096} . The results are presented in Table 3.7. For each data set, we test dimension 1, 2, 3 and 4 and we show the estimated dimension (before rounding it to the closest integer).

Note that when using the parameters of dimension 1 for estimating the dimension of a higher dimensional set, we usually get no points at distance $\varepsilon/2$, as $(1.5/0.19)^2$ is higher than 30. When this happens, the estimator (3.1) cannot be computed, we can only conclude that the dimension is probably bigger than 1.

Hypothesis	Data set 1	Data set 2	Data set 3
dimension = 1	1.13	> 1	> 1
dimension = 2	1.59	2.02	3.54
dimension = 3	1.64	2.03	3.33
dimension = 4	1.45	2.08	3.40

Table 3.7 – Testing the dimension of the 3 "real" data sets. To test dimension=1 (resp. dimension=2, resp. dimension=3) we set ϵ_1 so that $|PX(\epsilon_1)| = 30$ (resp. 122, resp. 249). Cells with a pink background represent the tests in which the estimated dimension is consistent with the tested dimension. The cells with the value "> d " mean that there were no pairs at distance ε_2 for the corresponding number of pairs at distance ε_1 .

We can see in Table 3.7 that there were no Type I errors (meaning that we never rejected the true hypothesis). On the other hand we found a Type II error, when the data set of dimension 1 passed the test for dimension 2.

3.6 Comparison with other estimators

In this paper we study the estimator (3.1), but a variety of other algorithms exist. For example, instead of doing statistics using only the distances between points, we could also use the angles between points, or other more complex features. These different approaches can be compared using the previous heuristic model. We refer to [Cam03; CS16] and specially to [LV07, Chapter 3] for a review of different dimension estimators. We will compare (3.1) with the estimators ANOVA, local PCA, and with the implementation of

(3.1) where one tries to avoid picking scales, reading instead the dimension from a log-log chart.

ANOVA

Diaz, Quiroz and Velasco propose in [DQV19b] a method based on the angles referred as ANOVA in the literature. Their idea is to estimate the local dimension around a point x by considering its k nearest neighbors and the $\binom{k}{2}$ angles at x formed by these points. From the variance of these $\binom{k}{2}$ angles, we can identify the closest β_d and deduce the dimension d , where β_d is defined as follows.

$$\beta_d = \frac{1}{\text{vol}(S^{d-1})^2} \int_{S^{d-1} \times S^{d-1}} \left(\angle(\theta, \eta) - \frac{\pi}{2} \right)^2 d\theta d\eta$$

The global dimension can then be recovered by taking the median, the mode or the mean of the local dimensions.

To be able to compare ANOVA to the estimator (3.1) we will instead take the variance of all angles. More precisely, consider every ordered triple of points in which all three points are at distance at most ε_1 from each other. Each such triple (x, y, z) determines an angle $\angle(y - x, z - x)$. We compute the variance of the so-obtained angles we can locate the closest β_d . We take that d to be the ANOVA dimension of our data set.

Putting ourselves again in the ideal situation that we are working in a scale at which curvature can be ignored and assuming (and that is a lot of assuming in this case) that the angles we find are independent of each other, we can argue as earlier in the discussion of the heuristic bound and we get that, in dimension 4, we would need at least 652 angles to achieve a confidence level of 90% with the ANOVA estimator.

Suppose now that we sample 1958 points from the 4-dimensional Clifford torus $M = \mathbb{T}^4$. That number was chosen so that we expect to have 516 pairs within $\varepsilon_1 = 0.54$ of each other, the heuristic bound for (3.1) in dimension 4. On the other hand we expect to have

$$\text{number of unordered triples} = \binom{1958}{3} \cdot \left(\frac{\text{vol}(B^{\mathbb{R}^4}(\varepsilon_1))}{\text{vol}(M)} \right)^2 \simeq 91.$$

Each unordered pair gives 3 angles, meaning can expect to find about 273 angles. This is much less than what we estimated that would be needed.

And this phenomenon gets worse when the volume of the underlying manifold grows:

if the manifold has volume $64\pi^4$ and if we have a data set with 516 pairs of points within $\epsilon_1 = 0.54$ then we expect to only have 45 triples, that is about 135 angles. The reason for this is that the number of points needed to have a given number of pairs grows with $\text{vol}(M)^{\frac{1}{2}}$ while it grows as $\text{vol}(M)^{\frac{2}{3}}$ when we fix the number of triples instead. This means that for large volumes the algorithm will only become worse as finding triples of points will become more and more difficult.

We numerically compared the rates of success of estimator (3.1) and of ANOVA for the Clifford torus with the same scales. The results are presented in Table 3.8. This experiment shows that, for this example and for these scales, estimator (3.1) gives significantly better results than ANOVA. However, we cannot conclude that estimator (3.1) has better performance in general. In particular, for manifolds of smaller volume or for different choices of scales, ANOVA could in principle give better results.

Manifold	d	number of points	estimator (3.1)	ANOVA
Clifford torus	2	76	93%	65%
Clifford torus	3	347	93%	67%

Table 3.8 – Comparison of the experimental rates of success between estimator (3.1) and ANOVA on the Clifford torus.

Remark 5. *Recall that we have considered a variation of ANOVA—our conclusions should also apply to the original algorithm as well.*

Local PCA

Principal Component Analysis aims to find the best linear space containing a given data set. According to [Bre+18b], PCA is the gold standard of dimension estimation. It works as follows. To our given data set $X = (x_1, \dots, x_n) \subset \mathbb{R}^N$ we associate first the mean

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

and then the $n \times N$ matrix A whose rows are the vectors $u_i - \bar{u}$, and one computes then the singular values $s_1 \geq s_2 \geq \dots \geq s_{\min\{n, N\}}$. If X is contained in a linear subspace of dimension d then $s_k = 0$ for all $k \geq d + 1$. Accordingly, one can declare that X has PCA-dimension k if the gap $s_k - s_{k+1}$ is maximal. Another possibility would be to fix

a threshold ϵ and declare the PCA-dimension of X to be the largest k with $s_k \geq \epsilon$. For example, if what one wants to do is to test the hypothesis that X is contained in a d -dimensional subspace then one can check for example if s_{d+1} is below some threshold $\epsilon < \frac{1}{d+2}$ —the number $\frac{1}{d+2}$ arises because it is the expected value for s_d if X is uniformly sampled out of the unit ball in \mathbb{R}^d .

In any case, that was PCA. The idea of *local PCA*, or *Nonlinear PCA*, is to apply PCA to certain subsets of the data set and then, for good measure, average the so obtained numbers. As we see, one does not only need to agree on what one calls PCA, but also on what subsets does one wants to subject to the PCA treatment. For example, as in [Bre+18b] one can cluster the data set using single linkage clustering¹ and then apply PCA to each cluster. This is, in our humble and uneducated opinion, a very reasonable choice if, for data sets X uniformly sampled out of a submanifold $M \subset \mathbb{R}^N$, what one wants to do is to find linear spaces (approximately) containing each connected component of M . On the other hand, if what one wants to do is to recover the dimension of M then it seems reasonable to rather apply PCA to (some of) the sets $X \cap \mathbb{B}(x, \epsilon)$ for $x \in X$ and for some ϵ chosen so that $M \cap \mathbb{B}(x, \epsilon)$ is well-approximated by its tangent space.

Note now that every set consisting of $d + 1$ points is contained in a d -dimensional affine subspace of \mathbb{R}^s . This means that if we want to distinguish dimension d from dimension $d + 1$ we need at the very least to find $d + 2$ tuples of nearby points. Now, if we once again sample 1958 points out of the Clifford torus $M = \mathbb{T}^4$ so that we can expect the magic number of 516 pairs of points at most at distance 0.54 from each other, then we can expect to find

$$\begin{array}{l} \text{number of unordered} \\ \text{6-tuples at scale 0.54} \end{array} = \binom{1958}{6} \cdot \left(\frac{\text{vol}(B^{\mathbb{R}^4}(0.54))}{\text{vol}(M)} \right)^5 \simeq 0.110373....$$

In other words, the expectation is to not have any 6-tuples, meaning that, if we are working at the same scales as we were implementing (3.1), then we have no chance of distinguishing our 4-manifold from a 5-manifold using local PCA. Note that even if we work at a scale of 2, scale at which we are acting as if the round sphere were totally flat, then we can expect

1. In single linkage clustering the clusters are, for some ϵ , the connected components of the graph with vertex set X and where two vertices are joined by an edge if they are within ϵ of each other.

to only have 2 6-tuples

$$\begin{array}{l} \text{number of unordered} \\ \text{6-tuples at scale 2} \end{array} = \binom{1958}{6} \cdot \left(\frac{\text{vol}(B^{\mathbb{R}^4}(2))}{\text{vol}(M)} \right)^5 \simeq 2.043944....$$

As we see, even if we work at scales at which we are flat earthers, we do not have by far enough sufficiently populated clusters for local PCA to be meaningful.

log-log plots

Another method to derive the estimator (3.1) when the reach is unknown, is to take a lot of scales

$$\epsilon_1 < \epsilon_2 < \dots < \epsilon_k,$$

and to consider the graph of the piecewise linear function with corners

$$(\log(\epsilon_i), \log(|PX(\epsilon_i)|))$$

and try to read out some sort of meaningful slope of that function (see [LV07, Chapter 3]). We implemented this procedure for 1 000 points uniformly sampled from the product of two 2-dimensional rotation tori of reach 1. We obtained the graph presented on Figure 3.3.

Three different parts can be observed on this graph:

1. A flat part, when the scale is smaller than the minimal distance between the points (represented with the value -1 on the y-axis).
2. A mostly linear part whose slope should approximate the dimension of M .
3. A flat part when the scale becomes greater than the diameter of M (which plateaus at $\log(1000 \cdot 999/2) \simeq 13.12$).

It is notable that in this example the second part looks linear even after twice the reach, even if the behavior of $\log(|PX(\epsilon_i)|)$ beyond this scale is unpredictable. Around the middle of this graph, the slope of the linear part seems to be close to 4, as predicted.

We also applied this method from the 3 "real" datasets of pictures of *Suzanne* with random rotations. The results are presented in Figure 3.4. For dimension 1 and 2, it worked surprisingly well and the log-log plots clearly show what is the right dimension. The plot for the data set of dimension 3 is not as clear, but the result seems consistent.

Let us conclude with two further comments on the log-log procedure:







Figure 3.3 – Log-log plot for a set of 1000 points sampled out of the product of two rotation tori. The graph also features three lines of slope 3, 4 and 5, for comparison.

(1) The log-log procedure seems to allow us to estimate the dimension of M without any assumption on its reach. It is however difficult to define a precise algorithm for it. Consequently, we cannot really compare its performance with that of the other estimators.

(2) Additionally, this method gives us insight on the choice of scales that could be relevant. Indeed, the scales determined in Section 3.4.4 have no reason to be optimal for the heuristic model. And as this model does not take the curvature nor the reach into account, it is impossible to use it to find optimal scales. When considering Figure 3.3, it seems natural to choose two scales that are far apart to measure the slope as precisely as possible, but close enough so that they do not get too close to the extremities. Here, two scales that would seem relevant could be $\exp(0.5) \simeq 1.6$ and $\exp(-0.5) \simeq 0.6$. In any case, the choice of scales $0.54 \simeq \exp(-0.6)$ and $0.23 \simeq \exp(-1.5)$ seem to be far from optimal. Using the estimator (3.1) with scales 1.6 and 0.6 on computational examples, we obtain a rate of success of 99% with only 1000 points for the product of two tori. For the scales 0.54 and 0.23, the heuristic model gives us a number of points of 3916 for a rate of success of 90%. This rate falls to 41% with only 1000 points on computational examples. However, these scales have been chosen after the study of the results to obtain the desired answer on a given manifold. But this does not give us a method to choose optimal scales a priori. For

example, the scales 1.6 and 0.6 one give a rate of success of 45% for 100 points on the 4-sphere, while the scale 1 and 0.6 give a rate of 92%.



PART III

3D printing

During my PhD, I experimented some methods of visualization of the objects I am working on. In particular, I tried to find the best way to represent the Alexander horned sphere and the Fox-Artin sphere so that they can be understood easier by the audience during my talks.

The best method I found is to print 3D models of these objects so that anyone can look at them from various angles and visualize them better. I created 3D model the Alexander horned sphere and the Fox-Artin sphere with a 3D modeling software and printed them with the 3D printer available in the University, with the help of Rémi Coulon.

This part features a short paper written for **Bridges Aalto 2022**, a conference linking art an mathematics. This paper presents the two 3D models I developed during my thesis.

3D Printed Models of Wild Spheres

Lucien Grillet

University of Rennes 1, France ; lucien.grillet@univ-rennes1.fr

Abstract

We discuss two wild embeddings of the sphere in the Euclidean space: the Alexander horned sphere and the Fox-Artin sphere. After explaining the mathematical properties carried by these two embeddings, we present 3D printed models of these objects. The models have been designed to be decorative table-top sculptures while being simple enough to serve as supports for mathematical discussions.

1 The Alexander horned sphere

In topology, the famous Jordan-Brouwer separation theorem states that if one draws a closed line on the plane which does not intersect itself, this line separates the plane into two disjoint regions.

Theorem (Jordan-Brouwer separation theorem). *Let C be a continuous simple closed curve in \mathbb{R}^2 . Then $\mathbb{R}^2 \setminus C$ has exactly two connected components.*

This theorem can seem a bit weak: we know that the line separates the plane into two regions but we do not know what the topology of these two regions are. In fact, the Jordan-Schoenflies theorem tells us that such a line can be continuously transformed to a standard circle.

Theorem (Jordan-Schoenflies theorem). *Let C be a continuous simple closed curve in \mathbb{R}^2 . There is an homeomorphism of \mathbb{R}^2 into itself that sends C on the standard circle $\{x \in \mathbb{R}^2 \mid \|x\|_2 = 1\}$.*

This theorem is stronger than the Jordan-Brouwer separation theorem because it completely describes the situation. In particular it shows that the region bounded by the line must be homeomorphic to an open disk.

Going one dimension higher, one can wonder if the situation is similar for embeddings of the sphere S^2 in the Euclidean space \mathbb{R}^3 . A first positive result is the Alexander duality, which can be thought of as a generalization of the Jordan-Brouwer separation theorem, it can be used to show that any continuous injective map from the sphere S^{n-1} into the sphere \mathbb{R}^n cuts the latter into two pieces.

However, the Jordan-Schoenflies theorem does not generalize to higher dimensions. *The Alexander horned sphere* (Represented in Figure 1) is a subspace of \mathbb{R}^3 homeomorphic to a sphere but whose exterior (together with the point at infinity) is not homeomorphic to an open ball. Indeed, its exterior is not even simply connected (the curve γ on Figure 1 cannot be shrunk to a point).




Figure 1: The Alexander horned sphere

This sphere was first described by J. W. Alexander in [1] and is an example of what we call a *wild embedding*: it is a continuous embedding such that there is no homeomorphism of \mathbb{R}^3 into itself that sends it to a smooth embedding. In fact, in this representation, the sphere is smooth everywhere, except on a Cantor set around which the sphere makes an infinite number of knots.

To construct this object, start with a sphere and grow two horns out of it. Divide each of these horns into two other horns and interlace one horn of the first division with one horn of second division. Repeat this process with each new pair of horns over and over infinitely many times with smaller and smaller horns until they accumulate on a Cantor set.

The resulting object is homeomorphic to a sphere, but if one defines a loop around one of the first horns, the infinite number of interlacing makes it impossible to shrink this loop to a point.

To make this object easier to understand, I decided to recreate it in the real world. Thanks to a 3D modeling software, I designed a model of the Alexander horned sphere. I chose a smooth embedding (except on the singular Cantor set) for the visual aspect and I chose to represent four levels of the pattern, with a scaling factor of 43% between each level, to obtain a model with as many levels as possible while being printable with the precision and the size of the 3D printer. The result is the 20 centimeters object shown in Figure 2. The main difficulty in the printing process was the fact that the printer needed to print a support together with the sphere, because the PLA filament always needs to be deposited on a solid structure. I chose to print this support with a soluble PVA filament to obtain a smooth result without any trace of this support.




Figure 2: 3D model and 3D printed sculpture of the Alexander horned sphere

2 The Fox-Artin sphere

The Alexander horned sphere is well-known, but a wide diversity of other wild spheres exist without having the same reputation. A reason for this is the lack of visualization of these spheres. Among all of them, there is one that I find particularly interesting: *the Fox-Artin sphere* (represented on Figure 3). I first encountered this sphere when confronted to the following question during my research work.

Question. Let K be a compact subset of \mathbb{R}^3 , and suppose that K verifies the two following properties.

- The interior $\overset{\circ}{K}$ of K is homeomorphic to an open 3-ball
- The boundary ∂K of K is homeomorphic to a 2-sphere

Is K homeomorphic to a closed ball ?

The Fox-Artin sphere offers a negative answer to this question: it is a subset of S^3 homeomorphic to a 2-sphere and it cuts S^3 into two components homeomorphic to open 3-balls. However, the union of the exterior component (according to Figure 3) with this 2-sphere does not yield a closed ball.

Returning to the situation in dimension 2, the Jordan-Schoenflies theorem told us two things:

- (A) The interior of a simple closed curve is homeomorphic to an open disk.
- (B) The union of a simple closed curve with its interior is homeomorphic to a closed disk.

In other words, it says that the homeomorphism from the open disk to the interior of C can be chosen so that it extends homeomorphically to C .

The Alexander horned sphere contradicts the point A (and thus the point B) while the Fox-Artin sphere only contradicts the point B. The answer to our question is thus negative: it is strangely possible to "glue" a 2-sphere on the boundary of an open 3-ball without obtaining a closed 3-ball.

However, as in the Alexander horned sphere, the interior component together with the sphere yield a subset homeomorphic to a closed ball.




Figure 3: *The Fox-Artin sphere*

This sphere is defined as a thickening of a wild arc, with a thickness converging to zero as we approach the endpoint. The arc used here is an infinite concatenation of loops where each loop is knotted with the previous one. This arc has been defined by R. H. Fox and E. Artin in [2]. In fact, the term "Fox-Artin sphere" can refer to slightly different objects, as several wild arcs have been defined in the aforementioned article. Here, we only consider the simplest of these arcs, which is infinitely knotted only on one side, but we could construct other wild spheres with other properties by working with other arcs.

Compared to the Alexander horned sphere, this sphere is usually easier to comprehend for people because it has a single singular point. It is clearer that it is homeomorphic to a sphere, but its properties are more difficult to understand and to prove.

I made a 15 centimeters sculpture of this Fox-Artin sphere with the same method used for the Alexander horned sphere. The model is smooth everywhere except at its singular endpoint. To reduce the height of the sculpture and to make a base for the sculpture to stand on, I decided to flatten the original sphere from which the horn is grown. With the precision of the 3D printer, I could print seven levels of the pattern with a scaling factor of 67% between each level. Figure 4 represents the 3D model and the final result.

The files used in this article can be freely accessed on <http://lugri.net/maths/wildspheres.php>.




Figure 4: 3D model and 3D printed sculpture of the Fox-Artin sphere

Acknowledgements

Thanks to Rémi Coulon who helped me with the technical difficulties involved in the process of printing such knotted objects.

References

- [1] J. W. Alexander. “An Example of a Simply Connected Surface Bounding a Region which is not Simply Connected.” *Proceedings of the National Academy of Sciences of the United States of America*, 1924.
- [2] R. H. Fox and Emil Artin. “Some wild cells and spheres in three-dimensional space.” *Annals of Mathematics, Second Series*, 1948.

BIBLIOGRAPHY

- [Aam+19] Eddie Aamari et al., « Estimating the reach of a manifold », *in: Electronic journal of statistics* 13.1 (2019), pp. 1359–1399.
- [Ale23] James W Alexander, « On the deformation of an n cell », *in: Proceedings of the National Academy of Sciences* 9.12 (1923), pp. 406–407.
- [Ale24] J. W. Alexander, « An example of a simple connected surface bounding a region which is not simply connected. », English, *in: Proc. Natl. Acad. Sci. USA* 10 (1924), pp. 8–10, ISSN: 0027-8424, DOI: 10.1073/pnas.10.1.8.
- [Ber03] Marcel Berger, *A panoramic view of Riemannian geometry*, Springer-Verlag, Berlin, 2003, pp. xxiv+824, ISBN: 3-540-65317-1, DOI: 10.1007/978-3-642-18245-7, URL: <https://doi.org/10.1007/978-3-642-18245-7>.
- [BH13] Martin R Bridson and André Haefliger, *Metric spaces of non-positive curvature*, vol. 319, Springer Science & Business Media, 2013.
- [Bin52] R. H. Bing, « A homeomorphism between the 3-sphere and the sum of two solid horned spheres », *in: Ann. of Math. (2)* 56 (1952), pp. 354–362, ISSN: 0003-486X, DOI: 10.2307/1969804, URL: <https://doi.org/10.2307/1969804>.
- [Bin59] R. H. Bing, « Conditions under which a surface in E^3 is tame », *in: Fund. Math.* 47 (1959), pp. 105–139, ISSN: 0016-2736, DOI: 10.4064/fm-47-1-105-139, URL: <https://doi.org/10.4064/fm-47-1-105-139>.
- [BLP05] Michel Boileau, Bernhard Leeb, and Joan Porti, « Geometrization of 3-dimensional orbifolds », *in: Ann. of Math. (2)* 162.1 (2005), pp. 195–290, ISSN: 0003-486X, DOI: 10.4007/annals.2005.162.195, URL: <https://doi.org/10.4007/annals.2005.162.195>.
- [BLW19] Jean-Daniel Boissonnat, André Lieutier, and Mathijs Wintraecken, « The reach, metric distortion, geodesic convexity and the variation of tangent spaces », *in: Journal of applied and computational topology* 3.1 (2019), pp. 29–58.

-
- [BNS06] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani, « Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. », *in: Journal of machine learning research* 7.11 (2006).
- [Bre+18a] Paul Breiding et al., « Learning algebraic varieties from samples », English, *in: Rev. Mat. Complut.* 31.3 (2018), pp. 545–593, ISSN: 1139-1138, DOI: 10.1007/s13163-018-0273-6.
- [Bre+18b] Paul Breiding et al., « Learning algebraic varieties from samples », *in: Revista Matemática Complutense* 31.3 (2018), pp. 545–593.
- [Bro60] Morton Brown, « A proof of the generalized Schoenflies theorem », *in: Bull. Amer. Math. Soc.* 66 (1960), pp. 74–76, ISSN: 0002-9904, DOI: 10.1090/S0002-9904-1960-10400-4, URL: <https://doi.org/10.1090/S0002-9904-1960-10400-4>.
- [Bro62] Morton Brown, « Locally flat imbeddings of topological manifolds », *in: Ann. of Math. (2)* 75 (1962), pp. 331–341, ISSN: 0003-486X, DOI: 10.2307/1970177, URL: <https://doi.org/10.2307/1970177>.
- [Cam03] Francesco Camastra, « Data dimensionality estimation methods: a survey », *in: Pattern recognition* 36.12 (2003), pp. 2945–2954.
- [CB08] Craig Calcaterra and Axel Boldt, « Lipschitz flow-box theorem », *in: J. Math. Anal. Appl.* 338.2 (2008), pp. 1108–1115, ISSN: 0022-247X, DOI: 10.1016/j.jmaa.2007.06.001, URL: <https://doi.org/10.1016/j.jmaa.2007.06.001>.
- [CEE75] Jeff Cheeger, David G Ebin, and David Gregory Ebin, *Comparison theorems in Riemannian geometry*, vol. 9, North-Holland Amsterdam, 1975.
- [CS16] Francesco Camastra and Antonino Staiano, « Intrinsic dimension estimation: Advances and open problems », *in: Information Sciences* 328 (2016), pp. 26–41.
- [Dav86] Robert J. Daverman, *Decompositions of manifolds*, vol. 124, Pure and Applied Mathematics, Academic Press, Inc., Orlando, FL, 1986, pp. xii+317, ISBN: 0-12-204220-4.

-
- [DQV19a] Mateo Díaz, Adolfo J. Quiroz, and Mauricio Velasco, « Local angles and dimension estimation from data on manifolds », English, *in: J. Multivariate Anal.* 173 (2019), pp. 229–247, ISSN: 0047-259X, DOI: 10.1016/j.jmva.2019.02.014.
- [DQV19b] Mateo Diaz, Adolfo J Quiroz, and Mauricio Velasco, « Local angles and dimension estimation from data on manifolds », *in: Journal of Multivariate Analysis* 173 (2019), pp. 229–247.
- [Edw80] Robert D. Edwards, « The topology of manifolds and cell-like maps », *in: Proceedings of the International Congress of Mathematicians (Helsinki, 1978)*, Acad. Sci. Fennica, Helsinki, 1980, pp. 111–127.
- [Eil49] Samuel Eilenberg, « On the problems of topology », *in: Ann. of Math. (2)* 50 (1949), pp. 247–260, ISSN: 0003-486X, DOI: 10.2307/1969448, URL: <https://doi.org/10.2307/1969448>.
- [ER92] J-P Eckmann and David Ruelle, « Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems », *in: Physica D: Nonlinear Phenomena* 56.2-3 (1992), pp. 185–187.
- [FA48] Ralph H. Fox and Emil Artin, « Some wild cells and spheres in three-dimensional space », *in: Ann. of Math. (2)* 49 (1948), pp. 979–990, ISSN: 0003-486X, DOI: 10.2307/1969408, URL: <https://doi.org/10.2307/1969408>.
- [Fed59] Herbert Federer, « Curvature measures », *in: Transactions of the American Mathematical Society* 93.3 (1959), pp. 418–491.
- [FMN16] Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan, « Testing the manifold hypothesis », *in: Journal of the American Mathematical Society* 29.4 (2016), pp. 983–1049.
- [Fre13] Michael Freedman, *Bing Topology and Casson Handles*, <https://www.math.uni-bielefeld.de/~sbehrens/files/Freedman2013.pdf>, 2013.
- [FS22] Michael Freedman and Michael Starbird, *The Geometry of the Bing Involution*, 2022, arXiv: 2209.07597 [math.GT].
- [GP04] Peter Grassberger and Itamar Procaccia, « Measuring the strangeness of strange attractors », *in: The theory of chaotic attractors*, Springer, 2004, pp. 170–189.

-
- [Gra86] Peter Grassberger, « Do climatic attractors exist? », *in: Nature* 323.6089 (1986), pp. 609–612.
- [Gri22] Lucien Grillet, *The associated program*, <https://github.com/lgrillet/dim-estimation>, 2022.
- [Ham08] D. H. Hamilton, « QC Riemann mapping theorem in space », *in: Complex analysis and dynamical systems III*, vol. 455, Contemp. Math. Amer. Math. Soc., Providence, RI, 2008, pp. 131–149, DOI: 10.1090/conm/455/08851, URL: <https://doi.org/10.1090/conm/455/08851>.
- [Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002, pp. xii+544, ISBN: 0-521-79160-X; 0-521-79540-0.
- [Jor87] Camille Jordan, *Cours d'analyse de l'École polytechnique. Tome III*, 1887.
- [KL88] Sławomir Kwasik and Kyung Bai Lee, « Locally linear actions on 3-manifolds », *in: Math. Proc. Cambridge Philos. Soc.* 104.2 (1988), pp. 253–260, ISSN: 0305-0041, DOI: 10.1017/S0305004100065427, URL: <https://doi.org/10.1017/S0305004100065427>.
- [Lem18] Alexander Lemmens, *A local Jordan-Brouwer separation theorem*, 2018, arXiv: 1810.13221 [math.AT].
- [LV07] John A Lee and Michel Verleysen, *Nonlinear dimensionality reduction*, vol. 1, Springer, 2007.
- [Mes77] Robert Messer, « Three dimensional manifolds with finitely generated fundamental groups », English, *in: Trans. Am. Math. Soc.* 226 (1977), pp. 119–145, ISSN: 0002-9947, DOI: 10.2307/1997944.
- [Mun60] James Munkres, « Obstructions to the smoothing of piecewise-differentiable homeomorphisms », *in: Ann. of Math. (2)* 72 (1960), pp. 521–554, ISSN: 0003-486X, DOI: 10.2307/1970228, URL: <https://doi.org/10.2307/1970228>.
- [MZ54] Deane Montgomery and Leo Zippin, « Examples of transformation groups », *in: Proc. Amer. Math. Soc.* 5 (1954), pp. 460–465, ISSN: 0002-9939, DOI: 10.2307/2031959, URL: <https://doi.org/10.2307/2031959>.
- [NM10] Hariharan Narayanan and Sanjoy Mitter, « Sample complexity of testing the manifold hypothesis », *in: Advances in neural information processing systems* 23 (2010).

-
- [NSW08a] Partha Niyogi, Stephen Smale, and Shmuel Weinberger, « Finding the homology of submanifolds with high confidence from random samples », English, *in: Discrete Comput. Geom.* 39.1-3 (2008), pp. 419–441, ISSN: 0179-5376, DOI: 10.1007/s00454-008-9053-2.
- [NSW08b] Partha Niyogi, Stephen Smale, and Shmuel Weinberger, « Finding the homology of submanifolds with high confidence from random samples », *in: Discrete & Computational Geometry* 39.1 (2008), pp. 419–441.
- [OP19] Jani Onninen and Pekka Pankka, *Bing meets Sobolev*, 2019, arXiv: 1908.09193 [math.MG].
- [Par20] John Pardon, *Smoothing finite group actions on three-manifolds*, 2020, arXiv: 1901.11127 [math.GT].
- [Pro88] Itamar Procaccia, « Complex or just complicated? », *in: Nature* 333.6173 (June 1988), pp. 498–499, ISSN: 1476-4687, DOI: 10.1038/333498a0, URL: <https://doi.org/10.1038/333498a0>.
- [Sch06] A. Schoenflies, « Beiträge zur Theorie der Punktmengen. III », *in: Math. Ann.* 62.2 (1906), pp. 286–328, ISSN: 0025-5831, DOI: 10.1007/BF01449982, URL: <https://doi.org/10.1007/BF01449982>.
- [Sim98] Dominique Simpelaere, « Correlation dimension », *in: Journal of statistical physics* 90.1 (1998), pp. 491–509.
- [Smi39] P. A. Smith, « Transformations of finite period. II », *in: Ann. of Math.* (2) 40 (1939), pp. 690–711, ISSN: 0003-486X, DOI: 10.2307/1968950, URL: <https://doi.org/10.2307/1968950>.
- [Tak83] Floris Takens, « Invariants related to dimension and entropy », *in: Atas do* 13 (1983), pp. 353–359.
- [Tak85] Floris Takens, « On the numerical determination of the dimension of an attractor », *in: Dynamical systems and bifurcations*, Springer, 1985, pp. 99–106.
- [The90] James Theiler, « Estimating fractal dimension », *in: JOSA A* 7.6 (1990), pp. 1055–1073.
- [Tuc74] Thomas W. Tucker, « Non-compact 3-manifolds and the missing-boundary problem », English, *in: Topology* 13 (1974), pp. 267–273, ISSN: 0040-9383, DOI: 10.1016/0040-9383(74)90019-6.

-
- [Wei14] Shmuel Weinberger, « The complexity of some topological inference problems », *in: Foundations of Computational Mathematics* 14.6 (2014), pp. 1277–1285.

Titre : La Conjecture de Smith en faible régularité

Mot clés : Topologie et géométrie en petites dimensions, 3-variétés, homéomorphismes sauvages, conjecture de Smith

Résumé : En 1939, Paul Althaus Smith démontra que l'ensemble des points fixes d'une application continue d'ordre fini de la 3-sphère dans elle-même était homéomorphe à une sphère de dimension inférieure. Ses résultats ne renseignent cependant pas sur la nature du plongement de cet ensemble de points fixes. En 1952, R. H. Bing donna un exemple d'une involution continue de la 3-sphère dont l'ensemble des points fixes est homéomorphe à une 2-sphère plongée de manière "sauvage". Suite aux travaux de nombreux mathématiciens tels que John Morgan, Hyman Bass, William Thurston et Grigori Perelman, nous savons aujourd'hui que, s'il s'agit d'une ap-

plication lisse, une telle application d'ordre fini est nécessairement conjuguée à une isométrie. Dans une série de conférences données en 2013 à Santa Barbara, Michael Freedman conjectura que cette dernière affirmation devrait également être vérifiée pour des applications de régularité intermédiaire telles que des applications lipschitziennes. Nous démontrons qu'une application lipschitzienne d'ordre fini d'une 3-variété et de constante de Lipschitz proche de 1 est nécessairement conjuguée à une application lisse, répondant partiellement à la question de Michael Freedman.

Title: The Smith conjecture in low regularity

Keywords: Low-dimensional geometry and topology, 3-manifolds, wild homeomorphisms, Smith conjecture

Abstract: In 1939, Paul Althaus Smith proved that the fixed set of a continuous self-map of finite order of the 3-sphere was homeomorphic to a lower dimensional sphere. However, his results do not give any information about the nature of the embedding of this fixed set. In 1952, R. H. Bing gave an example of a continuous involution of the 3-sphere whose fixed set is homeomorphic to a "wild" embedding of the 2-sphere. Following the work of many mathematicians such as John Morgan, Hyman Bass, William Thurston, and Grigori Perelman, we

know that, if such a finite order map is smooth, it is necessarily conjugate to an isometry. In a series of lectures given in 2013 in Santa Barbara, Michael Freedman conjectured that this property should also be verified for maps of intermediate regularity such as Lipschitz maps. We show that a Lipschitz map of finite order of a 3-manifold and of Lipschitz constant close to 1 is necessarily conjugate to a smooth map, partially answering Michael Freedman's question.