Université Pierre et Marie Curie

Universidad de la República

École Doctorale Paris Centre et PEDECIBA

Thèse de doctorat

Discipline : Mathématiques

présentée par

Pablo Lessa

Brownian motion on stationary random manifolds

---

dirigée par François Ledrappier et Matilde Martínez

Soutenue le 18 mars 2014 devant le jury composé de :

M. Jesús Álvarez	U. de Santiago de Compostela	rapporteur
M. Yves Coudene	Université de Bretagne Occidentale	examinateur
M. Gilles Courtois	Université Paris 6	examinateur
M. Vadim Kaimanovich	Université d'Ottawa	rapporteur
M. Raphaël Krikorian	Université Paris 6	examinateur
M. François Ledrappier	Université Paris 6	directeur
Mme. Matilde Martínez	Universidad de la República	directeur

Institut de Mathématiques de Jussieu
175, rue du chevaleret
75 013 Paris

École doctorale Paris centre Case 188
4 place Jussieu
75 252 Paris cedex 05

Centro de Matemática
Facultad de Ciencias
Iguá 4225
11400 Montevideo

Dedicada a Marcela por los mejores años de mi vida.

Serendipity: The occurrence and development of events by chance in a happy or beneficial way.

Acknowledgements

François Ledrappier and Matilde Martinez have been the perfect advisors, simultaneously challenging and enabling me to do better mathematics, I thank them deeply.

I also owe a lot to Martín Sambarino who among other things introduced me to François and arranged for me to be coadvised by Matilde.

Fernando Alcalde invited me to give my first ever talk and also handed me a copy of Benjamini and Curien's work on stationary random graphs. It turned out to be a key piece for this work and I would like to thank him especially for bringing it to my attention.

I learned many interesting things thanks to conversations with Bertrand Deroin, Christian Bonatti, Sébastien Álvarez, Liviu Nicolaescu, Renato Bettiol, Karsten Grove, Miguel Paternain and Rafael Potrie.

I would like to thank Jairo Bochi and Samuel Senti for inviting me to speak at Seminario Edaí, Barbara Shapira for inviting me to the workshop on non-positive curvature at Luminy, and Alberto Verjovski for several interesting conversations.

I have been lucky to have had ideal rapporteurs in Jesús Álvarez Lopez and Vadim Kaimanovich, I thank them for their comments and corrections.

Funding for my travels was provided by IFUM and CSIC, I also benefited from a CSIC scholarship, and have been partially supported by ANII as an SNI candidate researcher. I'm also grateful to the University of Notre Dame for having me as a visitor on several occasions.

Special thanks go to my friends in Rio: Guarino, Vivis, Luciano and Yuri. To the Uruguayans in Paris especially Andrés Sambarino and Matías Carrasco. To Safia and Patrice for their hospitality. To the fellow students at LPMA especially Cyril, Guillaume, Bastien, Nelo, Xan, Alexandre, and Nikos. To the fluid dynamics and lunch-coffee-chocolate group at Notre Dame, Liviu, Misha, François, Roxana, Nero, and Richard. To the young guys (and gal) at Notre Dame, Quinn, James, Ryan, Renato, and Amy. To the group back in Montevideo especially Gordo, Diego, Leva, Juliana, Rata and Frodo. To my parents Enrique and María Paz, my sister Carolina, and my fiancée Marcela.

Abstract

Brownian motion on stationary random manifolds

Abstract

We introduce the concept of a stationary random manifold with the objective of treating in a unified way results about manifolds with transitive isometry group, manifolds with a compact quotient, and generic leaves of compact foliations. We prove inequalities relating linear drift and entropy of Brownian motion with the volume growth of such manifolds, generalizing previous work by Avez, Kaimanovich, and Ledrappier among others. In the second part we prove that the leaf function of a compact foliation is semicontinuous, obtaining as corollaries Reeb's local stability theorem, part of Epstein's local structure theorem for foliations by compact leaves, and a continuity theorem of Álvarez and Candel.

Keywords

Ergodic theory, Random manifolds, Brownian motion, Entropy, Liouville property.

Mouvement brownien sur les variétés aléatoires stationnaires

Résumé

On introduit le concept d'une variété aléatoire stationnaire avec l'objectif de traiter de façon unifiée les résultats sur les variétés avec un group d'isométries transitif, les variétés avec quotient compact, et les feuilles génériques d'un feuilletage compact. On démontre des inégalités entre la vitesse de fuite, l'entropie du mouvement brownien et la croissance de volume de la variété aléatoire, en généralisant des résultats d'Avez, Kaimanovich, et Ledrappier. Dans la deuxième partie on démontre que la fonction feuille d'un feuilletage compact est semicontinue, en obtenant comme conséquences le théorème de stabilité local de Reeb, une partie du théorème de structure local pour les feuilletages à feuilles compactes d'Epstein, et un théorème de continuité d'Álvarez et Candel.

Mots-clefs

Théorie ergodique, Variétés aléatoires, Mouvement brownien, Entropie, Propriétés de Liouville.

Movimiento Browniano en variedades aleatorias estacionarias

Resumen

Introducimos el concepto de variedad aleatoria estacionaria con el fin de probar en forma unificada resultados sobre variedades con grupo de isometría transitivo, variedades con cociente compacto, y hojas genéricas de foliaciones compactas. Probamos desigualdades relacionando la velocidad de escape del movimiento Browniano con la entropía y el crecimiento de volumen de dichas variedades generalizando trabajos anteriores de Avez, Kaimanovich, y Ledrappier entre otros. En la segunda parte mostramos que la función hoja de una foliación compacta es semicontinua, obteniendo como corolarios el teorema de estabilidad local de Reeb, parte del teorema de estructura local de Epstein para foliaciones por hojas compactas, y el teorema de continuidad de Álvarez y Candel.

Palabras claves

Teoría ergódica, Variedades aleatorias, Movimiento Browniano, Entropía, Propiedad de Liouville.

Contents

Introduction

This thesis has two parts. In the first we are concerned with harmonic measures of foliations, the entropy theory of Riemannian manifolds, Liouville properties, and how they relate to the behavior of Brownian motion. The main influence for this part of our work are the two very rich papers of Kaimanovich, [Kaĭ86] and [Kaĭ88]. By analogy with the, by then, well established theory of entropy for random walks on discrete groups (see [KV83]), Kaimanovich defined an entropy for Riemannian manifolds and outlined how this entropy would relate to the Liouville properties, algebraic properties of the fundamental group of the manifold, and volume growth among other things. The main idea is that the Riemannian metric on the manifold must have some sort of recurrence in order for this statistical approach to work. The three main cases where one suspects that this recurrence condition is satisfied are: manifolds with transitive isometry group, manifolds with a compact quotient, and generic leaves of compact foliations.

The entropy theory of discrete groups has seen its results successively generalized to more general types of graphs then just Cayley graphs. So for example in [KW02] the theory is worked out for graphs whose isometry group is transitive. Benjamini and Curien [BC12] have introduced the concept of a stationary random graph which simultaneously generalizes and includes the case of Cayley graphs and graphs with transitive isometry. Following their lead we introduce the concept of a stationary random manifold, which is a rooted random manifold whose distribution is invariant under re-rooting by Brownian motion, and develop the basic theory of entropy for them. This concept allows for unified proofs of results about manifolds with transitive isometry group, compact quotient, or generic leaves of compact foliations.

The first part of the thesis is divided into three chapters. In the first we deal with results about a single manifold, roughly the relationship between the Liouville property and the behavior of Brownian motion on the manifold. In the second we introduce stationary random manifolds and three invariants for them: Kaimanovich entropy, linear drift (the mean rate of displacement of Brownian motion from the basepoint), and volume growth (which measures the exponential growth rate of the volume of balls as a function of their radii). The main results are that entropy exists, is non-negative, and is zero if and only if the manifold is almost surely Liouville; and the basic inequalities relating the three asymptotic quantities which are

$$\frac{1}{2}\ell(M)^2 \leq h(M) \leq \ell(M)v(M)$$

where $\ell(M), h(M)$ and $v(M)$ are the drift, entropy and volume growth respectively. These inequalities have several interesting consequences which had been established independently before. The third chapter deals with Brownian motion on stationary random manifolds and can be seen as a more detailed look at the results in the second chapter (in particular we improve the lower bound for entropy above, following the results of

Kaimanovich and Ledrappier for manifolds with compact quotient [Kaĭ86, Theorem 10], [Led10, Theorem A]).

In the second part of the thesis we prove geometric results about how the leaves of a compact foliation vary from point to point. The only thing we had needed to know about compact foliations in the first part of the thesis was that the leaf of a random point whose distribution is harmonic in the sense of Garnett [Gar83] is an example of a stationary random manifold. In this part of the thesis we look at the continuity properties of the leaf function, i.e. the function associating to each point its leaf considered as a Riemannian manifold with basepoint. The main influence here is the work of Álvarez and Candel [ÁC03] where they introduced the leaf function and outlined how it could be used to study the quasi-isometry invariants of topologically generic leaves of foliations (we do not work on their theory explicitly but the concept of the leaf function and their statement that it is continuous on the set of leaves without holonomy have been the main seeds for our research). We establish that the leaf function is semicontinuous in the sense that any limit of leaves of a converging sequence of points is a covering space of the leaf of the limit point. Furthermore we provide an upper bound for the largest covering space which can be obtained in this way, the holonomy cover. As a consequence we obtain, Álvarez and Candel's theorem that the leaf function is continuous on the set of leaves without holonomy, Reeb's local stability theorem, and part of Epstein's local structure theorem for foliations by compact leaves. The main tools for this part of the work come from the convergence theory of Riemannian manifolds of Cheeger, Gromov, Anderson, etc; in particular we use the C^k-compactness theorem of [Pet06, Theorem 72].

Part I

Ergodic theory of stationary random manifolds

Chapter 1

Liouville properties and Zero-one laws on Riemannian manifolds

Introduction

In this chapter we establish some results which can be considered folklore of the boundary theory of Markov chains. Our motivation here has been to clarify and provide proofs as well as to introduce the concept of mutual information which will be important in our study of entropy in the next chapter (see Theorem 2.11).

We begin by recalling the definition and basic properties of the heat kernel and Brownian motion on a complete and stochastically complete Riemannian manifold. We then establish the correspondance between bounded tail measurable functions on the space of Brownian paths and bounded solutions to the backward heat equation on a Riemannian manifold (see Lemma 1.5). In particular this shows that there are no non-constant bounded solutions to the backward heat equation if and only if Brownian motion satisfies the zero-one law and, similarly, a manifold will satisfy the Liouville property (i.e. there are no non-constant bounded harmonic functions) if and only if its Brownian motion is ergodic (see Theorem 1.4). A treatment of these results in the case of discrete time Markov chains can be found in [Kai92].

We continue by providing an example (due to Kaimanovich, see [Kai92] and also [AT11, Lemma 1.1, Remark 4.9]) of a manifold where not every bounded solution to the backward heat equation is a harmonic functions (i.e. there is a tail event which does not coincide with any invariant event even after modification on a set of zero probability). We then show that such examples do not occur among manifolds with bounded geometry (this is a particular case of Derrienic's zero-two law, see [Der85]).

Finally, we introduce the concept of mutual information and show that it can be used to characterize when Brownian motion satisfies the zero-one law. This idea was used by Varopoulos to show that any Riemannian manifold with a compact quotient and subexponential volume growth satisfies the Liouville property (see [Var86]). We also use mutual information to provide proofs of some results on Kaimanovich entropy announced in [Kaĭ86] and [Kaĭ88] (notably existence of entropy and equivalence of the Liouville property to it being zero on stationary random manifolds, see Theorem 2.11). For example, it follows from the properties of mutual information established in this chapter that that on a bounded geometry manifold M the limit

$$\epsilon(x) = \lim_{t\to+\infty} \int \log\left(\frac{q(t-1,y,z)}{q(t,x,z)}\right) q(1,x,y)q(t-1,y,z)\mathrm{d}y\mathrm{d}z$$

exists and is non-negative for all $x \in M$ where $q(t, x, y)$ is the transition probability density of Brownian motion, and that the Liouville property is equivalent to $\epsilon(x)$ being 0 for some x (see [Kaĭ88, Lemma 1] and [Kai92, Section 3]).

1.1 Brownian motion and the backward heat equation

1.1.1 Laplacian, heat semigroup, and heat kernel

Consider a connected complete d-dimensional Riemannian manifold M. We will begin by recalling the basic properties of the Laplacian, the heat semigroup, and the heat kernel on M. Detailed treatment can be found in [Gri09].

The Laplacian $\Delta f(x)$ of a smooth function $f : M \to \mathbb{R}$ at a point $x \in M$ is defined as the sum of second derivatives of f along d perpendicular geodesics through the point x. This coincides with the Euclidean Laplacian at $0 \in \mathbb{R}^d$ of the pullback of f under a normal parametrization around the point x (in particular, since the Euclidean Laplacian is invariant under rotations, our definition is independent of the choice of geodesics through p).

It can be shown (see [Gri09, Chapter 3]) that the manifold Laplacian satisfies integral formulas analogous to those satisfied by the Euclidean Laplacian on $\mathbb{R}^d$. Integration by parts takes the form

$$\int f(x)\Delta g(x)\mathrm{d}x = -\int \langle \nabla f(x), \nabla g(x)\rangle \mathrm{d}x$$

for all smooth $f, g : M \to \mathbb{R}$ with compact support where integration is with respect to the Riemannian volume (proofs involve local calculations plus partitions of the unity).

The above integral formula implies that the Laplacian Δ is non-positive definite and coincides with its adjoint Δ^* when restricted to the subspace of $L^2(M)$ consisting of smooth functions with compact support. The domain of Δ^* is strictly larger than the subspace of smooth functions with compact support. However , it is possible (using the notion of weak derivatives) to find a subspace of $L^2(M)$ containing the smooth functions with compact support such that Δ^* is self adjoint when restricted to this subspace (see [Gri09, Chapter 4]). Abusing notation we denote the self-adjoint extension of the Laplacian by Δ.

The spectral theorem now implies that Δ is conjugate via an isometry to multiplication by a non-positive function $\phi : X \to \mathbb{R}$ on the space $L^2(X, \mu)$ of square integrable functions on some measure space (X, μ). In particular one can define for each $t \geq 0$ the operator $P^t = \exp(t\Delta)$ so as to be conjugate to multiplication by $\exp(t\phi)$. This defines a semigroup of bounded operators (i.e. $P^{t+s} = P^t P^s$ and P^0 is the identity) with norm less than or equal to 1 (because $\exp(t\phi) \leq 1$) which can therefore be extended to all of $L^2(M)$ (instead of only the dense subspace on which Δ was self-adjoint). This is the so-called heat semigroup.

The heat semigroup is continuous in the sense that $t \mapsto P^t f$ is continuous on $t \geq 0$ with respect to the L^2 norm for any $f \in L^2(M)$. Furthermore one has

$$\partial_t P^t f = \Delta f$$

for all $t > 0$ and all $f \in L^2(M)$ (in particular it is implied that $P^t f$ belongs to the domain of definition of the self-adjoint extension of Δ) where the limit on the left hand side and the equality are interpreted in $L^2(M)$ (see [Gri09, Theorem 4.9]).

Using a local argument involving Sobolev's embedding theorem one obtains that $P^t f$ is a smooth function on M for any $t > 0$ and $f \in L^2(M)$. Furthermore the function

$u(t,x) = P^t f(x)$ is smooth at all $t > 0$ and $x \in M$ and satisfies the heat equation

$$\partial_t u(t,x) = \Delta_x u(t,x).$$

The Riesz representation theorem yields the existence for each $t > 0$ and $x \in M$ of a function $p(t,x,\cdot) \in L^2(M)$ such that

$$P^t f(x) = \int p(t,x,y) f(y) \mathrm{d}y$$

for all $f \in L^2(M)$.

The semigroup property yields

$$p(t+s,x,z) = \int p(t,x,y)p(s,y,z)\mathrm{d}y$$

so that $u(t,y) = p(t,x,y)$ satisfies the heat equation and (by the regularity properties above) is smooth. At this point one may replace $p(t,x,y)$ by the symmetric integral $\int p(t/2,x,z)p(t/2,y,z)\mathrm{d}z$ so that $p(t,x,y)$ is smooth with respect to all three variables.

The function $p(t,x,y)$ is called the heat kernel of M. The maximum principle for parabolic equations implies that one always has $p(t,x,y) > 0$ and one can show that $\int p(t,x,y)\mathrm{d}y \le 1$ for all $t > 0$. If the last integral is always equal to 1 then one says that M is stochastically complete.

The Euclidean plane minus one point $\mathbb{R}^2 \setminus \{0\}$ is an example of a stochastically complete manifold which is not complete. An example of a complete but not stochastically complete manifold can be obtained by endowing the plane $\mathbb{R}^2$ with the Riemannian metric given in polar coordinates by

$$\mathrm{d}s^2 = \mathrm{d}r^2 + p(r)^2 \mathrm{d}\theta^2,$$

for some function p satisfying $p(r) = e^{r^3}$ for all r large enough.

1.1.2 Brownian motion

On $\mathbb{R}$ one has $p(t,x,y) = (4\pi t)^{-1/2} \exp\left(-\frac{(x-y)^2}{4t}\right)$. We notice that the density of the time t of a standard Brownian motion starting at $x \in \mathbb{R}$ can be written as $p(t/2,x,y)$. When passing to a Riemannian manifold we have decided to keep this factor of $\frac{1}{2}$ which distinguishes the heat kernel from the transition density function of Brownian motion. With this choice Brownian motion on a Riemannian motion solves the simplest possible stochastic differential equation driven by a standard Brownian motion on $\mathbb{R}^d$. To avoid confusion we keep the notation $q(t,x,y) = p(t/2,x,y)$.

Given a stochastically complete manifold M and $x \in M$ as above we define Weiner measure $\mathbb{P}_x$ starting at x on the space $\Omega = C([0,+\infty), M)$ of continuous paths from $[0,+\infty)$ to M as the unique Borel measure (the topology being that of uniform convergence on closed intervals) such that for all Borel sets $A_1, \ldots, A_n \subset M$ and all positive times $t_1 < \cdots < t_n$ the probability

$$\mathbb{P}_x(\omega_{t_1} \in A_1, \cdots, \omega_{t_n} \in A_n)$$

of the set of paths $\omega \in \Omega$ which visit each A_i at the corresponding time t_i is given by the integral

$$\int\limits_{A_1 \times \cdots \times A_n} q(t_1,x,x_1)q(t_2-t_1,x_1,x_2)\cdots q(t_n-t_{n-1},x_{n-1},x_n)\mathrm{d}x_1\cdots\mathrm{d}x_n.$$

A Brownian motion with initial distribution μ (a Borel probability on M) is defined to be an M valued stochastic process whose distribution is given by

$$\int \mathbb{P}_x \mathrm{d}\mu(x).$$

With the above definition one can prove the existence of manifold valued Brownian motion via Kolmogorov's continuity theorem using further properties of the heat kernel (i.e. upper bounds in terms of distance).

Perhaps the most elegant construction of manifold Brownian motion (usually attributed to Eells, Elsworthy, and Malliavin, e.g. see [Hsu02, pg. 75]) is as a diffusion on the orthogonal frame bundle $O(M)$.

Consider the smooth vector fields $V_i, i = 1, \ldots, d$ on $O(M)$ such that the flow of V_i applied to a frame $X = (x, v_1, \ldots, v_d) \in O(M)$ (here $x \in M$ and the v_i form an orthonormal basis of the tangent space at x) moves the basepoint along the geodesic with initial condition v_i and transports the frame horizontally. Then any solution to the Stratonovich stochastic differential equation

$$\mathrm{d}X_t = \sum_{i=1}^{d} V_i(X_t) \circ \mathrm{d}W_t^i$$

driven by a standard Brownian motion $(W_t^1, \ldots, W_t^d)$ in $\mathbb{R}^d$, projects to a Brownian motion on M.

The equivalence of these two approaches is established in [Hsu02, Propositions 3.2.2 and 4.1.6].

1.1.3 Zero-one laws

For each $t \geq 0$ define the σ-algebra $\mathcal{F}_t$ of events occurring before time t as the Borel subsets of $\Omega = C([0, +\infty), M)$ generated by the open sets of the topology of uniform convergence on the interval $[0, t]$. Similarly we let $\mathcal{F}^t$ be the σ-algebra of events occurring after time t which is generated by the open sets of the topology of uniform convergence on compact subsets of the interval $[t, +\infty)$. Events belonging to all $\mathcal{F}^t$ are called tail events and form the tail σ-algebra defined by

$$\mathcal{F}^\infty = \bigcap_{t \geq 0} \mathcal{F}^t.$$

We notice that since Ω is separable and completely metrizable any probability on Ω is tight meaning we can find a compact subset having probability $1 - \epsilon$ for each $\epsilon > 0$ (see [Bil99, Theorem 1.3]). Compact subsets of Ω are characterized by the Arsela-Ascoli theorem as consisting of families of curves which are uniformly bounded and equicontinuous on each interval $[a, b]$, in particular on such subsets pointwise convergence coincides with local uniform convergence. Combining these two facts one sees that any Borel subset in Ω can be approximated (meaning the probability of the symmetric difference can be made arbitrarily small) by a finite disjoint unions of events of the form

$$\{\omega \in \Omega : \omega_{t_1} \in A_1, \ldots, \omega_{t_n} \in A_n\}$$

where $t_1 < \ldots < t_n$ and the sets A_i are Borel subsets of M. Similarly each set in $\mathcal{F}_t$ can be approximated by finite disjoint unions of events of the above form with $t_n \leq t$ and each set in $\mathcal{F}^t$ by events of the above form with the restriction $t_1 \geq t$.

The Markov property allows one to express the probability of a tail event with respect to the measure $\mathbb{P}_x$ as averages over y of the probabilities with respect to $\mathbb{P}_y$ of a 'shifted' event. More concretely let $\text{shift}^t : \Omega \to \Omega$ be defined for $t \geq 0$ by

$$(\text{shift}^t \omega)_s = \omega_{t+s},$$

one has the following property.

Lemma 1.1. *Let M be a complete connected and stochastically complete Riemannian manifold. For each tail event A the function*

$$u(t,x) = \mathbb{P}_x(\text{shift}^t(A))$$

solves the backward heat equation

$$\partial_t u(t,x) = -\frac{1}{2}\Delta u(t,x).$$

Proof. Fix $T > 0$ and set $v(t,x) = u(T-t,x)$ for each $t \in (0,T)$ and $x \in M$. By applying the Markov property one obtains

$$v(t,x) = \int q(t,x,y)\mathbb{P}_y\left(\text{shift}^T \text{shift}^{T-t} A\right) \mathrm{d}y = \int p(t/2,x,y)v(0,y)\mathrm{d}y$$

which implies that $\partial_t v(t,x) = \frac{1}{2}\Delta v(t,x)$ from which the desired result follows. □

We say that an event $A \subset \Omega$ is trivial if it has probability 0 or 1 with respect to all measures $\mathbb{P}_x$. Brownian motion on M is said to satisfy the zero-one law if all tail events are trivial. Lemma 1.1 allows one to show that triviality of a tail event for $\mathbb{P}_x$ is independent of the choice of $x \in M$ (in particular the zero-one law can be verified at a single $x \in M$).

Corollary 1.2. *Let M be a complete connected and stochastically complete Riemannian manifold and $A \subset \Omega$ be a tail event. Then A is trivial if and only if* $\text{shift}^t(A)$ *has probability* 0 *or* 1 *with respect to some $\mathbb{P}_x$ for some $t \geq 0$.*

Proof. Apply the maximimum principle to $u(t,x)$ defined in Lemma 1.1. □

An event A is said to be invariant if $(\text{shift}^t)^{-1}(A) = A$ for all $t \geq 0$ (this implies $\text{shift}^t(A) = A$ since the shift maps are surjective). The σ-algebra of all invariant events is denoted by $\mathcal{F}^{\text{inv}}$. Since invariant events are also tail events one may apply Lemma 1.1 to obtain the following.

Corollary 1.3. *Let M be a complete connected and stochastically complete Riemannian manifold. For each invariant event A the function*

$$v(x) = \mathbb{P}_x(A)$$

is harmonic (i.e. $\Delta v(x) = 0$ for all x).

We say Brownian motion is ergodic on M if all invariant events are trivial. By Corollary 1.2 ergodicity is equivalent to triviality of all invariant events with respect to a single probability $\mathbb{P}_x$.

1.1.4 Liouville properties

A manifold M is said to satisfy the Liouville property (some times we just say M is Liouville) if it admits no non-constant bounded harmonic functions. Similarly we say M is backward-heat Liouville if it admits no non-constant bounded solutions to the backward heat equation (defined for all $t \geq 0$).

Theorem 1.4. *Let M be a complete connected and stochastically complete Riemannian manifold. Then M is backward-heat Liouville if and only if its Brownian motion satisfies the zero-one law. Similarly, M is Liouville if and only if its Brownian motion is ergodic.*

Proof. Suppose M is backward-heat Liouville and A is a tail event. Then by Lemma 1.1 the function
$$u(t,x) = \mathbb{P}_x(\text{shift}^t A)$$
solves the backward equation and by hypothesis must be constant.

Given times $t_1 < \cdots < t_n$ and Borel sets $A_1, \ldots A_n \subset M$ we calculate using the Markov property (which is possible because $A \in \mathcal{F}^{t_n}$) to obtain that the probability

$$\mathbb{P}_x\left(\omega_{t_i} \in A_i \text{ for } i = 1, \ldots, n \text{ and } \omega \in A\right)$$

of the trajectory belonging to A while hitting each A_i at the corresponding time t_i is given by
$$\int\limits_{A_1 \times \cdots \times A_n} q(t_1, x, x_1) \cdots q(t_n - t_{n-1}, x_{n-1}, x_n) \mathbb{P}_{x_n}(\text{shift}^{t_n} A) \mathrm{d}x_1 \cdots \mathrm{d}x_n$$

which since $u(t,x)$ is constant yields

$$\mathbb{P}_x\left(\omega_{t_i} \in A_i \text{ for } i = 1, \ldots, n\right) \mathbb{P}_x\left(\omega \in A\right).$$

This implies that A is independent from $\mathcal{F}_t$ for all t so that A is independent from itself and must have probability 0 or 1. We conclude that if M is backward-heat Liouville then its Brownian motion satisfies the zero-one law (notice that the proof mimics that of the classical zero-one law).

The same argument shows that if M is Liouville then its Brownian motion is ergodic.

On the other hand if there is a bounded backward solution $u(t,x)$ defined for all $t \geq 0$ then
$$u(t, \omega_t)$$
is a bounded martingale with respect to any $\mathbb{P}_x$. Since $u(t, \cdot)$ is not constant (otherwise u would be constant) the random variable $u(t, \omega_t)$ is not almost-surely constant with respect to $\mathbb{P}_x$. On the other hand the martingale convergence theorem implies that the limit

$$f(\omega) = \lim_{t \to +\infty} u(t, \omega_t)$$

exists almost surely with respect to $\mathbb{P}_x$ and that its conditional expectation to $\mathcal{F}_t$ is $u(t, \omega_t)$. This shows that f is not almost-surely constant with respect to $\mathbb{P}_x$ and, since L is tail measurable, there are non-trivial tail events.

In the case where one assumes that there is a non-constant bounded harmonic function $v(x)$ one has that $u(t,x) = v(x)$ is a bounded backward solution independent of t. The same argument above works with the additional fact that the limit f is shift invariant and hence yields non-trivial invariant events. $\square$

We conclude this subsection reexamining the last part of the previous proof (i.e. the construction of bounded tail measurable function $f : \Omega \to \mathbb{R}$ starting from a bounded backward solution $u(t,x)$). In view of Corollary 1.2 all the measures $\mathbb{P}_x$ are mutually absolutely continuous when restricted to the tail σ-algebra $\mathcal{F}^\infty$. We call the measure class of any and all $\mathbb{P}_x$ the harmonic measure class on $\mathcal{F}^\infty$. We say a tail measurable function $f : \Omega \to \mathbb{R}$ is invariant if $f \circ \mathrm{shift}^t = f$ for all $t \geq 0$.

Lemma 1.5. *Let M be a complete connected and stochastically complete Riemannian manifold. There is a one to one correspondence associating to each bounded solution $u(t,x)$ to the backward equation $\partial_t u(t,x) = -\frac{1}{2}\Delta u(t,x)$ the bounded tail measurable function*

$$f_u(\omega) = \lim_{t\to+\infty} u(t,\omega_t)$$

considered up to modifications on zero-measure sets with respect to the harmonic measure class. Furthermore f_u can be modified on a null set with respect to the harmonic measure class so that it is shift invariant if and only if $u(t,x) = v(x)$ for some bounded harmonic function $v : M \to \mathbb{R}$.

Proof. First of all we fix $x \in M$ and notice that $u(t,\omega_t)$ is a bounded martingale with respect to $\mathbb{P}_x$ so that the limit $f_u(\omega)$ exists $\mathbb{P}_x$-almost surely. Since the existence of the limit f_u is a tail event this implies that f_u is well defined almost surely with respect to the harmonic measure class on $\mathcal{F}^\infty$.

We will now show that $u \mapsto f_u$ is injective.

For this purpose suppose $f_u = f_v$ almost surely with respect to $\mathbb{P}_x$. By the martingale convergence theorem the conditional expectation of f_u to $\mathcal{F}_t$ with respect to $\mathbb{P}_x$ is given by

$$\mathbb{E}_x\left(f_u|\mathcal{F}_t\right) = u(t,\omega_t)$$

and similarly for f_v so that one has for each $t \geq 0$ that

$$u(t,\omega_t) = v(t,\omega_t)$$

for $\mathbb{P}_x$ almost every $\omega \in \Omega$. Since ω_t has a strictly positive density $q(t,x,\cdot)$ under $\mathbb{P}_x$ and the functions $u(t,\cdot)$ and $v(t,\cdot)$ are continuous this implies that $u(t,\cdot) = v(t,\cdot)$ for each t so that $u = v$ as claimed.

If $u(t,x) = v(x)$ for some harmonic function v then

$$f_u(\omega) = \lim_{t\to+\infty} v(\omega_t) = \lim_{t\to+\infty} v(\omega_{t+s}) = f_u(\mathrm{shift}^s\omega)$$

almost surely with respect to the harmonic measure class so f_u can be modified on a zero measure set to be invariant.

Reciprocally assume that f_u is shift invariant. One has

$$\lim_{t\to+\infty} u(t,\omega_t) = f_u(\omega) = f_u(\mathrm{shift}^s\omega) = \lim_{t\to+\infty} u(t,\omega_{t+s}) = \lim_{t\to+\infty} u(t-s,\omega_t).$$

Setting $u_s(t,x) = u(t-s,x)$ [1] one obtains that $f_u = f_{u_s}$ so that by the previously established injectivity $u = u_s$. Since this works for all s we obtain that $u(t,x) = v(x)$ for some harmonic function v.

It remains only to show that the map $u \mapsto f_u$ is surjective.

1. One can extend u_s to $t \leq s$ uniquely using the heat equation.

By Lemma 1.6 below for each t and x there is a probability $\mathbb{P}_{(t,x)}$ on $\mathcal{F}^t$ which satisfies

$$\mathbb{P}_{(t,x)}(A) = \mathbb{P}_x\left(\mathrm{shift}^t(A)\right).$$

Denoting by $\mathbb{E}_{(t,x)}$ the expectation with respect to $\mathbb{P}_{(t,x)}$ and setting

$$u(t,x) = \mathbb{E}_{(t,x)}(f(\omega))$$

one has by the martingale convergence theorem and Lemma 1.6 that $f = f_u$. Hence $u \mapsto f_u$ is surjective as claimed. $\square$

Lemma 1.6. *Let M be a complete connected and stochastically complete Riemannian manifold. For each $t \geq 0$ the map shift^t is a bijection between the σ-algebras $\mathcal{F}^T$ and $\mathcal{F}^{T-t}$ on Ω for all $T \geq t$. In particular each shift^t is a bijection on $\mathcal{F}^\infty$.*

Furthermore, denoting by $\mathbb{P}_{(t,x)}$ the unique probability on $\mathcal{F}^t$ which satisfies

$$\mathbb{P}_{(t,x)}(A) = \mathbb{P}_x\left(\mathrm{shift}^t(A)\right)$$

for all $A \in \mathcal{F}^t$ one has that the conditional expectation of any bounded and tail measurable function $f : \Omega \to \mathbb{R}$ to the σ-algebra $\mathcal{F}_t$ relative to the probability $\mathbb{P}_{x_0}$ (x_0 being any chosen point in M) is given by

$$\mathbb{E}_{x_0}(f(\omega)|\mathcal{F}_t) = u(t,\omega_t)$$

where $u(t,x) = \mathbb{E}_{(t,x)}(f(\omega))$ is the expectation of f relative to $\mathbb{P}_{(t,x)}$ for all $t \geq 0$ and $x \in M$.

Proof. We had glossed over this point earlier (e.g. in Lemma 1.1) but the continuity of shift^t does not imply that if $A \in \mathcal{F}^t$ then $\mathrm{shift}^t(A)$ is Borel.

However, if $\omega \in A$ for some $A \in \mathcal{F}^t$ then all continuous paths which coincide with ω after time t also belong to A. This property implies (valid for all $t \geq 0$) that shift^t is a bijection between $\mathcal{F}^T$ and $\mathcal{F}^{T-t}$ (even though shift^t certainly is not injective as a function on Ω) for all $T \geq t$.

The second claim amounts to establishing the fact that

$$\mathbb{E}_{x_0}(f(\omega)1_A(\omega)) = \mathbb{E}_{x_0}(u(t,\omega_t)1_A(\omega)) \tag{1.1}$$

for all $A \in \mathcal{F}_t$.

Suppose first that $f = 1_B$ for some $B \in \mathcal{F}^\infty$ and

$$A = \{\omega \in \Omega : \omega_{s_i} \in A_i, i = 1, \ldots, m\}$$

where the A_i are Borel subsets of M and $t_1 < \cdots < t_n \leq t$.

Then one has

$$\begin{aligned}
\mathbb{E}_{x_0}(f(\omega)1_A(\omega)) &= \mathbb{P}_{x_0}(A \cap B) \\
&= \int\limits_{A_1 \times \cdots \times A_m} q(s_1, x_0, x_1) \cdots q(t - s_m, x_m, y)\mathbb{P}_y\left(\mathrm{shift}^t B\right) \mathrm{d}x_0 \cdots \mathrm{d}x_m \mathrm{d}y \\
&= \mathbb{E}_{x_0}(u(t,\omega_t)1_A(\omega)).
\end{aligned}$$

Since any $A \in \mathcal{F}_t$ can be approximated (with respect to $\mathbb{P}_{x_0}$) by finite disjoint unions of events of the above form we have established the claim for bounded tail measurable functions that are indicators of a tail set.

For the general case notice that given two functions for which Equation 1.1 holds one has that the equation holds for any linear combination of them. Furthermore, if f is the monotone limit of a sequence of non-negative functions for which Equation 1.1 is known to hold then by the monotone convergence theorem the equation holds for f as well. This proves that the claim holds for all bounded tail measurable functions. □

1.2 A bounded backward heat solution

Since the heat equation regularizes functions one expects that most solutions to the backward heat equation should explode in finite time. In particular it seems plausible that the existence of bounded solutions $u(t,x)$, defined for all $t \in \mathbb{R}$, to the backward equation should be rather rare.

One way in which one can obtain a bounded solution to the backward heat equation is to set $u(t,x) = v(x)$ where v is a bounded harmonic function. In particular on the hyperbolic plane there exist many such bounded solutions.

With the above comments in mind one might conjecture that all bounded solutions to the backward heat equation come from bounded harmonic functions. Our purpose in this section is to show that this is not always the case.

We will construct a metric on the plane which is rotationally symmetric around the origin and show that there are non-trivial tail events with respect to the radial part of its Brownian motion which are not shift invariant.

This idea was suggested to us by Vadim Kaimanovich (see also [Kai92, pg. 23]).

Theorem 1.7. *Consider the smooth Riemannian metric g on the plane $\mathbb{R}^2$ which in polar coordinates has the form*

$$\mathrm{d}s^2 = dr^2 + p(r)^2 \mathrm{d}\theta^2$$

with $p(r) = re^{\frac{1}{2}r^2}$. There exists a smooth bounded function $u(t,x)$ which solves the backward heat equation with respect to this metric and such that $u(t,\cdot)$ is not harmonic for any $t \in \mathbb{R}$.

Proof. To see that such an expression in polar coordinates yields a smooth metric at the origin of $\mathbb{R}^2$ we calculate explicitly the coefficients of the metric (letting e_1, e_2 be the canonical basis of $\mathbb{R}^2$ and $(x,y) = (r\cos(\theta), r\sin(\theta))$) and obtain

$$g_{11} = g(e_1,e_1) = 1 + y^2(p(r)^2/r^2 - 1)/r^2$$

$$g_{12} = g(e_1,e_2) = -xy(p(r)^2/r^2 - 1)/r^2$$

$$g_{22} = g(e_2,e_2) = 1 + x^2(p(r)^2/r^2 - 1)/r^2,$$

so the claim follows because $(p(r)^2/r^2 - 1)/r^2$ can be extended analytically to $r = 0$ (just consider the power series of $p(r)$).

Consider a solution r_t to the Ito differential equation

$$\begin{cases} r_0 = 1 \\ \mathrm{d}r_t = \mathrm{d}X_t + f(r_t)\mathrm{d}t \end{cases}$$

where $f(r) = \frac{1}{2}p'(r)/p(r) = (r + 1/r)/2$ and X_t is a standard Brownian motion on $\mathbb{R}$. If one sets

$$\tau_T = \int_0^T \frac{1}{f(r_t)^2}\mathrm{d}t$$

and

$$\theta_t = Y_{\tau_t}$$

where Y_t is an Euclidean Brownian motion independent from X_t, then $(r_t\cos(\theta_t), r_t\sin(\theta_t))$ is a Brownian motion for the metric g (see [Hsu02, Example 3.3.3]).

We will show that there is a non-trivial tail event for the process r_t which is not shift invariant.

For this purpose notice that the fact that $f(r) \geq 1$ implies that

$$r_T = 1 + X_T + \int_0^T f(r_t)\mathrm{d}t \geq (1 + X_T + T)^+$$

for all T where $x^+ = x$ if $x > 0$ and 0 otherwise.

Next set $H(r) = \log(1 + r^2)$ and notice that $H'(r) = h(r) = 1/f(r)$ if $r > 0$. By the Ito formula one has

$$\mathrm{d}H(r_t) = h(r_t)\mathrm{d}X_t + (1 + \frac{1}{2}h'(r_t))\mathrm{d}t.$$

We will show that the limit

$$L = \lim_{t\to+\infty} H(r_t) - t$$

exists almost surely. Clearly L is tail measurable with respect to the filtration associated to r_t and is not shift invariant (replacing r_t by r_{t+s} changes the value of L by s as well). If we show that L is not almost surely constant then there are non-trivial tail events (of the form $\{L > a\}$) which are not shift invariant.

Notice that

$$H(r_T) - T = \int_0^T h(r_t)\mathrm{d}X_t + \frac{1}{2}\int_0^T h'(r_t)\mathrm{d}t.$$

Using the inequality $r_t \geq (1 + X_t + t)^+$ one obtains that for almost all trajectories one eventually has $r_t > t/2$. Combined with the fact that $|h'(r)| = O(1/r^2)$ when $r \to +\infty$ one obtains that

$$\int_0^{+\infty} |h'(r_t)|\mathrm{d}t < +\infty$$

almost surely.

To show that the martingale part of $H(r_T) - T$ converges it suffices to show that its variance is bounded. By the Ito isometry one has

$$\mathbb{E}\left[\left(\int_0^T h(r_t)\mathrm{d}X_t\right)^2\right] = \int_0^T \mathbb{E}\left[h(r_t)^2\right]\mathrm{d}t.$$

To bound the integrand we separate into two cases according to whether $|X_t| > t/2$ or not and obtain (for $t > 2$ using that $h \leq 1$ and that h is decreasing on $r > 1$)

$$\mathbb{E}\left[h(r_t)^2\right] \leq \mathbb{P}\left[|X_t| > t/2\right] + h(t/2)^2 = \mathbb{P}\left[|X_1| > \sqrt{t}/2\right] + h(t/2)^2.$$

The right hand side is integrable because the first term decreases exponentially while the second is of order $O(1/t^2)$.

Hence we have established that the limit L of $H(r_t) - t$ exists almost surely when $t \to +\infty$. To complete the proof it remains to show that the random variable L is not almost surely constant (see Figure 1.1 below for evidence supporting this claim).

Suppose that L were almost surely equal to a constant C. Let the stopping time σ for r_t be minimal among those with the property that $r_\sigma = 1$ and $r_t = 2$ for some $t < \sigma$. One always has $\sigma > 0$ and, by the Varadhan-Stroock support theorem, there is a positive probability that σ is finite. The Markov property implies that on the set with $\sigma < \infty$ one has

$$C = \lim_{t \to +\infty} H(r_t) - t = \lim_{t \to +\infty} H(r_{\sigma+t}) - t = C + \sigma$$

contradicting the fact that σ is positive. □

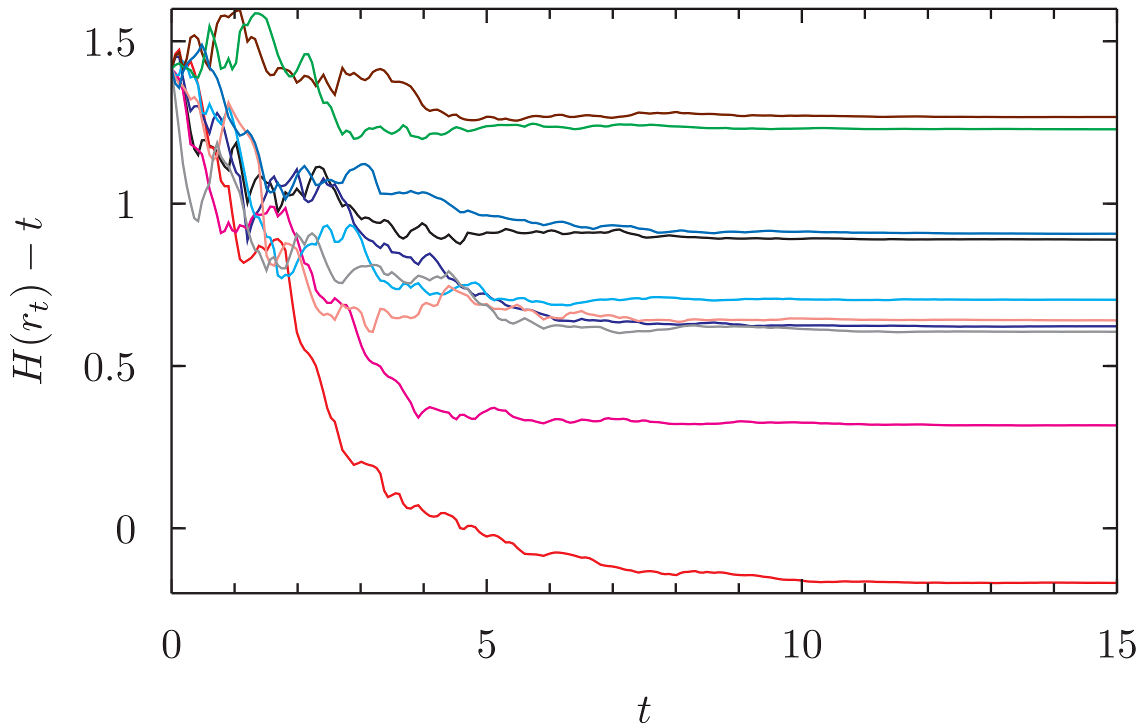

Figure 1.1: Ten trajectories of the process $H(r_t) - t$.

In the above example the radial process r_t grows super-linearly so that τ_t converges almost surely as $t \to +\infty$ and hence so does θ_t. Events of the form $\theta_\infty = \lim_{t\to+\infty} \theta_t \in [a, b]$ are invariant and therefore may be used to define non-constant bounded harmonic functions.

The existence of a manifold which satisfies the Liouville property but none the less admits non-constant bounded solutions to the backward heat equation was announced in [Kai92, pg. 23].

1.3 Steadyness of Brownian motion

In the previous section we gave an example of a radially symmetric Riemannian metric on $\mathbb{R}^2$ such that the corresponding Brownian motion had a non-trivial tail event which was not invariant. The curvature at distance r from the origin in this example can be calculated to be $-(3+r^2)$, in particular it is unbounded. We will show in this section that examples of this kind with bounded curvature and positive injectivity radius do not exist.

Following Kaimanovich we say Brownian motion on M is steady if every tail event can be modified on a null set with respect to the harmonic measure class on $\mathcal{F}^\infty$ to be invariant. This is equivalent (via Lemma 1.5) to the property that every bounded solution to the backward heat equation is of the form $u(t,x) = v(x)$ for some bounded harmonic function v.

Recall that a Riemannian manifold is said to have bounded geometry if its injectivity radius is positive and its sectional curvature is bounded in absolute value. In particular such a manifold is complete since unit speed geodesics starting at any point are always defined up to a time at least equal to the injectivity radius, and hence are defined for all time.

The following result was proved in the case M has a compact quotient under isometries by Varopoulos (see [Var86, pg. 359]). A more general result with no assumption on the injectivity radius of M was announced by Kaimanovich with a proof sketch (see [Kaĭ86, Theorem 1]).

Theorem 1.8. *Let M be connected Riemannian manifold with bounded geometry. Then M is stochastically complete and Brownian motion on M is steady. In particular every bounded solution $u(t,x)$ to the backward heat equation defined for all $t \geq 0$ is of the form $u(t,x) = v(x)$ for some harmonic function v.*

The so-called zero-two law is a sharp criteria for equivalence of the tail and invariant σ-algebras of Markov chains (see [Der76]). In our situation it amounts to the statement that

$$\sup_{x\in M}\left\{\lim_{t\to+\infty}\int |p(t+\tau,x,y) - p(t,x,y)|\mathrm{d}y\right\}$$

is either equal to 0 or to 2 for all $x \in M$ and all $\tau > 0$ and furthermore the limit is 0 if and only if Brownian motion is steady.

We will verify that the above limit cannot be 2 in Lemma 1.9 below. From this, steadiness of Brownian motion follows from the zero-two law. A proof which does not rely on the zero-two law will be given at the end of this subsection.

Lemma 1.9. *Let M be a connected Riemannian manifold with bounded geometry. For each $\tau > 0$ there exists $\epsilon_\tau > 0$ such that*

$$\int |q(t+\tau,x,y) - q(t,x,y)|\mathrm{d}y \leq 2 - \epsilon_\tau$$

for all $x \in M$ and $t \geq \tau$. In particular, if $u(t,x)$ is a solution to the backward equation $\partial_t u = -\frac{1}{2}\Delta_x u$ bounded by 1 in absolute value then

$$|u(t+\tau,x) - u(t,x)| \leq 2 - \epsilon_\tau$$

for all $t \geq 0$ and all $x \in M$.

Proof. Let $K > 0$ be a finite bound for the absolute value of all the sectional curvatures of M and $\rho > 0$ be strictly less than the injectivity radius at all points of M and the diameter of the d-dimensional sphere of constant curvature K.

Fix $x \in M$ and let $\psi : \mathbb{R}^d \to M$ be a normal parametrization at x, i.e. $\psi(v) = \exp_x \circ L(v)$ where $\exp_x : T_xM \to M$ is the Riemannian exponential map at x and $L : \mathbb{R}^d \to T_xM$ is a linear isometry between $\mathbb{R}^d$ (endowed with the usual inner product) and the tangent space T_xM (with the inner product given by the Riemannian metric on M).

Consider the metric of constant curvature $-K$ ball $B_\rho(0)$ of radius ρ centered at 0 in $\mathbb{R}^d$ of the form $\mathrm{d}s^2 = \mathrm{d}r^2 + p(r)\mathrm{d}\theta^2$ where $\mathrm{d}\theta^2$ is the standard Riemannian metric on the unit sphere $S^{d-1} \subset \mathbb{R}^d$ and one sets

$$p(r) = \sinh(\sqrt{K}r).$$

We denote by $\varphi(t,\cdot)$ the probability density of the time t of Brownian motion started at 0 and killed upon first exit from B_ρ with respect to the constant curvature metric above. The only fact about φ we need is that it is everywhere positive on B_ρ for all t.

Let $q_\rho^K(t,x,y)$ be defined for y in the open ball $B_\rho(x)$ of radius ρ centered at x by

$$q_\rho^K(t,x,y) = \varphi(t,\psi^{-1}(y))$$

where $\psi^{-1}(y)$ is the unique preimage of y in $B_\rho(0)$.

Theorem 1 of [DGM77] states that for all $y \in B_\rho(x)$ one has

$$q_\rho^K(t,x,y) \le q_\rho(t,x,y)$$

where $q_\rho(t,x,\cdot)$ is the probability density of the time t of Brownian motion on M started at x and killed upon first exit from the ball of radius ρ centered at x.

Also one has $q_\rho(t,x,y) \le q(t,x,y)$ since the probability of Brownian motion on M going from x to a small neighborhood of y in time t diminishes if one demands that it never exit the ball of radius ρ centered at x before that. Therefore one has

$$q_\rho^K(t,x,y) \le q(t,x,y)$$

for all $y \in B_\rho(x)$.

Define $\epsilon_\tau(x)$ by the equation

$$\epsilon_\tau(x) = \int\limits_{B_\rho(x)} \min(q_\rho^K(\tau,x,y), q_\rho^K(2\tau,x,y))\mathrm{d}y.$$

Let ω be the Euclidean volume form on B_ρ and $\lambda(p)\omega$ be the pullback of the volume form of M under ψ. Since the sectional curvature of M is bounded from above by K by [Pet06, Theorem 27] one has $\lambda(p) \ge \sin(\sqrt{K}r)^{d-1}$ for all p at distance r from 0 in B_ρ. Since $\epsilon_\tau(x)$ can be calculated by integrating a fixed positive function on B_ρ with respect to the form $\lambda(p)\omega$ one obtains that $\epsilon_\tau = \inf\{\epsilon_\tau(x), x \in M\}$ is positive.

Since $q(\tau,x,\cdot) \ge q_\rho^K(\tau,x,\cdot)$ and $q(2\tau,x,\cdot) \ge q_\rho^K(2\tau,x,\cdot)$ one obtains the following

$$\int |q(\tau,x,y)-q(2\tau,x,y)|\mathrm{d}y \le \int q(\tau,x,y)+q(2\tau,x,y)-\min(q(\tau,x,y),q(2\tau,x,y))\mathrm{d}y \le 2-\epsilon_\tau.$$

From this it follows for all $t \ge 0$ that

$$\int |q(t+2\tau,x,y) - q(t+\tau,x,y)|\mathrm{d}y \le \int |q(2\tau,x,z) - q(\tau,x,z)|q(t,z,y)\mathrm{d}z\mathrm{d}y \le 2 - \epsilon_\tau$$

as claimed.

To conclude we observe that if u satisfies the backward equation and is bounded in absolute value by 1 then one has

$$|u(t+\epsilon,x) - u(t,x)| = |\int (q(\tau,x,y) - q(2\tau,x,y))u(t+2\tau,y)\mathrm{d}y| \le 2 - \epsilon_\tau$$

which concludes the proof. □

As mentioned above one can prove Theorem 1.8 from the previous lemma using the zero-two law. The proof below relies instead on the bijection between bounded tail measurable functions and solutions to the backward equation (see Lemma 1.5).

Proof of Theorem 1.8. Let $v_r(x)$ denote the volume of the ball of radius r centered at a point $x \in M$. The lower curvature bound implies that v_r is less than or equal to the volume of a ball of radius r in hyperbolic space of constant curvature $-K$ (see [Pet06, Lemma 35]). In particular one has

$$\int_1^{+\infty} \frac{r}{\log(v_r)} \mathrm{d}r = +\infty$$

so that M is stochastically complete by [Gri09, Theorem 11.8].

Suppose that Brownian motion on M is not steady. Then one can find a non-trivial non-invariant (even up to modifications on null-sets with respect to the harmonic measure class) tail set A and $\tau > 0$ such that $B = \text{shift}^\tau(A)$ is disjoint from A. It follows from Lemma 1.2 that B is also non-trivial.

Consider the tail function defined by $f(\omega) = 1_A(\omega) - 1_B(\omega)$. By Lemma 1.5 there exists a bounded solution u to the backward equation such that $f = f_u$. By Lemma 1.6 one knows that u is bounded by 1 in absolute value almost eveywhere and by continuity of u this holds everywhere.

Notice that for almost every ω with respect to the harmonic measure class one has:

$$\lim_{t \to +\infty} u(t, \omega_t) = f(\omega)$$

and

$$\lim_{t \to +\infty} u(t - \tau, \omega_t) = \lim_{t \to +\infty} u(t, \omega_{t+\tau}) = f(\text{shift}^\tau(\omega)).$$

In particular by choosing such a generic path in B one obtains that there exists $\omega \in \Omega$ such that

$$\lim_{t \to +\infty} u(t, \omega_t) = -1$$

and

$$\lim_{t \to +\infty} u(t - \tau, \omega_t) = 1.$$

This implies that there exist values of t and x such that $u(t-\tau, x) - u(t, x)$ is arbitrarily close to 2, contradicting Lemma 1.9. □

1.4 Mutual information

Suppose M is a stochastically complete manifold whose Brownian motion satisfies the zero-one law. Then given $x \in M$, $A \in \mathcal{F}_t$ and $B \in \mathcal{F}^\infty$ one has that A and B are independent under $\mathbb{P}_x$ i.e. $\mathbb{P}_x(A \cap B) = \mathbb{P}_x(A)\mathbb{P}_x(B)$. The converse is also true, i.e. if each tail event B is independent from the events in $\mathcal{F}_t$ for all t then the Brownian motion on M satisfies the zero-one law (Proof: As in the proof of the classical zero-one law, one approximates B by events in $\mathcal{F}_t$ to show that it is independent from itself and hence trivial).

A, perhaps convoluted, but useful way of rephrasing this is the following: Consider the function $\omega \mapsto (\omega, \omega)$ from Ω to $\Omega \times \Omega$. Since this function is continuous one can push forward $\mathbb{P}_x$ to obtain a probability $\widehat{\mathbb{P}}_x$ on $\Omega \times \Omega$. The measure $\widehat{\mathbb{P}}_x$ describes the joint distribution of two copies of the same Brownian motion on M. On the other hand the probability $\mathbb{P}_x \times \mathbb{P}_x$ on $\Omega \times \Omega$ describes the joint distribution of two independent Brownian motions on M starting at x. The two probabilities $\widehat{\mathbb{P}}_x$ and $\mathbb{P}_x \times \mathbb{P}_x$ are very different (e.g. they are mutually singular). However, assuming the zero-one law is satisfied, if one restricts them both to the σ-algebra $\sigma(\mathcal{F}_t \times \mathcal{F}^\infty)$ generated by sets of the form $A \times B$

with $A \in \mathcal{F}_t$ and $B \in \mathcal{F}^\infty$ then they coincide. In fact, Brownian motion on M satisfies the zero-one law if and only if $\widehat{\mathbb{P}}_x$ and $\mathbb{P}_x \times \mathbb{P}_x$ coincide when restricted to $\sigma(\mathcal{F}_t \times \mathcal{F}^\infty)$ for all $t \geq 0$.

The mutual information between two random variables is a non-negative number which is zero if and only if they are independent. Given a σ-algebra $\mathcal{F}$ of Borel sets in Ω one may consider the identity map $\omega \mapsto \omega$ as a random variable from Ω endowed with the Borel σ-algebra to Ω endowed with $\mathcal{F}$, and hence one may define mutual information between σ-algebras.

Concretely, given $x \in M$ we define the mutual information between $\mathcal{F}_t$ and $\mathcal{F}^T$ (where $0 \leq t \leq T$ and possibly $T = \infty$) under $\mathbb{P}_x$ as

$$I_x\left(\mathcal{F}_t, \mathcal{F}^T\right) = \sup\left\{\sum_{i=1}^n \log\left(\frac{\widehat{\mathbb{P}}_x(A_i)}{\mathbb{P}_x \times \mathbb{P}_x(A_i)}\right)\widehat{\mathbb{P}}_x(A_i)\right\}$$

where the supremum is taken over all finite partitions $A_1, \ldots, A_n$ of $\Omega \times \Omega$ with each A_i belonging to $\sigma\left(\mathcal{F}_t \times \mathcal{F}^T\right)$. One may interpret the result as a measure of how much the behavior of Brownian motion after time T (or the tail behavior if $T = \infty$) depends on what happened before time t.

The fact that $I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ is always non-negative and is zero if and only if $\widehat{P}_x$ and $\mathbb{P}_x \times \mathbb{P}_x$ coincide on $\sigma\left(\mathcal{F}_t \times \mathcal{F}^T\right)$ follows from Jensen's inequality applied to the strictly convex function $-\log$ (see [Gra11, Lemma 3.1] for details).

Mutual information was used to unify results about random walks on discrete and continuous groups by Derriennic, in particular he established several results analogous to the Theorem below in that context (e.g. see [Der85, Section III]). In the case of a manifold with a compact quotient under isometries similar results to those below where established by Varopoulos (see [Var86, Part I.5]). Results of this type where also announced by Kaimanovich both in the case when M has a compact quotient and when M is a generic leaf of a compact foliation (e.g. [Kaĭ86, Theorem 2] and [Kaĭ88, Lemma 1]). In the context of discrete time Markov chain similar results are discussed in detail in [Kai92, Section 3].

Theorem 1.10. *Let M be a complete connected and stochastically complete Riemannian manifold. Then Brownian motion on M satisfies the zero-one law if and only if $I_x(\mathcal{F}_t, \mathcal{F}^\infty) = 0$ for some $t > 0$ and $x \in M$. Furthermore, the following properties hold for all $x \in M$ and $0 < t \leq T < \infty$:*

1. $I_x(\mathcal{F}_t, \mathcal{F}^T) = \int \log\left(\frac{q(T-t,x_1,x_2)}{q(T,x,x_2)}\right) q(t,x,x_1)q(T-t,x_1,x_2)\mathrm{d}x_1\mathrm{d}x_2.$

2. *The function $T \mapsto I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ is non-increasing and satisfies the inequality $I_x(\mathcal{F}_t, \mathcal{F}^\infty) \leq \lim\limits_{T\to+\infty} I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ with equality if some $I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ is finite.*

Proof. If Brownian motion on M satisfies the zero-one law then $\mathcal{F}_t$ is independent from $\mathcal{F}^\infty$ under $\mathbb{P}_x$ for all x and therefore $I_x(\mathcal{F}_t, \mathcal{F}^\infty) = 0$.

Assume now that there is some $x \in M$ with $I_x(\mathcal{F}_t, \mathcal{F}^\infty) = 0$ and fix $B \in \mathcal{F}^\infty$. We must show that $\mathbb{P}_x(B)$ is either 0 or 1.

For this purpose fix $s < t$ and and an open subset U of M and notice that

$$\mathbb{P}_x(\omega_s \in U, \omega \in B) = \int\limits_U q(s,x,y)\mathbb{P}_{(s,y)}(B)\mathrm{d}y.$$

On the other hand by hypothesis the above is also equal to

$$\mathbb{P}_x(\omega_s \in U)\mathbb{P}_x(B) = \int\limits_U q(s,x,y)\mathbb{P}_x(B)\mathrm{d}y$$

from which one obtains that

$$\mathbb{P}_{(s,y)}(B) = \mathbb{P}_x(B)$$

for almost all $y \in M$. Since $u(s,y) = \mathbb{P}_{(s,y)}(B)$ is a solution to the backward equation it must be constant and equal to $\mathbb{P}_x(B)$ for all s and y.

Consider now the set of paths where with $\omega_{t_i} \in A_i$ for all $i = 1,\ldots,n$ where $t_1 < \cdots < t_n$ and the A_i are Borel subsets of M. One may calculate using the above to obtain

$$\begin{aligned}\mathbb{P}_x\left(A \cap B\right) &= \int\limits_{A_1\times\cdots\times A_n} q(t_1,x,x_1)\cdots q(t_n - t_{n-1}, x_{n-1}, x_n)\mathbb{P}_{(t_n,x_n)}(B)\mathrm{d}y \\ &= \mathbb{P}_x(A)\mathbb{P}_x(B)\end{aligned}$$

so that B is independent from all events A of this form. Since B may be approximated with respect to $\mathbb{P}_x$ by finite disjoint unions of events of the form A above this shows that B is independent from itself and hence has probability equal to 0 or 1 as claimed.

We will now establish the integral formula for $I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ (property 1 above).

The so-called Gelfand-Yaglom-Perez Theorem (see [Pin64, Theorem 2.1.2] and the translator's notes on page 23 or [Gra11, Lemma 7.4] for further detail) implies that

$$I_x\left(\mathcal{F}_t, \mathcal{F}^T\right) = \int f_T(\omega^1,\omega^2)\log(f_T(\omega^1,\omega^2))\mathrm{d}(\mathbb{P}_x \times \mathbb{P}_x)(\omega^1,\omega^2)$$

where f_T is the Radon-Nikodym derivative of $\widehat{P}_x$ restricted to $\sigma\left(\mathcal{F}_t, \mathcal{F}^T\right)$ with respect to $\mathbb{P}_x \times \mathbb{P}_x$ restricted to the same σ-algebra. The formula then follows by substituting the explicit formula for f_T that we will establish below in Lemma 1.11. Notice that, because $x \mapsto x\log(x)$ is bounded from below, the integral formula always makes sense regardless of convergence considerations, but may assume the value $+\infty$.

We will now prove property 2 of the statement.

To begin notice that when T increases the set of partitions used to define $I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ decreases, hence the supremum taken over all such partitions decreases as well. This implies that $T \mapsto I_x\left(\mathcal{F}_t, \mathcal{F}^T\right)$ is decreasing and also that

$$I_x\left(\mathcal{F}_t, \mathcal{F}^\infty\right) \le \lim_{T\to+\infty} I_x\left(\mathcal{F}_t, \mathcal{F}^T\right).$$

Now assume that $I_x(\mathcal{F}_t, \mathcal{F}^{T_0})$ is finite and set $f = f_{T_0}$. Notice that by definition of the Radon-Nikodym derivative one has

$$\widehat{\mathbb{P}}_x(A) = \int\limits_A f(\omega^1,\omega^2)\mathrm{d}(\mathbb{P}_x \times \mathbb{P}_x)(\omega^1,\omega^2)$$

for all $A \in \sigma(\mathcal{F}_t, \mathcal{F}^{T_0})$. In particular the same equation is valid for all A in $\sigma(\mathcal{F}_t, \mathcal{F}^T)$ if $T > T_0$. This implies that whenever $T > T_0$ the function f_T coincides with the conditional expectation of f to the σ-algebra $\sigma(\mathcal{F}_t, \mathcal{F}^T)$ with respect to $\mathbb{P}_x \times \mathbb{P}_x$. Hence f_T is a reverse martingale (all statements of this type are relative to the measure $\mathbb{P}_x \times \mathbb{P}_x$ from now on) when $T \to +\infty$ and converges almost surely to f_∞ which is the Radon-Nikodym derivative of $\widehat{\mathbb{P}}_x$ with respect to $\mathbb{P}_x \times \mathbb{P}_x$ on $\sigma(\mathcal{F}_t, \mathcal{F}^\infty)$ (see [Doo01, pg. 483]).

It follows that $f_T \log(f_T)$ converges almost surely to $f_\infty \log(f_\infty)$ when T goes to $+\infty$ and it remains to show only that these functions are uniformly integrable in order to obtain that

$$\lim_{T\to+\infty} \int f_T \log(f_T) \mathrm{d}(\mathbb{P}_x \times \mathbb{P}_x) = \int f_\infty \log(f_\infty) \mathrm{d}(\mathbb{P}_x \times \mathbb{P}_x)$$

and conclude (by the Gelfand-Yaglom-Perez Theorem as above) that

$$\lim_{T\to+\infty} I_x(\mathcal{F}_t, \mathcal{F}^T) = I_x(\mathcal{F}_t, \mathcal{F}^\infty)$$

as claimed.

To simplify notation set $\varphi(x) = x\log(x)$ and $\mathcal{G}_T = \sigma(\mathcal{F}_t \times \mathcal{F}^T)$ (including the case $T = \infty$), and denote integration and conditional expectation with respect to $\mathbb{P}_x \times \mathbb{P}_x$ by $\mathbb{E}$. We notice that $x \mapsto \varphi(x)$ is convex and always larger than or equal to $-e^{-1}$ on $x \geq 0$.

Setting $g = \varphi(f)$ and $g_T = \varphi(f_T)$ one has by Jensen's inequality

$$-e^{-1} \leq g_T = \varphi(f_T) = \varphi\left(\mathbb{E}\left(f|\mathcal{G}_T\right)\right) \leq \mathbb{E}\left(\varphi(f)|\mathcal{G}_T\right).$$

By the reverse martingale convergence theorem (see [Doo01, pg. 483]) the right hand side converges in L^1 to $\mathbb{E}\left(\varphi(f)|\mathcal{G}_\infty\right)$. From this it follows that the functions g_T are uniformly integrable which concludes the proof of claim 2. $\square$

We now establish the result on the Radon-Nikodym derivative of $\widehat{\mathbb{P}}_x$ with respect to $\mathbb{P}_x \times \mathbb{P}_x$ which was used in the previous proof (see also [Var86, pg. 354]).

Lemma 1.11. *Let M be a complete connected and stochastically complete Riemannian manifold. Then for all x and $0 < t < T < +\infty$ the measure $\widehat{\mathbb{P}}_x$ restricted to $\sigma\left(\mathcal{F}_t \times \mathcal{F}^T\right)$ is absolutely continuous with respect to $\mathbb{P}_x \times \mathbb{P}_x$ restricted to the same σ-algebra and the corresponding Radon-Nikodym derivative is given by*

$$f_T(\omega^1, \omega^2) = \frac{q(T-t, \omega^1_t, \omega^2_T)}{q(T, \omega^1_0, \omega^2_T)}.$$

Proof. Consider two subsets of Ω defined by

$$A = \{\omega \in \Omega : \omega_{s_1} \in A_1, \ldots, \omega_{s_m} \in A_m\}$$

$$B = \{\omega \in \Omega : \omega_{t_1} \in B_1, \ldots, \omega_{t_n} \in B_n\}$$

where $s_1 < \cdots < s_m = t$, $T = t_1 < \cdots t_n$, and the sets A_i and B_j are Borel subsets of M.

By direct calculation using the definition of f_T we obtain that

$$\int_{A\times B} f_T(\omega^1,\omega^2)\mathrm{d}\mathbb{P}_x \times \mathbb{P}_x(\omega^1,\omega^2) = \int_{A\times B} \frac{q(T-t,\omega^1_t,\omega^2_T)}{q(T,\omega^1_0,\omega^2_T)} \mathrm{d}\mathbb{P}_x(\omega^1)\mathrm{d}\mathbb{P}_x(\omega^2).$$

The right hand side coincides (via the definition of $\mathbb{P}_x$) with the integral over $A_1 \times \cdots A_m \times B_1 \times \cdots \times B_n$ of

$$\frac{q(T-t,x_m,y_1)}{q(T,x,y_1)} q(s_1,x,x_1)\cdots q(s_m - s_{m-1}, x_{m-1}, x_m) q(t_1, x, y_1) \cdots q(t_n - t_{n-1}, x_{n-1}, x_n)$$

which after cancellation yields

$$\mathbb{P}_x(\omega_{s_1} \in A_1, \ldots, \omega_{s_m} \in A_m, \omega_{t_1} \in B_1, \ldots, \omega_{t_n} \in B_n).$$

This last probability is seen to be equal to $\widehat{P}_x(A \times B)$ by definition of $\widehat{\mathbb{P}}_x$.

Hence we have established that the integral of f_T with respect to $\mathbb{P}_x \times \mathbb{P}_x$ over any set of the form $A \times B$ as above is $\widehat{\mathbb{P}}_x(A \times B)$. Since any set in $\mathcal{G} = \sigma(\mathcal{F}_t \times \mathcal{F}^T)$ can be approximated by finite disjoint unions of such sets we have that the integral of f_T on any set of this σ-algebra with respect to $\mathbb{P}_x \times \mathbb{P}_x$ is equal to the probability of the set with respect to $\widehat{\mathbb{P}}_x$. As f_T is $\mathcal{G}$-measurable this implies that f_T is (a version of) the Radon-Nikodym derivative of $\mathbb{P}_x$ with respect to $\mathbb{P}_x \times \mathbb{P}_x$ on $\mathcal{G}$ as claimed. $\square$

Chapter 2

Entropy of stationary random manifolds

Introduction

In this chapter we introduce the notion of a stationary random manifold, which is a random Riemannian manifold with basepoint whose distribution is invariant under re-rooting by moving the basepoint a fixed time along a Brownian motion. The typical examples of stationary random manifolds are the following:

1. A single manifold with basepoint whose isometry group is transitive (e.g. a Lie group with a left invariant Riemannian metric).

2. A manifold M admitting a compact quotient with a random basepoint whose distribution is uniform on a fundamental domain.

3. The leaf of a random point in a foliation whose distribution is a harmonic measure in the sense of Lucy Garnett (see [Gar83]).

The point of the definition is that theorems for the above three special cases of manifolds can be dealt with in a uniform way. Our definition is analogous to the concept of a 'stationary random graph' of Benjamini and Curien (see [BC12]), which simultaneously generalizes and includes known results about random walks on discrete groups (see [Ave74] and [KV83]) and graphs with transitive isomorphism groups (see [KW02]).

We study three asymptotic quantities associated to stationary random manifolds: linear drift, which measures the mean displacement of Brownian motion from the origin per unit of time; entropy which measures the growth of differential entropy of the distribution of Brownian motion relative to the Riemannian volume measure; and volume growth which measures the exponential growth rate of the volume of balls in terms of their radii. All these quantities have been previously studied in different contexts, for example the linear drift of group random walks was studied by Guivarc'h (see [Gui80]), entropy of group random walks appears in the work of Avez (see [Ave74]), and the entropy of Riemannian manifolds was first defined by Kaimanovich (see [Kaĭ86]).

We show (see Theorem 2.11) that the entropy of a stationary random manifold is zero if and only if this manifold is almost surely Liouville, this result was announced by Kaimanovich in the case of manifolds with a compact quotient (see [Kaĭ86]) and generic leaves with respect to a harmonic measure in a foliation (see [Kaĭ88]).

Our main result (see Theorem 2.15) is that the following inequalities hold for ergodic stationary random manifolds

$$\frac{1}{2}\ell(M)^2 \leq h(M) \leq \ell(M)v(M)$$

where $\ell(M), h(M)$ and $v(M)$ are the drift, entropy, and volume growth respectively.

As a toy case is obtained by setting M constant equal to the Euclidean plane (with any fixed basepoint) one has that volume growth is subexponential (i.e. $v(M) = 0$) and hence entropy is zero. This implies that there are no non-constant bounded harmonic functions on the Euclidean plane, i.e. Liouville's theorem.

The same proof works for manifolds with a transitive isometry group. We obtain Avez's theorem which states that if such a manifold has subexponential volume growth then it satisfies the Liouville property (see [Ave74]).

For manifolds admitting a compact quotient one obtains the same result, i.e. subexponential volume growth implies the Liouville property. This was proved by Varopoulos (see [Var86]) and announced by Kaimanovich among a wealth of other results (see [Kaĭ86]). In particular Kaimanovich announced that the upper inequality $h(M) \leq \ell(M)v(M)$ held in this context (for group random walks this is essentially contained in the work of Guivarc'h, [Gui80]).

One might conjecture that all Riemannian manifolds with subexponential volume growth are Liouville. The counterexample given by $\mathbb{R}^2$ with the metric $\mathrm{d}s^2 = \mathrm{d}x^2 + (1+x^2)^2\mathrm{d}y^2$ (which admits the harmonic function $\arctan(x)$) was attributed to O. Chung (possibly Lung Ock Chung?) by Avez.

Even though they do not necessarily admit a compact quotient, generic leaves of a compact foliation are recurrent and therefore their Riemannian metric has a loosely repeating pattern (we might say that the metric is quasi-periodic and that if a compact quotient exists it is exactly periodic). It was announced by Kaimanovich in [Kaĭ88] that this recurrence is enough to show that if almost every leaf has subexpontial volume growth then almost every leaf is Liouville. Our result also implies this as a corollary.

The sharper lower bound for entropy $2\ell(M)^2 \leq h(M)$ was proved by Kaimanovich in the case of negatively curved manifolds admitting a compact quotient, and this was later generalized to any manifold admitting a compact quotient by Ledrappier (see [Kaĭ86, Theorem 10] and [Led10, Theorem A]). We will prove this inequality for a special type of stationary manifold in the next chapter (see 3.23).

The consequence that linear drift is positive if and only if entropy is, was established by Ledrappier and Karlsson in the case of a manifold with a compact quotient (see [KL07]).

To conclude we point out some of the main technical difficulties which distinguish the theory of stationary random manifolds from that of stationary random graphs.

An initial problem is that, while a Polish space of isometry classes of graphs is relatively easy to define, an analogous definition for a space of manifolds does not come so easily. For this purpose we use the idea of the Gromov space, a space whose points correspond to isometry classes of proper pointed metric spaces, which was introduced to the study of foliations in the work of Álvarez and Candel (see [ÁC03]). In order to show the the leaf of a random point in a foliation yields an example of a stationary random manifold we are lead to study the regularity properties of the leaf function establishing measurability in Lemma 2.8 and semicontinuity in the second part of the thesis (see Theorem 4.3).

Once this approach is adopted one must overcome the fact that convergence on the Gromov space is essentially a C^0 notion but one is interested in quantities such as the heat kernel which depend on derivatives of the Riemannian metric. We bridge this gap by restricting ourselves to subspaces of the Gromov space consisting of manifolds with 'uniformly bounded geometry' (see Theorem 2.4). On these subspaces one may use compactness theorems (in particular we use [Pet06, Theorem 72]) to show that convergence on the Gromov space is equivalent to higher order 'smooth convergence'. We can then use the continuous dependence of the heat kernel under smooth convergence (see [Lu12]) to establish the existence of harmonic measures and the necessary regularity of the quantities used to define linear drift and entropy.

2.1 The Gromov space and harmonic measures

2.1.1 The Gromov space

In this subsection we construct a model of 'the Gromov space' which is a complete separable metric space whose points represent the isometry classes of all proper (i.e. closed balls are compact) pointed metric spaces. The topology on the Gromov space is that of pointed Gromov-Hausdorff convergence (see [BBI01, Chapter 8]).

Our main point is that one can construct the Gromov space using well defined sets (i.e. avoiding use of 'the set of all metric spaces') and without using the axiom of choice (see [BBI01, Remark 7.2.5] and the paragraph preceding it). We will later be interested in certain probability measures on the Gromov space.

A sequence of pointed metric spaces (X_n, o_n) (here o_n is the basepoint of the space which we will sometimes abuse notation by omitting; also, we use d to denote the distance on different metric spaces simultaneously) is said to converge in the pointed Gromov-Hausdorff sense to a pointed metric space (X, o) if for each $r > 0$ and $\epsilon > 0$ there exists n_0 and for all $n > n_0$ a function $f_n : B_r(o_n) \to X$ (we use $B_r(x)$ to denote the ball of radius r centered at x in a metric space) satisfying the following three properties:

1. $f_n(o_n) = o$
2. $\sup\{|d(f_n(x), f_n(y)) - d(x, y)| : x, y \in B_r(o_n)\} < \epsilon$
3. $B_{r-\epsilon}(o) \subset \bigcup_{x \in B_r(o_n)} B_\epsilon(f_n(x))$.

Given two metric spaces X and Y we say a distance on the disjoint union $X \sqcup Y$ is admissible if it coincides with the given distance on X when restricted to $X \times X$ and similarly for Y.

Following Gromov (see [Gro81, Section 6]) we metricize pointed Gromov-Hausdorff convergence by defining the distance $\mathrm{d}_{\mathcal{GS}}(X_1, X_2)$ between two pointed metric spaces (X_1, o_1) and (X_2, o_2) as the infimum of all $\epsilon \in (0, \frac{1}{2})$ such that there exists an admissible distance d on the disjoint union $X_1 \sqcup X_2$ which satisfies the three inequalities $d(o_1, o_2) < \epsilon$, $d(B_{1/\epsilon}(o_1), X_2) < \epsilon$, and $d(X_1, B_{1/\epsilon}(o_2)) < \epsilon$; or $\frac{1}{2}$ if no such admissible distance exists (this truncation is necessary in order for $\mathrm{d}_{\mathcal{GS}}$ to satisfy the triangle inequality as noted by Gromov in the above-mentioned reference). The author is grateful to Jan Cristina for the following proof (see [Cri08]).

Lemma 2.1. *The distance $\mathrm{d}_{\mathcal{GS}}$ metricizes pointed Gromov-Hausdorff convergence.*

Proof. If a sequence (X_n, o_n) converges to (X, o) then given $\delta > 0$ and setting $\epsilon = \delta/2$ and $r = 2/\delta$ there exists for each sufficiently large n a function f_n satisfying the three properties in the definition above for r and ϵ.

Using this one can define $d : (X_n \sqcup X)^2 \to [0, +\infty)$ so that it coincides with the given distance on each X_i and satisfies

$$d(x, y) = d(y, x) = \inf\{d(x, x') + \delta + d(f(x'), y)\}$$

for all $x \in X_n$ and $y \in X$ (where the infimum is over $x' \in B_r(o_n)$).

One verifies that d is a distance on $X_n \sqcup X$ which shows that $\mathrm{d}_{\mathcal{GS}}(X_n, X) < \delta$.

For the converse statement suppose that $\mathrm{d}_{\mathcal{GS}}(X_n, X) < \delta$. Then setting $\epsilon = 2\delta$ and $r = 1/\delta$ there exists an admissible distance d on the disjoint union $X_n \sqcup X$ satisfying $d(o_n, o) < \delta$, $d(B_r(o_n), X) < \delta$, and $d(X_n, B_r(o)) < \delta$.

Using this we define $f_n : B_r(o_n) \to X$ so that $f_n(o_n) = o$ and $d(x, f_n(x)) < \delta$ for all $x \in B_r(o_n)$. By the triangle inequality we obtain

$$|d(f_n(x), f_n(y)) - d(x, y)| \le d(f_n(x), x) + d(y, f_n(y)) < \epsilon.$$

Given $y \in B_{r-\epsilon}(o)$ there exists $x \in X_n$ with $d(x, y) < \delta$. In particular $d(o_n, x) < d(o_n, o) + d(o, y) < r$. Hence f_n is defined at x and one has $d(y, f_n(x)) \le d(y, x) + d(x, f_n(x)) < \epsilon$. So we have established that

$$B_{r-\epsilon}(o) \subset \bigcup_{x \in B_r(o_n)} B_\epsilon(f(x))$$

which concludes the proof. □

We now consider the set of finite pointed metric spaces of the form (X, o) where $X = \{0, \dots, n\}$ for some non-negative integer n and $o = 0$. Consider two such pointed metric spaces to be equivalent if they are isometric via a basepoint preserving isometry (each equivalence class has finitely many elements) and let $\mathcal{FGS}$ (read 'finite Gromov space') be the set of all equivalence classes. One verifies that $(\mathcal{FGS}, \mathrm{d}_{\mathcal{GS}})$ is a separable metric space.

Definition 2.2. We define the Gromov space $(\mathcal{GS}, \mathrm{d}_{\mathcal{GS}})$ as the metric completion of $(\mathcal{FGS}, \mathrm{d}_{\mathcal{GS}})$.

With the above definition it follows immediately that the Gromov space is a complete separable metric space with $\mathcal{FGS}$ as a dense subset. It remains to show that each of its points 'represents' an isometry class of proper pointed metric spaces and that all such classes are represented by some point.

Theorem 2.3. *For each point p in $\mathcal{GS}$ there exists a unique (up to pointed isometry) proper pointed metric space (X, o) with the property that all sequences (X_n, o_n) of representatives of Cauchy sequences in $\mathcal{FGS}$ converging to p converge in the pointed Gromov-Hausdorff sense to (X, o).*

Furthermore, the thus defined correspondence between isometry classes of pointed proper metric spaces and points in the Gromov space is bijective.

Proof. Consider a sequence (X_n, o_n) of finite metric spaces representing some Cauchy sequence in $\mathcal{FGS}$. By taking a subsequence we assume that the distance between X_n and X_{n+1} is less than 2^{-n} for all n.

By definition there exists an adapted metric d_n on $X_n \sqcup X_{n+1}$ with the property that $d_n(o_n, o_{n+1}) < 2^{-n}$ and the ball of radius 2^n centered at the basepoint of either half is at distance less than 2^{-n} from the other half.

Let Y be the countable disjoint union of all X_n. We define a distance on Y by letting $d(x, y)$ in the case $x \in X_n$ and $y \in X_{n+k}$ be the infimum of

$$d_n(x_0, x_1) + \cdots + d_{n+k-1}(x_{k-1}, x_k)$$

over all sequences of k elements with $x = x_0, \ldots, x_k = y$ and $x_i \in X_{n+i}$ for all i. The other case is determined by symmetry.

Set $X = \widehat{Y} \setminus Y$ and $o = \lim_{n\to+\infty} o_n$ where $\widehat{Y}$ is the metric completion of Y. We claim (X, o) is proper and is the limit of (X_n, o_n) in the pointed Gromov-Hausdorff sense. Once this claim is established uniqueness of (X, o) up to pointed isometry is given by [BBI01, Theorem 8.1.7]. And, since pointed Gromov-Hausdorff convergence is characterized by $\mathrm{d}_{\mathcal{GS}}$ (see Lemma 2.1), the triangle inequality implies that (X, o) is also the limit of any Cauchy sequence equivalent to the one determined by (X_n, o_n).

We will now establish the claim.

Fix $r > 0$ and let B be the closed ball of radius r centered at o in X. We must show that B is compact.

For this purpose notice that for all n and all $k \geq 0$ one has that ball of radius 2^{n-1} centered at o_{n+k} is at distance less than $2^{-(n-1)}$ from X_n. If $2^{n-1} > r$ then one can approximate any x by a sequence x_k in Y with the property that eventually $d(x_k, o_{n_k}) < 2^{n-1}$ (where one chooses n_k so that $x_k \in X_{n_k}$). It follows that $n_k \to +\infty$ (otherwise infinitely many x_k would belong to the same X_N which is finite, and ultimately one obtains that $x \in X_N$) and one obtains that the distance between B and X_n is less than or equal to $2^{-(n-1)}$ as well. In particular since X_n is finite this shows that B can be covered by a finite number of balls of radius $2^{-(n-2)}$. This establishes that B is compact as claimed.

We have shown that each equivalence class of Cauchy sequences in Y determines a unique isometry class of pointed proper metric spaces. Now let (X, o) be a pointed proper metric space and for each n let $X_n = \{o_n = x_{n,0}, x_{n,1}, \ldots, x_{n,k_n}\}$ be a finite subset of $B_{2^n}(o)$ which is 2^{-n}-dense. There is a unique point $p_n \in \mathcal{FGS}$ such that all of its representatives are isometric to the pointed metric space (X_n, o_n). Since (X_n, o_n) converges in the pointed Gromov-Hausdorff sense to (X, o) it follows that any sequence of such representatives converges to (X, o) as well. From this one obtains that p_n converges to some point p in $\mathcal{GS}$ which represents the isometry class of (X, o). Hence the correspondence between points in $\mathcal{GS}$ and isometry classes of pointed proper metric spaces is bijective, which concludes the proof. □

In view of the above theorem we will no longer distinguish between a point in $\mathcal{GS}$ and a pointed proper metric spaces (X, o) in the isometry class represented by it.

2.1.2 Spaces of manifolds with uniformly bounded geometry

We say a complete connected Riemannian manifold M has geometry bounded by $(r, \{C_k\})$ (where r is a positive radius and C_k is a sequence of positive constants indexed on $k = 0, 1, \ldots$) if its injectivity radius is larger than or equal to r and one has

$$|\nabla^k R| \leq C_k$$

at all points, where R is the curvature tensor, $\nabla^k R$ is its k-th covariant derivate, and the tensor norm induced by the Riemannian metric is used on the left hand side.

We denote by $\mathcal{M}(d, r, \{c_k\})$ the subset of the Gromov space representing isometry classes of d-dimensional pointed Riemannian manifolds with geometry bounded by $(r, \{c_k\})$.

Following [Pet06, 10.3.2] we say a sequence (M_n, o_n, g_n) of pointed connected complete Riemannian manifolds (here g_n is the Riemannian metric an o_n the basepoint) converges smoothly to a pointed connected complete Riemannian manifold (M, o, g) if for each $r > 0$ there exists an open set $\Omega \supset B_r(o)$ and for n large enough a smooth pointed (i.e. $f_n(o) = o_n$) embedding $f_n : \Omega \to M$ such that the pullback metric $f_n^* g_n$ converges smoothly to g on compact subsets of Ω.

We will prove the following geometric result in the second part of the thesis (see Theorem 4.11).

Theorem 2.4. *Let $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ for some choice of dimension d, radius r, and sequence C_k. Then $\mathcal{M}$ is a compact subset of the Gromov space on which smooth convergence is equivalent to convergence in the pointed Gromov-Hausdorff sense.*

We will say a subset of the Gromov space 'consists of manifolds with uniformly bounded geometry' if it is contained in some set of the form $\mathcal{M}(d, r, \{C_k\})$. Elements of such subsets are represented by triplets (M, o_M, g_M) (o_M being the basepoint and g_M the Riemannian metric). We usually write just M leaving the other two elements implicit and will refer to them as o_M and g_M when needed.

Recall that a complete Riemannian manifold M is said to be stochastically complete if the integral of its heat kernel $p(t, x, y)$ with respect to y equals 1 for all $t > 0$ and $x \in M$. We denote by $q(t, x, y) = p(t/2, x, y)$ the transition probability density of Brownian motion on such a manifold M. With this convention one has that $q(t, x, y)$ is $(2\pi t)^{-1/2} e^{-(x-y)^2/2t}$ on $\mathbb{R}$ and $(2\pi t)^{-3/2} e^{-t/2 - d(x,y)^2/2t} d(x, y)/\sinh(d(x, y))$ on three dimensional hyperbolic space (see [DM88, pg. 185]).

We will need the following uniform upper bound on the transition density q for manifolds with uniformly bounded geometry.

Theorem 2.5. *Let $\mathcal{M}$ be a subset of the Gromov space consisting of n-dimensional manifolds with uniformly bounded geometry. Then for each $t_0 > 0$ and $D > 2$ there exist a positive constant C such that the inequality*

$$q(t, x, y) \leq C \exp\left(-\frac{d(x, y)^2}{Dt}\right)$$

hold for all $t \geq t_0$ and all pairs of points x, y belonging to any manifold M of $\mathcal{M}$.

Proof. The on diagonal bound given by Theorem 8 of [Cha84, pg. 198] (setting $r = \sqrt{t}$) yields a constant c_1 depending only on n such that any complete manifold of dimension n satisfies

$$q(t, x, x) \leq \frac{c_1}{\operatorname{vol}(B_{\sqrt{t/2}}(x))} t^{-\frac{n}{2}}.$$

Because of the uniform bounds on curvature and injectivity radius one may bound the volume of the ball of radius $\sqrt{t/2}$ from below uniformly on $\mathcal{M}$ by some multiple of $t^{n/2}$ for small t and by a constant for large t.

Using this one obtains that

$$q(t, x, x) \leq \frac{1}{\gamma(t)}$$

for all $t > 0$ and all x in a manifold of $\mathcal{M}$ where $\gamma(t)$ is of the form

$$\gamma(t) = \max(c_2 t^n, c_3).$$

One verifies that there exists c_4 such that $\gamma(t) \leq \gamma(2t) \leq c_4\gamma(t)$ for all $t > 0$ after which by Corollary 16.4 of [Gri09] one obtains that for each $D > 2$ there exist $c_5, c_6 > 0$ (depending on $\mathcal{M}$) such that

$$q(t,x,y) \leq \frac{c_5}{\gamma(c_6 t)} \exp\left(-\frac{d(x,y)^2}{Dt}\right)$$

for all $x, y \in M$, $t > 0$ and $M \in \mathcal{M}$.

Restricting to $t \geq t_0$ on obtains for each $D > 2$ a constant C such that

$$q(t,x,y) \leq C \exp\left(-\frac{d(x,y)^2}{Dt}\right)$$

for all $x, y \in M$, $t \geq t_0$ and $M \in \mathcal{M}$, as claimed. □

2.1.3 Harmonic measures

We say a probability measure μ on the Gromov space is harmonic if gives full measure to some set $\mathcal{M}$ of manifolds with uniformly bounded geometry and is invariant under re-rooting by Brownian motion, i.e. one has

$$\int f(M,o,g)\mathrm{d}\mu(M,o,g) = \int f(M,x,g)q(t,o,x)\mathrm{d}x\mathrm{d}\mu(M,o,g)$$

for all $t > 0$ and all bounded measurable functions $f : \mathcal{M} \to \mathbb{R}$.

In order for the above equation to make sense one needs to know that the inner integral on the right hand side is Borel measurable on the Gromov space. We prove this in the following lemma together with further regularity properties which will be useful to construct harmonic measures. The key point is that the heat kernel depends continuously on the manifold in the smooth topology. This intuitive fact was used (for time-dependent metrics) by Perelman in his proof of the geometrization conjecture after which it has received careful treatment by several authors (see [Lu12] and the references therein). It had also been previously used by Lucy Garnett to prove the existence of harmonic measures on foliated spaces which we will consider in the next subsection (see [Gar83, Fact 1] and [Can03]).

Lemma 2.6. *Let $\mathcal{M}$ be a compact subset of the Gromov space consisting of manifolds of uniformly bounded geometry and for each $t > 0$,$r > 0$ and each function $f : \mathcal{M} \to \mathbb{R}$ define $P^t f$ and $P^t_r f$ on $\mathcal{M}$ by*

$$P^t_r f(M,o,g) = \int\limits_{B_r(o)} f(M,x,g)q(t,o,x)\mathrm{d}x$$

$$P^t f(M,o,g) = \int\limits_{M} f(M,x,g)q(t,o,x)\mathrm{d}x.$$

Then the following properties hold:

1. *If f is continuous then $P^t_r f$ is continuous.*
2. *If f is continuous then $P^t_r f$ converges uniformly to $P^t f$ when $r \to +\infty$. In particular, $P^t f$ is continuous.*

3. *If f is bounded and Borel measurable then $P^t f$ is also Borel measurable and satisfies* $\sup|P^t f| \le \sup|f|$.

Proof. We begin by showing that if f is continuous then $P_r^t f$ is as well.

Consider a manifold $(M, o, g) \in \mathcal{M}$ and a sequence $(M_n, o_n, g_n) \in \mathcal{M}$ converging to it. By Theorem 2.4 the convergence is smooth so there exists an exhaustion U_n of M by precompact open sets and a sequence of diffeomorphisms $\varphi_n : U_n \to V_n \subset M_n$ such that $\varphi_n(o) = o_n$ and the pullback metric $\varphi_n^* g_n$ converges smoothly on compact subsets to g.

In this situation Theorem 2.1 of [Lu12] (applied in the case where the fields X_n and the potentials Q_n are equal to 0 and the metrics $g_n(\tau)$ are constant with respect to τ) guarantees that the sequence of pullbacks $q_n(t, o, x) = q^{M_n}(t, o_n, \varphi(x))$ of the transition probability densities of Brownian motion on each M_n converges uniformly on compact subsets of $[0, +\infty) \times M$ to a fundamental solution $\tilde{q}(t, o, \cdot)$ of the heat equation (that fact that we use $\Delta/2$ instead of Δ is clearly inessential) which satisfies

$$\int\limits_M \tilde{q}(t, o, x)\mathrm{d}x \le \sup\left\{\int_{M_n} q^{M_n}(t, o_n, x)\mathrm{d}x\right\} = 1.$$

By Theorem 4.1.5 [Hsu02] the transition density $q(t, o, x)$ is the minimal fundamental solution so one has $q(t, o, x) \le \tilde{q}(t, o, x)$. Combined with the fact that the integral of both kernels with respect to x is at most 1 and the $\int q(t, o, x)\mathrm{d}x = 1$ one obtains $q(t, o, x) = \tilde{q}(t, o, x)$ so that $q_n(t, o, \cdot)$ converges uniformly on compact sets to $q(t, o, \cdot)$.

Setting $F(x) = f(M, x, g)$ and $F_n(x) = f(M_n, \varphi_n(x), g_n)$ and using the fact that f is uniformly continuous (because $\mathcal{M}$ is compact) one obtains that $F_n \to F$ uniformly on compact subsets of M.

Finally because the pullback metrics $\varphi_n^* g_n$ converge smoothly to g, the Jacobian J_n of φ_n converges uniformly to 1 on compact subsets and also the open sets $\Omega_n = \varphi_n^{-1}(B_r(o_n))$ converge in the Hausdorff distance to $B_r(o)$.

Combining these four facts (local uniform convergence of $q_n(t, o, \cdot)$ to $q(t, o, \cdot)$, F_n to F, J_n to 1, and Hausdorff convergence of Ω_n to $B_r(o)$) with the fact that f is bounded one obtains

$$\begin{aligned}
\lim_{n\to+\infty} P_r^t f(M_n, o_n, g_n) &= \lim_{n\to+\infty} \int\limits_{B_r(o_n)} q(t, o_n, x) f(M_n, x, g_n)\mathrm{d}x \\
&= \lim_{n\to+\infty} \int\limits_{\Omega_n} q_n(t, o, x) F_n(x) J_n(x)\mathrm{d}x \\
&= \lim_{n\to+\infty} \int\limits_{B_r(o)} q_n(t, o, x) F_n(x) J_n(x)\mathrm{d}x \\
&= \lim_{n\to+\infty} \int\limits_{B_r(o)} q(t, o, x) F(x)\mathrm{d}x \\
&= \lim_{n\to+\infty} \int\limits_{B_r(o)} q(t, o, x) f(M, x, g)\mathrm{d}x \\
&= P_r^t f(M, o, g)
\end{aligned}$$

which implies that $P_r^t f$ is continuous as claimed.

We will now show that $P_r^t f$ converges uniformly to $P^t f$ when $r \to +\infty$ if f is continuous.

By Theorem 2.5 for each fixed $t > 0$ there exists a constant C such that

$$q(t, o, x) \le C \exp\left(-d(o,x)^2/C\right)$$

for all manifolds $(M, o, g) \in \mathcal{M}$ and all $x \in M$. Furthermore by the Bishop comparison theorem one may increase C above so that the volume of the ball of radius r is bounded from above by $\exp(Cr)$.

Combining these two facts one obtains that

$$\begin{aligned}
|(P_r^t - P^t)f(M,o,g)| &\le \int\limits_{d(o,x)>r} |f(M,x,g)|q(t,o,x)\mathrm{d}x \\
&\le C \sup|f| \int\limits_{d(o,x)>r} \exp\left(-d(o,x)^2/C\right)\mathrm{d}x \\
&\le C \sup|f| \sum_{n=1}^{+\infty} \exp(-(nr)^2/C + C(n+1)r) \\
&\le C \sup|f| \sum_{n=1}^{+\infty} \exp(-(nr)^2/C + C(n+1)r) \\
&= C \sup|f| \varphi(r)
\end{aligned}$$

where the last inequality is obtained by bounding the integral by the sum over anulii of the form $B_{(n+1)r}(o) \setminus B_{nr}(o)$ and the integral over each anulus by the maximum value time the volume of the ball of radius $(n+1)r$.

As soon as $r > 2C^2$ one has that $C(n+1)r - n^2r^2/C$ is decreasing with respect to r for all $n \ge 1$. Hence $\varphi(r) \to 0$ when $r \to +\infty$ and one obtains that $P_r^t f$ converges uniformly to $P^t f$ as claimed.

To conclude we will prove that $P^t f$ is Borel for all bounded Borel f and that it is bounded in absolute value by $\sup|f|$.

For this purpose consider for some $C > 0$ the family of functions $\mathcal{F}$ on $\mathcal{M}$ bounded in absolute value by C and such that $P^t f$ is Borel measurable. We have shown that $\mathcal{F}$ contains the continuous functions bounded in absolute value by C. By the dominated convergence theorem it is closed by pointwise limits (this is because Borel functions are the smallest class containing continuous functions and closed under pointwise limits, see for example [Kec95, Theorem 11.6]). Therefore it contains all Borel measurable functions bounded by C in absolute value. Since this works for all C one has that $P^t f$ is Borel measurable for all bounded Borel measurable f. The claim $\sup|P^t f| \le \sup|f|$ follows directly from the definition of $P^t f$ because one has $\int q(t,x,y)\mathrm{d}y = 1$ on all manifolds in $\mathcal{M}$. □

Here are three simple examples of harmonic measure:

1. The measure giving full mass to a single Riemannian manifold whose isometry group acts transitively (e.g. a Lie group with a left invariant Riemannian metric) is harmonic.

2. If (M, g) is a compact Riemannian manifold and x is a random point in M whose distribution is given by normalized volume measure, then the distribution of (M, x, g) is harmonic.

3. Let (M, g) be a Riemannian manifold admitting a discrete group of isometries G acting with a finite volume fundamental domain M_0. If x is a random point in M_0 whose distribution is given by normalized volume measure, then the distribution of (M, x, g) is harmonic.

The following theorem implies that one can associate at least one harmonic measure to each manifold of bounded geometry. By this we mean that if (M, g) has bounded geometry then the closure in the Gromov space of the set of pointed manifolds of the form (M, x, g) supports at least one harmonic measure. We will see other examples of harmonic measures in the next subsection.

Theorem 2.7. *If $\mathcal{M}$ is a compact subset of the Gromov space consisting of manifolds with uniformly bounded geometry then there exists at least one harmonic measure supported on $\mathcal{M}$.*

Proof. For each probability μ on $\mathcal{M}$ and $t > 0$ define the measure $P^t\mu$ on $\mathcal{M}$ using the Riesz representation theorem and Lemma 2.6 in such way that for all continuous $f : \mathcal{M} \to \mathbb{R}$ one has

$$\int_{\mathcal{M}} f \mathrm{d}P^t\mu = \int\limits_{\mathcal{M}} P^t f \mathrm{d}\mu.$$

The maps $\{P^t : t \geq 0\}$ form a commuting family of linear maps which leave the convex and weakly compact set of probability measures on $\mathcal{M}$ invariant. By the Markov-Kakutani fixed point theorem there is a common fixed point for all the P^t which must be a harmonic measure. $\square$

2.1.4 Foliations and leaf functions

Harmonic measures on foliations were introduced by Lucy Garnett in [Gar81] (see also [Gar83] and [Can03]). In this subsection we will explore how they relate to harmonic measures on the Gromov space in the sense of our definition.

To begin we must fix a definition of foliation. There are several definitions in the literature, the crucial feature for our purposes is that each leaf should be a Riemannian manifold. An important example is given by the foliation defined by an integrable distribution of tangent subspaces on a Riemannian manifold, in this case each leaf inherits a Riemannian metric from the ambient space.

A d-dimensional foliation is a metric space X partitioned into disjoint subsets called leaves. Each leaf is a continuously and injectively immersed d-dimensional connected complete Riemannian manifold. Furthermore, for each $x \in X$ there is an open neighborhood U, a Polish space T, and a homeomorphism $h : \mathbb{R}^d \times T \to U$ with the following properties:

1. For each $t \in T$ the map $x \mapsto h(x, t)$ is a smooth injective immersion of $\mathbb{R}^d$ into a single leaf.

2. For each $t \in T$ let g_t be the metric on $\mathbb{R}^d$ obtained by pullback under $x \mapsto h(x, t)$ of the corresponding leaf's metric. If a sequence t_n converges to $t \in T$ then the Riemannian metrics g_{t_n} converge smoothly on compact sets to g_t.

As part of a program to study the geometry of topologically generic leaves Álvarez and Candel introduced the 'leaf function' which is a natural function into the Gromov space associated to each foliation X (see [ÁC03]). It is defined as the function mapping each point x in the foliation to the leaf $L(x)$ containing it considered as a pointed Riemannian manifold with basepoint x.

We will now establish that the leaf function is Borel measurable. We do this by using a result of Solovay for which we must assume the existence of an inaccessible cardinal. The author believes a more direct proof without this assumption is attainable in the same vein as Theorem 4.3 of the second section where semicontinuity of the leaf function is established and related to Reeb type stability results.

Lemma 2.8. *Let X be a compact foliation and L its leaf function. Then the following holds:*

1. *L takes values in a compact subset of the Gromov space consisting of manifolds with uniformly bounded geometry.*

2. *L is measurable with respect to the completion of the Borel σ-algebra with respect to any probability measure on X.*

Proof. We will prove the the first claim in the second section of this thesis, see Theorem 4.1.

To establish the second claim suppose X is a compact foliation, μ is a probability on X and there exists an open set U in the Gromov space such that L^{-1} is not μ-measurable (where L is the leaf function of X).

By [dlR93, Théorème 4-3] (see also [Roh52]) there exists a full measure set $X' \subset X$ and a bi-measurable bijection $f : X' \to \mathbb{R}$ such that $f(X') = [0, m_0] \sqcup C$ where $m_0 \geq 0$ and C is a countable subset of $\mathbb{R}$ disjoint form $[0, m_0]$, such that $f_*\mu$ equals the sum of Lebesgue measure on $[0, m_0]$ with a probability measure on C.

If follows that $f(L^{-1}(U))$ is not Lebesgue measurable. And we have therefore constructed a non-Lebesgue measurable subset of $\mathbb{R}$ without using the axiom of choice.

Assuming the existence of an inaccessible cardinal this is not possible due to [Sol70, Theorem 1]. □

A probability measure m on a foliation X is said to be harmonic (see [Gar83, Fact 4]) if it satisfies

$$\int f(x)\mathrm{d}m(x) = \int q(t,x,y)f(y)\mathrm{d}y\mathrm{d}m(x)$$

for all bounded measurable functions $f : X \to \mathbb{R}$.

Every compact foliation admits at least one harmonic measure (see [Gar83] and [Can03]). The following theorem implies that any result which establishes properties of generic manifolds for harmonic measures on the Gromov space immediately implies a similar result for generic leaves of compact foliations.

Theorem 2.9. *Let X be a compact foliation with leaf function L and m a harmonic measure on X. Then the push-forward measure L_*m is harmonic measure on the Gromov space.*

Proof. Let $\mathcal{M}$ be a compact set of the Gromov space containing the image of L and consisting of manifolds with uniformly bounded geometry. If $f : \mathcal{M} \to \mathbb{R}$ is bounded and measurable then by definition of L_*m and harmonicity of m one has

$$\begin{aligned}
&\int\limits_{\mathcal{M}}\int\limits_{L(x)} q(t,x,y)f(M,y,g)\mathrm{d}y\mathrm{d}L_*m(M,x,g)\\
&\qquad = \int\limits_{X}\int\limits_{L(x)} q(t,x,y)f\circ L(y)\mathrm{d}y\mathrm{d}m(x)\\
&\qquad = \int\limits_{X} f\circ L(x)\mathrm{d}m(x) = \int\limits_{\mathcal{M}} f(M,x,g)\mathrm{d}L_*m(M,x,g)
\end{aligned}$$

so L_*m is harmonic as claimed. □

2.2 Asymptotics of random manifolds

2.2.1 Stationary random manifolds

We define a stationary random manifold to be a random element of the Gromov space whose distribution is a harmonic measure. A typical example is obtained as follows: let X be a compact foliation and m a harmonic measure on X, the leaf function $L : X \to \mathcal{GS}$ is a stationary random manifold defined on (X, m) (see 2.9). As noted in the previous section if a bounded geometry manifold (M, g) admits a finite volume quotient under isometries and one takes a random point o in a fundamental domain of this action distributed according to the normalized volume measure then (M, o, g) is a stationary random manifold. This poses the following question:

Question 2.1 (Andrés Sambarino)**.** Let (M, g) be a bounded geometry manifold and suppose there exists a random point o of M such that (M, o, g) is stationary. Does it follow that M admits a finite volume quotient by isometries?

We will next introduce and study several asymptotic quantities associated to each stationary random manifold M.

2.2.2 Entropy

We introduce an asymptotic quantity 'Kaimanovich entropy' associated to each stationary random manifold which measures the asymptotic behavior of the differential entropy between the time t distribution of Brownian motion and the Riemannian volume measure. Several alternate definitions for this quantity in different contexts as well as theorems and applications where announced in the interesting papers [Kaĭ86] and [Kaĭ88]. For the case of manifolds with a compact quotient some of these properties (e.g. the so-called Shannon-McMillan-Breiman type theorem) were later proved in [Led96].

The main point is that Kaimanovich entropy relates directly to the mutual information between σ-algebras $\mathcal{F}_t$ and $\mathcal{F}^T$ for Brownian motion. This allows one to show that Kaimanovich entropy is zero if and only if the random manifold satisfies the Liouville property almost surely. Later on we will will relate entropy to other asymptotic quantities.

The following technical lemma sets the basis for our study.

Lemma 2.10. *Let $\mathcal{M}$ be a compact subset of the Gromov space consisting of manifolds with uniformly bounded geometry. The following function is finite and continuous with* *respect to $t > 0$ and $M \in \mathcal{M}$:*

$$h_t(M) = -\int_M q(t, o_M, x) \log(q(t, o_M, x))\mathrm{d}x,$$

$$I_t^T(M) = \int_{M\times M} \log\left(\frac{q(T-t, x, y)}{q(T, o_M, y)}\right) q(t, o_M, x)q(T-t, x, y)\mathrm{d}x\mathrm{d}y.$$

Also, the following formula holds (where P^t is defined by Lemma 2.6):

$$I_t^T(M) = h_T(M) - (P^t h_{T-t})(M).$$

Proof. The proof is similar to that of Lemma 2.6. We define

$$h_t^r(M) = -\int_{B_r(o_M)} q(t, o_M, x) \log(q(t, o_M, x))\mathrm{d}x$$

with the purpose of showing that h_t^r is continuous and converges uniformly to h_t on $\mathcal{M}$ when $r \to +\infty$.

Assume $\mathcal{M}$ consists of n-dimensional manifolds with curvature greater than $-k^2$ and let $K(t, r)$ be the heat kernel at time t between two points at distance r in the n dimensional hyperbolic plane with constant curvature $-k^2$. Then by Theorem 2.2 of [Ich88] one has $q(t, o_M, x) \geq K(t, d(o_M, x))$ for all $M \in \mathcal{M}$. From [DM88, Theorem 3.1] one obtains that $\log(K(t, r))$ is bounded from below by a polynomial f in r on any compact interval of positive times.

On the other hand by Theorem 2.5 for each compact interval of times there exists $C > 0$ such that one has $q(t, o_M, x) \leq Ce^{-d(o_M,x)^2/3t}$ on all M of $\mathcal{M}$.

Combining these facts yields

$$-\log(C) \leq -\log(q(t, o_M, x)) \leq f(d(o_M, x))$$

for all $M \in \mathcal{M}$.

By Bishop's volume comparison theorem the volume of the ball of radius r in any manifold of $\mathcal{M}$ is bounded by that in n-dimensional hyperbolic space with curvature $-k^2$. Hence one has an upper bound for volume of the form $\mathrm{vol}(B_r(o_M)) \leq \exp(Cr)$ (notice that the previous inequality remains valid if one increases C so there is no problem in using the same constant for both bounds).

Similarly to the proof of Lemma 2.6 one obtains for each r a positive constant $\epsilon(r)$ which decreases to 0 as $r \to +\infty$ such that

$$|h_t(M) - h_t^r(M)| \leq \int_{M\setminus B_r(o_M)} q(t, o_N, x)|\log(q(t, o_M, x))|\mathrm{d}y \leq \epsilon(r)$$

for all manifolds in $\mathcal{M}$.

Hence $h_t(M)$ is the uniform limit on $\mathcal{M}$ of $h_t^r(M)$ when $r \to +\infty$ and it suffices to establish continuity of the later.

Before doing that we establish continuity with respect to t of h_t. Given $M \in \mathcal{M}$, $\epsilon > 0$ and $t > 0$, one can find $r > 0$ such that $|h_s - h_s^r| < \epsilon/3$ for all s in a compact neighborhood of t (notice that our bounds were obtained uniformly on such intervals). Since $q(t, o_M, x)$ is continuous with respect to t and x one has that $q(s, o_M, \cdot)$ converges uniformly to $q(t, o_M, \cdot)$ on $B_r(o_M)$ when $s \to t$. Hence $h_s^r(M) \to h_t^r(M)$. Combining the two facts one obtains that there exists a neighborhood of t on which

$$|h_s(M) - h_t(M)| \leq |h_s(M) - h_s^r(M)| + |h_s^r(M) - h_t^r(M)| + |h_t^r(M) - h_t(M)| < \epsilon$$

which yields continuity of h_t with respect to t as claimed.

We now establish continuity of $h_t^r(M)$ with respect to M. Assume the sequence (M_n, o_n, g) in $\mathcal{M}$ converges to (M, o, g). By Theorem 2.4 convergence is smooth so that there exists an exhaustion U_n of M by precompact open sets and a sequence of diffeomorphisms $\varphi_n : U_n \to V_n \subset M_n$ such that $\varphi_n(o) = o_n$ and the pullback metric $\varphi_n^* g_n$ converges smoothly on compact subsets to g.

From this it follows that the Jacobian $J(x)$ of φ_n at x converges to 1 uniformly on compact sets. And by the results of [Lu12] the functions $q_n(t, o, x) = q^{M_n}(t, o_n, \varphi_n(x))$ (q^{M_n} being the transition density of Brownian motion on M_n) converge uniformly on compact sets to the transition density $q(t, o, x)$ of Brownian motion on M. Using this the continuity of $h_t^r(M)$ follows as claimed.

We will establish the formula $I_t^T = h_T - P^t h_{T-t}$, from this the continuity of I_t^T follows from that of h_T and h_{T-t} by Lemma 2.6. The proof is the following computation using the property $\int q(t, x, y) q(s, y, z) \mathrm{d}y = q(t + s, x, z)$ of the heat kernel:

$$\begin{aligned}
I_t^T(M) &= \int \log \left(\frac{q(T - t, x, y)}{q(T, o, y)} \right) q(t, o, x) q(T - t, x, y) \mathrm{d}x \mathrm{d}y \\
&= \int \log \left(q(T - t, x, y) \right) q(t, o, x) q(T - t, x, y) \mathrm{d}x \mathrm{d}y \\
&- \int \log \left(q(T, o, y) \right) q(t, o, x) q(T - t, x, y) \mathrm{d}x \mathrm{d}y \\
&= - \int q(t, o, x) h_{T-t}(M, x, g) \mathrm{d}x - \int \log \left(q(T, o, y) \right) q(T, o, y) \mathrm{d}y \\
&= -P^t h_{T-t}(M) + h_T(M).
\end{aligned}$$

□

Recall that we say a Riemannian manifold is Liouville if it admits no non-constant bounded harmonic functions.

Theorem 2.11. *The following limit (Kaimanovich entropy) exists and is non-negative for any stationary random manifold M:*

$$h(M) = \lim_{t \to +\infty} \mathbb{E} \left(\frac{1}{t} h_t(M) \right).$$

Furthermore, $h(M) = 0$ if and only if M is almost surely Liouville.

Proof. Let $H_t = \mathbb{E}(h_t(M))$ and notice that by dominated convergence it is continuous with respect to $t > 0$, bounded by the maximum of h_t on $\mathcal{M}$, and by Lemma 2.10 one has

$$H_T - H_{T-t} = \mathbb{E}(h_T(M) - h_{T-t}(M)) = \mathbb{E}(h_T(M) - P^t h_{T-t}(M)) = \mathbb{E}(I_t^T(M)).$$

The mutual information I_t^T is non-negative and decreases to $I_t^\infty(M)$ when $T \to +\infty$ (see Theorem 1.10 and preceding paragraphs). By the monotone convergence theorem it follows that $T \mapsto H_T - H_{T-t}$ decreases to $\mathbb{E}(I_t^\infty(M)) \geq 0$ when $T \to +\infty$. From this one obtains that

$$h(M) = \lim_{T \to +\infty} \mathbb{E}(\frac{1}{T} h_T(M)) = \mathbb{E}(\frac{1}{t} I_t^\infty(M))$$

for all $t > 0$.

If N is a manifold with bounded geometry then by Lemma 2.10 one has that $I_t^T(N)$ is finite. It follows from Theorem 1.10 that $I_t^\infty(N) = 0$ for some $t > 0$ if and only if N is Liouville. Hence $h(M) = 0$ if and only if M is almost surely Liouville as claimed. □

2.2.3 Linear Drift

Given a pointed bounded geometry manifold (M, o, g) we define by

$$\ell_t(M) = \int d(o,x)q(t,o,x)\mathrm{d}x$$

the mean displacement of Brownian motion starting at o from its starting point. We are interested in the asymptotics of $\ell_t(M)$ when $t \to +\infty$ and M is a stationary random manifold. In particular we introduce the linear drift of a random manifold M via the following theorem.

Theorem 2.12. *Let M be a stationary random manifold. Then the linear drift*

$$\ell(M) = \lim_{t\to+\infty} \mathbb{E}\left(\frac{1}{t}\ell_t(M)\right)$$

exists and is finite.

Proof. Suppose that M takes values in a set of manifolds $\mathcal{M}$ with uniformly bounded geometry. We begin by establishing that $\ell_t(M)$ is continuous with respect to both M and t.

If (M_n, o_n, g_n) is a sequence in $\mathcal{M}$ converging to (M, o, g) then by Theorem 2.4 there exists an exhaustion U_n of M by relatively compact open sets and smooth embeddings $\varphi_n : U_n \to M_n$ with $\varphi_n(o) = o_n$ such that $\varphi_n^* g_n$ converges smoothly g on compact subsets of M.

It follows that the Jacobian J_n of φ_n converges uniformly to 1 on compact sets and $d_n(x) = d(o_n, \varphi_n(x))$ converges uniformly on compact sets to $d(o,x)$. Also, by Theorem 2.1 of [Lu12], one has that $q_n(t,x) = q(t, o_n, \varphi_n(x))$ converges uniformly on compact sets to $q(t,o,x)$.

Combining these facts one obtains that for each $r > 0$

$$\ell_t^r(M) = \int\limits_{B_r(o)} q(t,o,x)d(o,x)\mathrm{d}x$$

depends continuously on $(M, o, g) \in \mathcal{M}$.

Let $A, B > 0$ be given by Lemma 2.14 below. By Jensen's inequality one has for all $M \in \mathcal{M}$ that

$$\begin{aligned}\left(\ell(M) - \ell_t^{At}(M)\right)^2 &= \left(\int\limits_{M\setminus B_{At}(o)} q(t,o,x)d(o,x)\mathrm{d}x\right)^2 \\ &\le \int\limits_{M\setminus B_{At}(o)} q(t,o,x)d(o,x)^2\mathrm{d}x \le Be^{-At}t^2\end{aligned}$$

which establishes that ℓ_t^r converges uniformly to ℓ_t on $\mathcal{M}$ when $r \to +\infty$ for all $t \ge 1$. In particular ℓ_t is continuous with respect to M on $\mathcal{M}$ for $t \ge 1$ (in fact this is true for all t but we will not need it).

To establish continuity with respect to t assume $t, s > 1$ and notice that using Lemma 2.14 as above one obtains that for all $T > \max(s,t)$ the integrals of both $d(o,x)q(t,o,x)$ and $d(o,x)q(s,o,x)$ on $B_{AT}(o)$ are bounded by $Be^{-At}t^2$. Combining this with the fact that if $s \to t$ then $q(s, o, \cdot)$ converges uniformly on $B_{AT}(o)$ to $q(t, o, \cdot)$ yields the desired result.

It follows that if $s \to t > 1$ then $\ell_s(M)$ converges uniformly on $\mathcal{M}$ to $\ell_t(M)$ and hence

$$L_t = \mathbb{E}(\ell_t(M))$$

is continuous with respect to $t \geq 1$.

We will now establish that L_t is subadditive, i.e. satisfies $L_{t+s} \leq L_t + L_s$, from which the existence of the finite limit $\lim L_t/t$ follows [1].

For this purpose we calculate using the triangle inequality

$$\begin{aligned}
\ell_{t+s}(M) &= \int q(t+s, o, x)d(o, x)\mathrm{d}x \\
&= \int q(t, o, y)q(s, y, x)d(o, x)\mathrm{d}x\mathrm{d}y \\
&\leq \int q(t, o, y)q(s, y, x)(d(o, y) + d(y, x))\mathrm{d}x\mathrm{d}y = \ell_t(M) + P^t\ell_s(M)
\end{aligned}$$

which taking expectation yields the desired result. □

Define $\ell^+(M)$ for a stationary random manifold M as the infimum of all $L > 0$ such that

$$\lim_{t \to +\infty} \mathbb{E}\left(\int\limits_{B_{Lt}(o)} q(t, o, x)\mathrm{d}x\right) = 1$$

we wish to guarantee that $\ell(M) = \ell^+(M)$.

However, a counterexample is given by a random manifold M which is equal to the hyperbolic plane with constant curvature -1 or -2 each with probability $1/2$. In this example $\ell(M)$ equals 1.5 but $\ell^+(M)$ equals 2. The problem arises because the distribution of M is a convex combination of other harmonic measures (in this case Dirac deltas). We say a harmonic measure on a compact set $\mathcal{M}$ of manifolds with uniformly bounded geometry is ergodic if it is extremal among all harmonic measures on this set. A random manifold is said to be ergodic if its distribution is.

Lemma 2.13. *Let M be a stationary random manifold. Then $\ell(M) \leq \ell^+(M)$. Furthermore, if M is ergodic then $\ell(M) = \ell^+(M)$.*

Proof. Let A and B be given by Lemma 2.14. For any $L > 0$ one has

$$\frac{1}{t}\ell_t(M) \leq L + A \int\limits_{M \setminus B_{Lt}(o)} q(t, o, x)\mathrm{d}x + \int\limits_{M \setminus B_{At}(o)} \frac{d(o, x)}{t} q(t, o, x)\mathrm{d}x.$$

Using Jensen's inequality and Lemma 2.14 one obtains that the third term is bounded from above by $\sqrt{B}e^{-At/2}$. This implies in particular that $\ell^+(M) \leq A$. If $L > \ell^+(M)$ then the expectation of the second term goes to zero from which one obtains that $\ell(M) \leq \ell^+(M)$.

Proof of the converse inequality will be postponed until the next chapter (see Corollary 3.7). □

1. Notice that there exist subadditive functions S_t for which S_t/t has no limit. One such example can be obtained by considering a $\mathbb{Q}$-linear extension of the function defined for all rational a, b by $S_{a+b\sqrt{2}} = a + 2b\sqrt{2}$. This shows that continuity of L_t is essential to the proof.

The main result of [Ich88] is that one can compare the radial process of Brownian motion on a Riemannian manifold with that of a model space with constant curvature. Combining this result with upper bounds for the heat kernel yields the following technical lemma which we have used in the proofs above.

Lemma 2.14. *Let $\mathcal{M}$ be a compact subset of the Gromov space consisting manifolds with uniformly bounded geometry. Then there exist constants $A, B > 0$ such that for each Brownian motions X_t starting at the origin o of a manifold (M, o, g) in $\mathcal{M}$ one has*

$$\mathbb{E}\left(\frac{d(o,X_t)^2}{t^2}1_{\{d(o,X_t)>At\}}\right) \le Be^{-At}$$

for all $t \ge 1$.

Proof. Let n denote the dimension of the manifolds in $\mathcal{M}$ and let $-k^2$ be a lower bound for their curvature.

Letting Y_t be a Brownian motion starting at the origin of $\mathbb{R}^n$ endowed with a complete metric of constant curvature $-k^2$ one has by [Ich88, Theorem 2.1] that if X_t is any Brownian motion starting at the origin of a manifold in $\mathcal{M}$ then

$$\mathbb{P}(d(Y_0, Y_t) > x) \ge \mathbb{P}(d(X_0, X_t) > x)$$

for all x.

Notice that if V, W are non-negative random variables with $\mathbb{P}(V > x) \ge \mathbb{P}(W > x)$ for all x then $\mathbb{E}(V) \ge \mathbb{E}(W)$. In particular for any non-decreasing non-negative function f one has $\mathbb{E}(f(V)) \ge \mathbb{E}(f(W))$.

Applying this observation one obtains

$$\mathbb{E}\left(\frac{d(Y_0,Y_t)^2}{t^2}1_{\{d(Y_0,Y_t)>At\}}\right) \ge \mathbb{E}\left(\frac{d(X_0,X_t)^2}{t^2}1_{\{d(X_0,X_t)>At\}}\right)$$

for all $A > 0$ and all $t \ge 0$. So that it suffices to bound the expectation on the left hand side.

Letting $K(t, r)$ denote the probability transition density Brownian motion on the hyperbolic plane with constant curvature $-k^2$. One has explicitly

$$\mathbb{E}\left(d(Y_0,Y_t)1_{\{d(Y_0,Y_t)>At\}}\right) = \int_{At}^{+\infty} \frac{r^2}{t^2}\mathrm{vol}(S^{n-1})\frac{1}{k}\sinh(kr)^{n-1}K(t,r)\mathrm{d}r,$$

where S^{n-1} is the standard $n-1$-dimensional sphere

By Theorem 2.5 there exists a constant C such that one has $K(t, r) \le Ce^{-r^2/3t}$ for all $t \ge 1$. Applying this, and bounding $\sinh(r)$ by e^r, one obtains

$$\mathbb{E}\left(d(Y_0,Y_t)1_{\{d(Y_0,Y_t)>At\}}\right) \le \frac{C\mathrm{vol}(S^{n-1})}{k}\int_{At}^{+\infty} \frac{r^2}{t^2}e^{(n-1)kr-r^2/3t}\mathrm{d}r$$

which bounding $e^{-r^2/3t}$ by $e^{-Ar/3}$ and choosing $A = 3nk$ yields

$$\cdots \le \frac{C\mathrm{vol}(S^{n-1})}{kt^2}\int_{At}^{+\infty} r^2e^{-r}\mathrm{d}r = \frac{C\mathrm{vol}(S^{n-1})}{kt^2}(A^2t^2+2At+2)e^{-At}$$

which, choosing B appropriately, yields the desired bound for $t \ge 1$. □

2.2.4 Inequalities

For a manifold M with a compact quotient it can be shown that the limit

$$\lim_{r\to+\infty} \frac{1}{r}\log(\mathrm{vol}(B_r(x)))$$

exists and has the same value for all $x \in M$. If, for some Riemannian manifold possibly without a compact quotient, this limit above is zero then we say M has subexponential growth.

We define the volume growth of a stationary random manifold M as

$$v(M) = \liminf_{r\to+\infty} \mathbb{E}\left(\frac{1}{r}\log(\mathrm{vol}(B_r(o_M)))\right).$$

By Bishop's inequality one has a uniform exponential upper bound on the volume of the ball of radius r on any set of manifolds with uniformly bounded geometry. This implies by dominated convergence that M has subexponential growth almost surely then $v(M) = 0$. On the other hand a uniform lower bound on volume is given by the fact that M takes values in a space of manifolds with uniformly bounded geometry. Hence, by Fatou's Lemma one has that if $v(M) = 0$ then M satisfies

$$\liminf_{r\to+\infty} \frac{1}{r}\log(\mathrm{vol}(B_r(o_M))) = 0$$

almost surely.

We will bound entropy of a stationary random manifold from above and below in terms of its linear drift and volume growth. The upper bound was announced by Kaimanovich in the case of manifolds with a compact quotient (see [Kaĭ86, Theorem 6]). A sharper version of the lower bound, also for a manifolds with a compact quotient, was established by Ledrappier (see [Led10]). Analogous results for random walks on discrete groups also exist and have been sucessively improved by several authors, some of the first of these can be attributed to Varopoulos, Carne, and Guivarc'h (see [GMM12] and the references therein). An analogous theorem for stationary random graphs is due to Benjamini and Curien (see [BC12, Proposition 3.6]).

Theorem 2.15. *For all ergodic stationary random manifolds M the following holds:*

$$\frac{1}{2}\ell(M)^2 \le h(M) \le \ell(M)v(M).$$

Proof. We begin with the lower bound.

For this purpose fix $D > 2$ and let C be given by Theorem 2.5 so that

$$q(t,x,y) \le C\exp\left(-\frac{d(x,y)^2}{Dt}\right)$$

holds on all manifolds in the range of M for all $t \ge 1$.

Using this upper bound and Jensen's inequality we obtain

$$\begin{aligned}\frac{1}{t}h_t(M) &= -\frac{1}{t}\int q(t,o,x)\log(q(t,o,x))\mathrm{d}x\\ &\geq \frac{1}{t}\int q(t,o,x)(\frac{d(o,x)^2}{Dt} - \log(C))\mathrm{d}x\\ &= \int q(t,o,x)(\frac{d(o,x)}{\sqrt{Dt}})^2\mathrm{d}x - \frac{\log(C)}{t}\\ &\geq \left(\int q(t,o,x)\frac{d(o,x)}{\sqrt{Dt}}\mathrm{d}x\right)^2 - \frac{\log(C)}{t}\\ &= \frac{1}{D}(\frac{1}{t}\ell_t(M))^2 - \frac{\log(C)}{t}.\end{aligned}$$

Taking expectation and using Jensen's inequality once more one obtains

$$\mathbb{E}(\frac{1}{t}h_t(M)) \geq \frac{1}{D}\mathbb{E}\left((\frac{1}{t}\ell_t(M))^2\right) - \frac{\log(C)}{t} \geq \frac{1}{D}\mathbb{E}(\frac{1}{t}\ell_t(M))^2 - \frac{\log(C)}{t}$$

which by taking limit with $t \to +\infty$ yields

$$h(M) \geq \frac{1}{D}\ell(M)^2.$$

Letting D decrease to 2 one obtains $h(M) \geq \frac{1}{2}\ell(M)^2$ as claimed.

For the lower bound let $K(t,r)$ be the transition density of Brownian motion on n-dimensional hyperbolic space of constant curvature $-k^2$ where we assume all manifolds in the range of M have curvature greater than or equal to $-k^2$ and dimension n. By [Ich88, Theorem 2.2] one has

$$q(t,o_M,x) \geq K(t,d(o_M,x))$$

which combined with the upper bound given by Theorem 2.5 yields

$$|\log(q(t,o_M,x))| \leq \max(\log(C),\log(K(t,d(o_M,x)))$$

for some constant $C > 0$ depending only on $\mathcal{M}$, and all $t \geq 1$.

The density $K(t,r)$ is obtained by evaluating the heat kernel of hyperbolic space (curvature -1) at $t/2$ and r/k. Hence from the lower bounds for the hyperbolic heat kernel given by [DM88, Theorem 3.1] one obtains constants $a,b,c > 0$ depending only on $\mathcal{M}$ such that

$$|\log(q(t,o_M,x))| \leq a + b\log(t) + c(t + r + t/r)$$

for all $t \geq 1$.

Let A, B be given by Lemma 2.14. We obtain uniformly over the range of M (setting $r(x) = d(o_M,x)$) that

$$\begin{aligned}&\frac{1}{t}\int\limits_{M\setminus B_{At}} q(t,o_M,x)|\log(q(t,o_M,x))|\mathrm{d}x\\ &\qquad\leq \frac{a}{t} + \frac{b\log(t)}{t} + c\int\limits_{M\setminus B_{At}} q(t,o_M,x)(1 + \frac{r(x)}{t} + \frac{r(x)^2}{t^2})\mathrm{d}x\\ &\qquad\leq \frac{a}{t} + \frac{b\log(t)}{t} + c\int\limits_{M\setminus B_{At}} q(t,o_M,x)(\frac{r(x)^2}{A^2t^2} + \frac{r(x)^2}{At^2} + \frac{r(x)^2}{t^2})\mathrm{d}x\\ &\qquad\leq \frac{a}{t} + \frac{b\log(t)}{t} + c(A^{-2} + A^{-1} + A)Be^{-At}\end{aligned}$$

for all $t \geq 1$.

This immediately implies (taking expectation and limit) that

$$h(M) = \lim_{t \to +\infty} \mathbb{E}\left(-\frac{1}{t} \int\limits_{B_{At}(o_M)} q(t, o_M, x) \log(q(t, o_M, x)) \mathrm{d}x \right).$$

The same upper bound on $|\log(q(t, o_M, x)|$ can now be applied in the ball of radius At (where $r(x)/t$ can be bounded by A before integrating). One obtains that for any $L \leq A$ letting U_t be the annulus between radii Lt and At centered at o_M one has

$$\frac{1}{t}\int\limits_{U_t} q(t, o_M, x)|\log(q(t, o_M, x))|\mathrm{d}x \leq \frac{a}{t} + \frac{b\log(t)}{t} + c(1 + A + A^2) \int\limits_{M \setminus B_{Lt}(o_M)} q(t, o_M, x)\mathrm{d}x$$

for all $t \geq 1$.

Assuming $L > \ell(M)$ this last inequality implies

$$h(M) = \lim_{t \to +\infty} \mathbb{E}\left(-\frac{1}{t} \int\limits_{B_{Lt}(o_M)} q(t, o_M, x) \log(q(t, o_M, x)) \mathrm{d}x \right).$$

To conclude notice that $\varphi(z) = -z\log(z)$ is convex on $z > 0$. Setting $v_t = \mathrm{vol}(B_{Lt}(o_M))$ and letting p_t be the integral of $q(t, o_M, x)$ over x in $B_{Lt}(o_M)$ one obtains using Jensen's inequality applied to normalized volume on the ball that

$$v_t \frac{1}{v_t} \int\limits_{B_{Lt}(o_M)} \varphi(q(t, o_M, x))\mathrm{d}x \leq v_t\varphi(p_t/v_t) = p_t\log(v_t) - p_t\log(p_t) \leq \log(v_t).$$

Taking expectation now yields

$$h(M) \leq \liminf_{t \to +\infty} \mathbb{E}\left(\frac{1}{t}\log(\mathrm{vol}(B_{Lt}(o_M)))\right) = Lv(M)$$

for all $L > \ell(M)$. Which letting L decrease to $\ell(M)$ proves the claimed upper bound. □

Recall that $\mathbb{R}^2$ endowed with the metric $\mathrm{d}s^2 = \mathrm{d}x^2 + (1 + x^2)^2\mathrm{d}y^2$ has subexponential volume growth but admits the bounded harmonic function $\arctan(x)$ (this example was attributed to O. Chung by Avez). Avez proved in 1976 that for manifolds with a transitive isometry group such an example is impossible (see [Ave76]).

Corollary 2.16 (Avez)**.** *If (M, g) is a connected Riemannian manifold with subexponential volume growth whose isometry group acts transitively then M satisfies the Liouville property.*

A generalization of Avez's result to manifolds admitting a compact quotient under isometries was obtained by Varopolous (see [Var86, Theorem 3]).

Corollary 2.17 (Varopoulos)**.** *If (M, g) is a Riemannian manifold with subexponential volume growth which admits a compact quotient under isometries then M satisfies the Liouville property.*

An analogous result to the previous two for generic leaves of compact foliations (with respect to any harmonic measure) was announced by Kaimanovich and also follows from our theorem above (see Theorem 2 of [Kaĭ88] and the comments on page 307). A particular interesting case is the horospheric foliation on the unit tangent bundle of a compact negatively curved manifold.

Corollary 2.18 (Kaimanovich)**.** *If M is a negatively curved compact Riemannian manifold then almost every horosphere with respect to any harmonic measure for the horospheric foliation on the unit tangent bundle T^1M satisfies the Liouville property.*

Another interesting consequence of Theorem 2.15 is that $h(M) = 0$ if and only if $\ell(M) = 0$. The following case was established by Karlsson and Ledrappier using a discretization procedure to reduce the proof to an analysis of a random walk on a discrete group (see [KL07]).

Corollary 2.19 (Karlsson-Ledrappier)**.** *Let (M, g) be a manifold with bounded geometry that admits a compact quotient under isometries. Then the M satisfies the Liouville property if and only if its Brownian motion is non-ballistic (i.e. linear drift is zero).*

One might be tempted to conjecture that a stationary random manifold of exponential growth must have non-constant bounded harmonic functions. A counterexample is provided by Thurston's Sol-geometry (see [LS84, pg. 304] or consider the case with $p = q = 1$ and drift parameter $a = 0$ in the central limit theorem of [BSW12]).

Example 2.20 (Lyons-Sullivan)**.** Let $M = \mathbb{R}^3$ endowed with the Riemannian metric $\mathrm{d}s^2 = e^{2z}\mathrm{d}x^2 + e^{-2z}\mathrm{d}y^2 + \mathrm{d}z^2$. Then M has a transitive isometry group, exponential volume growth, and satisfies the Liouville property.

Chapter 3

Brownian motion on stationary random manifolds.

Introduction

In this chapter we construct Brownian motion on a stationary random manifold. The technical steps consist of defining the corresponding path space (where both the manifold and path can vary) and proving that the Weiner measures one has on the paths over each manifold vary with sufficient regularity to define a global measure in path space over any harmonic measure on the set of manifolds under consideration.

Once this object (Brownian motion on a stationary random manifold) is shown to exist several interesting consequences follow almost immediately.

First, applying Kingman's subadditive ergodic theorem (see [Kin68]), one obtains that linear drift can be defined pathwise, i.e. almost every path has a well defined rate of escape. This implies for example that linear drift of Brownian motion is well defined for almost every Brownian path on almost every leaf of a compact foliation (with respect to any harmonic measure). Also, it allows us to complete the proof of the fact, needed in the previous chapter, that the linear drift $\ell(M)$ of an ergodic stationary random manifold coincides with the growth rate of a ball containing almost all mass of the transition probability density $q(t, o, x)$ (see Lemma 2.13 and Corollary 3.7).

Second, using Birkhoff's ergodic theorem, one obtains that a non-compact ergodic stationary random manifold must almost surely contain infinitely many disjoint diffeomorphic copies of any finite radius ball. This is in the spirit of the much more detailed result of Ghys which reduces the possible topologies of non-compact generic leaves of a foliation by surfaces to six possible types (see [Ghy95]).

The original proof of Kingman's subadditive ergodic theorem (see [Kin68]) consisted in showing that for any stationary subadditive process $\{x_{s,t} : s < t\}$ with finite time constant there exists and additive process $\{y_{s,t} : s < t\}$ with the same time constant and satisfying $y_{s,t} \leq x_{s,t}$ almost surely. This reduces the subadditive ergodic theorem to Birkhoff's ergodic theorem. This 'subadditive decomposition' has been mostly forgotten (evidence for this is that, although the subadditive ergodic theorem has been reproved a great number of times, see for example [Ste89] and the references therein; the subadditive decomposition theorem has received only two additional proofs beside the original, see commentary by Burkholder in [Kin73] and the alternative proof by del Junco in [dJ77]).

A more or less independent sequence of works also relying on the idea of 'subadditive decomposition' exists beginning with Furstenberg's formula for the largest Lyapunov exponent of a sequence of random matrices (see [Fur63, Theorem 8.5] and [Led84, pg. 358]). Karlsson and Ledrappier have unified such results in a general framework by giving a formula expressing the rate of escape of random sequences in metric spaces using the expected increment of a Busemann function along the sequence (see [KL11, Theorem 18]). Notice that since Busemann functions are 1-Lipschitz their increments along a random sequence form an additive process which is dominated by the subadditive distance process of the sequence, hence these type of results amount to a more explicit version of the decomposition theorem for subadditive processes (where the additive processes has a natural geometric interpretation). In our context we obtain (see Theorem 3.14) a Furstenberg-type formula for the linear drift of Brownian motion on a stationary random manifold in terms of increments of a random Busemann function similar to [Led10, Proposition 1.1].

In the last subsection we improve the lower bound $\frac{1}{2}\ell(M)^2 \le h(M)$ for entropy obtained in Theorem 2.15 to $2\ell(M)^2 \le h(M)$ in the case of certain stationary random Hadamard manifold. In the case of a manifold with compact quotient this result was proved by Kaimanovich and Ledrappier, see [Kaĭ86, Theorem 10] and [Led10, Theorem A]. Equality implies that the gradient of almost every Busemann function at the origin must be collinear with that of a positive harmonic function, a condition which has strong rigidity consequences in the case of a single negatively curved manifold with compact quotient (see [Led10] and [LS12]). It is unknown to the author whether similar rigidity results can be obtained for stationary random manifolds.

3.1 Brownian motion on stationary random manifolds

3.1.1 Path space

We will construct a 'path space' over a given set $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ of manifolds with uniformly bounded geometry. Since later on we will be interested in time-reversal of Brownian motion we chose to consider paths whose domain is the entire real line instead of $[0, +\infty)$ as it was in previous chapters.

For this purpose let $\widehat{\mathcal{M}}'$ be the set of pairs (M, ω) where M is a manifold in $\mathcal{M}$ (denote by o_M its basepoint and by g_M its Riemannian metric) and $\omega : \mathbb{R} \to M$ is a continuous curve with $\omega_0 = o_M$. And by $\widehat{\mathcal{M}}$ denote the equivalence classes of elements of $\widehat{\mathcal{M}}'$ where (M, ω) and (M', ω') are equivalent if there is a pointed isometry $f : M \to M'$ such that $\omega' = f \circ \omega$. We will not be careful in distinguishing elements of $\widehat{\mathcal{M}}$ with their representatives (M, ω) in $\widehat{\mathcal{M}}'$ since all our definitions will be invariant under the defined equivalence relationship.

Recall that a metric on a disjoint union of two metric spaces is admissible if it coincides with the given metrics when restricted to each half. We mimick the definition of the distance on the Gromov space to turn $\widehat{\mathcal{M}}$ into a metric space.

Definition 3.1. Define the distance between two elements (M, ω) and (M', ω') in $\widehat{\mathcal{M}}$ as either $1/2$ or, if such an ϵ exists, the infimum among all $\epsilon \in (0, 1/2)$ such that that there exist an admissible metric on the disjoint union $M \sqcup M'$ with the following properties:

1. $d(o, o') < \epsilon$.

2. $d(B_{1/\epsilon}(o), M') < \epsilon$ and $d(M, B_{1/\epsilon}(o')) < \epsilon$.

3. $d(\omega_t, \omega'_t) < \epsilon$ for all $t \in [-1/\epsilon, 1/\epsilon]$.

Define for each $t \in \mathbb{R}$ the shift map $\mathrm{shift}^t : \widehat{\mathcal{M}} \to \widehat{\mathcal{M}}$ by

$$\mathrm{shift}^t(M, \omega) = (M, \mathrm{shift}^t \omega)$$

where $(\mathrm{shift}^t \omega)_s = \omega_{t+s}$ for all s and the basepoint of M on the right hand has been changed to ω_t (previously it was ω_0).

Also, one has a projection $\pi : \widehat{\mathcal{M}} \to \mathcal{M}$ which associates to each (M, ω) the unique pointed manifold in $\mathcal{M}$ isometric (with basepoint) to (M, ω_0).

$$\begin{tikzcd}
\widehat{\mathcal{M}} \arrow[loop right, "\mathrm{shift}^t"] \arrow[d, "\pi"'] \\
\mathcal{M}
\end{tikzcd}$$

Figure 3.1: The main objects of this section.

Recall that a sequence (M_n, o_n, g_n) of pointed manifolds is said to converge smoothly to (M, o, g) if there exists an exaustion U_n of M be increasing relatively compact open sets and a sequence of smooth embeddings $\varphi_n : U_n \to M_n$ with $\varphi_n(o) = o_n$ such that the pullback metrics $\varphi_n^* g_n$ converge smoothly on compact sets to g. We record the regularity properties of the spaces and arrows in the above diagram for future use.

Lemma 3.2. *Let $\mathcal{M}$ and $\widehat{\mathcal{M}}$ be as above. Then $\mathcal{M}$ is compact and metrizable when endowed with the topology of smooth convergence, $\widehat{\mathcal{M}}$ is a complete metric space with the distance defined above, and the projection $\pi : \widehat{\mathcal{M}} \to \mathcal{M}$ is continuous and surjective. Furthermore, for each $t \in \mathbb{R}$ the shift map shift^t is a self homeomorphism of $\widehat{\mathcal{M}}$.*

Proof. By Theorem 2.4 one has that $\mathcal{M}$ is compact with respect to pointed Gromov-Hausdorff convergence and that this convergence is equivalent to smooth convergence on $\mathcal{M}$. Since Gromov-Hausdorff convergence is metric this establishes the claim on $\mathcal{M}$.

The continuity of the projection follows because we have defined $\widehat{\mathcal{M}}$ to be larger than the pointed Gromov-Hausdorff distance between the projected manifolds. Continuity of the shift map can be verified directly from the definition. Hence it remains to establish that $\widehat{\mathcal{M}}$ is separable and complete.

We begin by establishing completeness. For this purpose take a sequence (M_n, ω^n) in $\widehat{\mathcal{M}}$ such that $d((M_n, \omega^n), (M_{n+1}, \omega^{n+1})) < 2^{-n}$ for all n. For each n let d_n be an admissible metric on the disjoint union $M_n \sqcup M_{n+1}$ satisfying the conditions of Definition 3.1 for $\epsilon = 2^{-n}$. Using these distances define a distance on the infinite disjoint union $X = \bigsqcup M_n$ such that if $x \in M_n$ and $y \in M_{n+p}$ the distance between them is the infimum over sequences $x = x_0, x_1, x_2, \ldots, x_p = y$ with $x_i \in M_{n+i}$ of $d_n(x_0, x_1) + \cdots + d_{n+p-1}(x_{n+p-1}, x_{n+p})$. Let $\widehat{X}$ be the completion of X, $M = \widehat{X} \setminus X$, and let ω be the local uniform limit of the curves ω^n in $\widehat{X}$.

One can verify that ω_t belongs to M for all t. Furthermore we have shown in Theorem 2.3 that the sequence M_n converges in the pointed Gromov-Hausdorff sense to M. It follows that M is isometric to a pointed manifold in $\mathcal{M}$. Since the distance between (M_n, ω^n) and (M, ω) is less than $2^{-(n-1)}$ this establishes that $\widehat{\mathcal{M}}$ is complete as claimed.

Take a countable dense subset D of $\mathcal{M}$ and consider a countable set $\widehat{D}$ consisting of pairs (M, ω) where M ranges over D and ω over a countable dense (with respect to local uniform convergence) subset of curves in each M. We claim that $\widehat{D}$ is dense in $\widehat{\mathcal{M}}$.

To establish this fact fix (M, ω) in $\widehat{\mathcal{M}}$ and $\epsilon > 0$. Since ω is continuous we may partition the interval $[-1/\epsilon, 1/\epsilon]$ into a finite number of disjoint intervals $I_i = [t_i, t_{i+1}]$ such that the diameter of $\{\omega_t : t \in I_i\}$ is less than $\epsilon/3$ for all i. Take $\delta \in (0, \epsilon/3)$ so that $1/\delta$ is larger than the diameter of $\{\omega_t : t \in [-1/\epsilon, 1/\epsilon]\}$ and notice that there exists (M', o', g') in D at distance less than than δ from (M, ω_0, g) (g being the metric on M). It follows that there exists an admissible distance on the disjoint union $M \sqcup M'$ such each point ω_{t_i} is at distance less than $\epsilon/3$ from some point $p_i \in M'$. Interpolating between the p_i with geodesic segments one obtains a curve ω'_t in M' at distance less than ϵ from ω_t for all $t \in [-1/\epsilon, 1/\epsilon]$. Arbitrarily close to $omega'$ in the topology of local uniform convergence there are curves ω'' such that (M', ω'') belongs to $\widehat{D}$. This proves that $\widehat{D}$ is dense as claimed. □

3.1.2 Brownian motion

By the topology of local uniform convergence on an interval $I \subset \mathbb{R}$ on $\widehat{\mathcal{M}}$ we mean the topology generated by replacing $t \in [-1/\epsilon, 1/\epsilon]$ in condition 3 of Definition 3.1 by $t \in [-1/\epsilon, 1/\epsilon] \cap I$. We denote by $\mathcal{F}_0^\infty$ the σ-algebra generated by the topology of local uniform convergence on $[0, +\infty)$ on $\widehat{\mathcal{M}}$.

Recall that $C(\mathbb{R}, M)$ denotes the space of continuous curves on a manifold M and denote by $\mathcal{F}_0^\infty(M)$ the σ-algebra on $C(\mathbb{R}, M)$ generated by the topology of local uniform convergence on $[0, +\infty)$. Each element ω of $C(\mathbb{R}, M)$ is naturally associated to the element (M, ω) of $\widehat{\mathcal{M}}$ though this association is not necessarily injective (e.g. applying any isometry to a curve yields the same element of $\widehat{\mathcal{M}}$).

Lemma 3.3. *The natural projection $\omega \mapsto (M, \omega)$ from the space of continuous curves $C(\mathbb{R}, M)$ on a manifold $M \in \mathcal{M}$ into $\widehat{\mathcal{M}}$ is continuous when both are endowed with the topology of local uniform convergence on the same interval $I \subset \mathbb{R}$. In particular each probability measure on $\mathcal{F}_0^\infty(M)$ yields a probability on $\left(\widehat{\mathcal{M}}, \mathcal{F}_0^\infty\right)$.*

Proof. The uniform distance between two curves ω, ω' in M on a compact interval $[a, b]$ is larger than or equal to the corresponding distance on $\widehat{\mathcal{M}}$ endowed with the topology of uniform convergence on $[a, b]$ (this follows directly from Definition 3.1). □

Given $(M, o, g) \in \mathcal{M}$ we denote by $P_{(M,o,g)}$ the pushforward of Wiener measure corresponding to Brownian motion starting at o to the σ-algebra $\mathcal{F}_0^\infty$ on $\widehat{\mathcal{M}}$. The Markov property of Brownian motion on M takes the following form:

Corollary 3.4. *Let $(M, o, g) \in \mathcal{M}$, then for any $t > 0$ and $A \in F_t^\infty$ one has*

$$\mathbb{P}_{(M,o,g)}(A) = \int q(t, o, x) P_{(M,x,g)}(A) \mathrm{d}x.$$

By a Brownian motion on a stationary random manifold we mean a random element (M, ω) of $\widehat{\mathcal{M}}$ with distribution $\widehat{\mu}$ given by the following theorem for some harmonic measure μ on $\mathcal{M}$.

Theorem 3.5. *For each probability measure μ on $\mathcal{M}$ there is a unique probability $\widehat{\mu}$ on $\left(\widehat{\mathcal{M}}, \mathcal{F}_0^\infty\right)$ defined by*

$$\widehat{\mu} = \int \mathbb{P}_{(M,o,g)} \mathrm{d}\mu(M, o, g).$$

Furthermore the following two properties hold

1. *If a probability μ on $\mathcal{M}$ is harmonic then its lift $\widehat{\mu}$ can be extended uniquely to a shift invariant Borel probability on $\widehat{\mathcal{M}}$.*

2. *If a probability μ on $\mathcal{M}$ is harmonic and ergodic then the unique shift invariant extension of $\widehat{\mu}$ is ergodic.*

Proof. We have shown in Lemma 3.3 that the Weiner measures on each (M, o, g) in $\mathcal{M}$ lifts to a probability $P_{(M,o,g)}$ on $\widehat{\mathcal{M}}$. If we show that

$$(M, o, g) \mapsto \int f(M, \omega)\mathrm{d}P_{(M,o,g)}(N, \omega) \tag{3.1}$$

is continuous for all continuous bounded $f : \widehat{\mathcal{M}} \to \mathbb{R}$ then $\widehat{\mu} = \int \mathbb{P}_{(M,o,g)}\mathrm{d}\mu$ is well defined as an element of the dual of the space of bounded continuous functions. By the Riesz representation theorem this means that $\widehat{\mu}$ is well defined as a probability on the Stone-Cech compactification of $\widehat{\mathcal{M}}$. After this it suffices to show that for each $\epsilon > 0$ there is a compact subset K of $\widehat{\mathcal{M}}$ with $P_{(M,o,g)}(K) \geq 1 - \epsilon$ for all (M, o, g) in $\mathcal{M}$ to obtain that in fact $\widehat{\mu}$ is a probability on $\widehat{\mathcal{M}}$.

To begin we define the space of orthonormal frames $O(\mathcal{M})$ as the set of tuples

$$(M, o, g, v_1, \ldots, v_d)$$

where (M, o, g) is in $\mathcal{M}$ and $v_1, \ldots, v_d$ is an orthonormal basis of the tangent space T_oM. This space is considered up to equivalence by isometries which respect the basepoints and orthonormal frames. A sequence $(M^n, o^n, g^n, v_1^n, \ldots, v_d^n)$ converges to $(M, o, g, v_1, \ldots, v_d)$ if there exists an exhaustion U_n of M and smooth embeddings $\varphi_n : U_n \to M_n$ satisfying the usual properties for smooth convergence plus that $|D\varphi_n v_i - v_i^n| \to 0$ for $i = 1, \ldots, d$. We omit the verification that $O(\mathcal{M})$ is separable, compact and the projection is continuous. Similarly we define the space of frames $F(\mathcal{M})$ by dropping the condition that $v_1, \ldots, v_d$ be orthonormal (this space is no longer compact).

On the frame bundle $F(M)$ of any manifold $(M, o, g) \in \mathcal{M}$ there are d unique smooth vector fields $H_1, \ldots, H_d$ with the property that the flow defined by H_i corresponds to parallel transport of each orthonormal frame along the geodesic whose initial condition is the i-th vector of the frame. Any solution to the Stratonovich differential equation

$$dX_t = \sum_{i=1}^{d} H_i(X_t) \circ \mathrm{d}W_t^i \tag{3.2}$$

driven by a standard Brownian motion $(W_t^1, \ldots, W_t^d)$ on $\mathbb{R}^d$ and starting at an orthonormal frame, projects to a Brownian motion on M. In particular the distribution of the solution starting at any orthonormal frame $v_1, \ldots, v_d$ over the basepoint o is the Weiner measure P_o and lifts to $P_{(M,o,g)}$ on $\widehat{\mathcal{M}}$. We will show that $P_{(M,o,g)}$ depends continuously on (M, o, g) in $\mathcal{M}$ by showing that it depends continuously on the points in $O(\mathcal{M})$. Since points projecting to the same element of $\mathcal{M}$ have the same associated measure in $\widehat{\mathcal{M}}$ this approach may seem unnecessarily complicated, however it allows us to use standard theorems on continuity of solutions to rough differential equations and simultaneously solve the problem of finding compact sets with large probability for all $P_{(M,o,g)}$.

Choose $p \in (2, 3)$ and consider for each $T > 0$ the Polish space Ω of geometric rough paths with locally finite p-variation (by which we mean that restricted to each interval

$I = [0,T]$ the path belongs to $C^{p-var}(I, G^2(\mathbb{R}^d))$ as defined in [FV10, Definition 9.15]). Let ν be the probability measure on $\Omega G_p(\mathbb{R}^d)$ which is the distribution of 'Stranovich enhanced Brownian motion' (see [FV10, Section 13.2]). Given $(M, o, g, v_1, \ldots, v_d)$ and a path α in Ω there is a unique curve $\omega : [0, +\infty) \to M$ which is the projection to M of the solution to Equation 3.2 driven by the rough path α (this is the existence and uniqueness theorem for rough differential equations of Terry Lyons, we will use the version given by [FV10, Theorem 10.26]). We will show that (M, ω) is continuous as a function from $\Omega \times O(\mathcal{M})$ to $\widehat{\mathcal{M}}$ (where the later is endowed with the topology of local uniform convergence).

Assuming the continuity claim has been established we notice that, because the push-forward of ν under the map $\alpha \mapsto \omega$ is the Weiner measure corresponding to o, it would follow by dominated convergence that the function defined by Equation 3.1 is continuous for all bounded f. Furthermore, by letting α vary in a compact subset of Ω with ν-probability greater than $1 - \epsilon$ one would obtain that there is a compact subset of $\widehat{\mathcal{M}}$ with probability greater than $1 - \epsilon$ for all $P_{(M,o,g)}$.

We now establish the continuity claim. Consider a sequence in $O(\mathcal{M})$ such that

$$(M^n, o^n, g^n, v_1^n, \ldots, v_d^n) \to (M, o, g, v_1, \ldots, v_d)$$

when $n \to +\infty$, and let $\varphi_n : U_n \to M^n$ be the corresponding embeddings of an exaustion of M. Given $\alpha \in \Omega$ and a compact interval $[0,T]$ we may take a compact d-dimensional submanifold with boundary $N \subset F(M)$ such that the solution to Equation 3.2 driven by α starting at $v_1, \ldots, v_d$ remains in N on $[0,T]$. We embed N into some $\mathbb{R}^N$ and consider smooth compactly supported extensions of the horizontal vector fields H_i to $\mathbb{R}^N$. The pullbacks of the horizontal vector fields H_i^n on $F(M_n)$ under φ_n restricted to N can be extended to all of $\mathbb{R}^N$ in such a way that they share a compact support and converge smoothly to the corresponding H_i and the pullbacks of the orthonormal frames v_i^n eventually belong to N and converge to $v_1, \ldots, v_d$. Hence by [FV10, Theorem 10.26] the solutions to Equation 3.2 driven by α and the vector fields H_i^n starting at the pullbacks of $v_1^n, \ldots, v_d^n$ converge to the corresponding solution driven by the vector fields H_i uniformly on $[0,T]$. This implies continuity of the solution with respect to $(M, o, g, v_1, \ldots, v_d)$ in $O(\mathcal{M})$ [1]. Continuity with respect to α for a fixed manifold with frame in $O(\mathcal{M})$ also follows from [FV10, Theorem 10.26].

We have thus far established the existence of $\widehat{\mu}$ for each probability μ on $\mathcal{M}$. It remains to interpret harmonicity and ergodicity of μ with properties of $\widehat{\mu}$ relative to the shift maps.

If μ is harmonic then by Corollary 3.4 one has for all $A \in \mathcal{F}_t^\infty$ that $\widehat{\mu}(A) = \text{shift}_*^t \widehat{\mu}(A)$ from this it follows that the pushforward measure $\text{shift}_*^t \widehat{\mu}$ (defined on $\mathcal{F}_{-t}^\infty$ coincides with μ on $\mathcal{F}_0^\infty$. Hence we may extend $\widehat{\mu}$ uniquely to a shift invariant probability $\mathcal{F}_t^\infty$ for all $t \in \mathbb{R}$. Since all these extensions are compatible Tulcea's extension theorem implies that there is a unique shift invariant Borel extension.

Suppose that μ is ergodic, harmonic, and that the unique shift invariant extension of $\widehat{\mu}$ to the Borel σ-algebra is not ergodic. Then by definition one can obtain two distinct shift invariant probabilities such that

$$\widehat{\mu} = \alpha \widehat{\mu}_1 + (1 - \alpha) \widehat{\mu}_2$$

for some $\alpha \in (0,1)$. Since $\widehat{\mu}$ projects to μ this allows one to express μ as a non-trivial convex combination of the projection μ_1 and μ_2 of $\widehat{\mu}_1$ and $\widehat{\mu}_2$ respectively. Using Corollary

1. Here we use the fact that if there is an exhaustion U_n of M and embeddings $\varphi_n : U_n \to M_n$ defining the conditions for smooth convergence, and furthermore there are curves ω on M and ω_n on M_n such that $\varphi_n^{-1} \circ \omega_n$ converges to ω uniformly on compact sets, then (M_n, ω_n) converges to (M, ω) as elements of $\widehat{\mathcal{M}}$. The proof amounts to repeating the arguments in Lemma 2.1.

3.4 one sees that because the $\widehat{\mu}_i$ are shift invariant μ_1 and μ_2 are harmonic. But because μ is ergodic one must have $\mu_1 = \mu_2 = \mu$ contradicting the fact that $\widehat{\mu}_1 \neq \widehat{\mu}_2$. Hence each ergodic harmonic measure μ on $\mathcal{M}$ lifts to (after taking the unique shift invariant extension) an ergodic shift invariant probability $\widehat{\mu}$ as claimed. □

3.1.3 Pathwise linear drift

Suppose X_t is a Brownian motion on a manifold with bounded geometry. By the results of Ichihara (see [Ich88, Example 2.1]) one has that $\limsup_{t \to +\infty} d(X_0, X_t)/t$ is finite. However, the limit

$$\lim_{t \to +\infty} \frac{d(X_0, X_t)}{t}$$

need not exists almost surely. Furthermore even if the above limit exists it might be random (i.e. take different values with positive probability). For a concrete example consider a metric on $\mathbb{R}^2$ which contains isometrically embedded copies of a half space with constant curvature -1 and a half space of constant curvature -2. For any Brownian motion with respect to such a metric the above limit will take two distinct values with positive probability.

The construction of the Brownian motion process on $\widehat{\mathcal{M}}$ implies that manifolds such as the previously discussed cannot be generic with respect to any ergodic harmonic measure.

Theorem 3.6. *Let (M, ω) be a Brownian motion an an ergodic stationary random manifold. Then there one has*

$$\lim_{t \to +\infty} \mathbb{E}\left(\frac{d(\omega_0, \omega_t)}{t}\right) = \ell(M)$$

and

$$\lim_{t \to +\infty} \frac{d(X_0, X_t)}{t} = \ell(M)$$

almost surely, where $\ell(M)$ is the linear drift of the stationary random manifold M.

Proof. Let μ be the distribution of M and consider the shift invariant lift $\widehat{\mu}$ of μ to $\widehat{\mathcal{M}}$ given by Theorem 3.5. We define the function $d_{s,t}$ on $\widehat{\mathcal{M}}$ for $s < t$ by

$$d_{s,t}(M, \omega) = d(\omega_s, \omega_t).$$

The triangle inequality implies that $d_{s,u} \leq d_{s,t} + d_{s,u}$, added to the fact that $\widehat{\mu}$ is shift invariant one obtains that the family $\{d_{s,t}\}$ is a stationary subadditive process with respect to the probability $\widehat{\mu}$. Therefore by Kingman's subadditive ergodic theorem (see Theorems 1 and 5 of [Kin68]) the limits

$$\lim_{t \to +\infty} \frac{d_{s,t}}{t}$$

exists $\widehat{\mu}$ almost surely and in L^1. Since the limit is almost surely shift invariant it must be constant if $\widehat{\mu}$ is ergodic. □

As a consequence of the previous theorem we can complete the proof that $\ell^+(M) = \ell(M)$ for all ergodic stationary random manifolds M, a fact which was used in the previous chapter.

Corollary 3.7. *Let M be an ergodic stationary random manifold then one has*

$$\lim_{t\to+\infty} \mathbb{E}\left(\int\limits_{B_{Lt}(o_M)} q(t,o_M,x)\mathrm{d}x\right) = 1$$

for all $L > \ell(M)$.

Proof. Let μ be the distribution of M, $\mathcal{M}$ the support of μ, and $\widehat{\mu}$ the lift of μ to $\widehat{\mathcal{M}}$. The quantity inside the limit is equal to

$$\widehat{\mu}\left(\{(M,\omega)\in\widehat{\mu} : d(\omega_0,\omega_t) < Lt\}\right)$$

which converges to 1 by Theorem 3.6. □

3.1.4 Recurrence and the correspondence principle

Szemeredi's theorem states that if a set of integers A has positive upper density then it contains arbitrarily long arithmetic progressions. Furstenberg's proof of this result begins by associating to A a shift invariant probability measure μ on the space of infinite strings of zeros and ones with the property that if the measure of the set of strings beginning with a particular finite pattern is positive then this pattern appears in A. In particular the set U of strings beginning with 1 has positive probability and Szermeridi's theorem is shown to be equivalent to the property that for each k there exist n such that $\mu(U\cap T^{-n}U\cap\cdots\cap T^{-nk}U) > 0$ where T is the shift transformation (see [EW11, pg. 178]).

The above idea of Furstenberg has been abstracted to a general albeit informal 'correspondence principle' which might be roughly stated as follows: The combinatorial properties of a concrete mathematical object can sometimes be codified by a measure preserving dynamical system. Hence ergodic theorems can yield proofs of combinatorial statements and vice-versa.

In this subsection we try to apply the above principle to a fixed bounded geometry manifold (M,g). The idea is that the closure of the set of manifolds (M,x,g) where x varies over all of M, supports at least one harmonic measure 'diffused from M'. If for some $o\in M$ the pointed manifold (M,o,g) belongs to the support of this measure then this imposes strong 'recurrence' properties on the Riemannian metric and topology of M forcing a certain finite radius 'pattern' in the manifold to appear infinitely many times. Perhaps the strongest result which has been obtained by this type of reasoning is the theorem of Ghys which states that a non-compact leaf of a compact foliation by surfaces which is generic with respect to a harmonic measure can only have one out of six possible topologies (see [Ghy95]). We will prove a weaker result which illustrates the same principle. We begin with a definition.

Definition 3.8. Let (M,o,g) be a pointed bounded geometry manifold in some set of manifolds with uniformly bounded geometry $\mathcal{M}$. A diffusion measure (M,o,g) is any weak limit when $t\to+\infty$ of a subsequence of the measures

$$\frac{1}{t}\int_0^t (\pi\circ\mathrm{shift}^s)_* P_{(M,o,g)}\mathrm{d}s$$

where π is the projection from $\widehat{\mathcal{M}}$ to $\mathcal{M}$ and the shift maps on $\widehat{\mathcal{M}}$ are denoted by shift^t.

We might say a manifold is 'recurrent' if it belongs to the support of at least one of its diffusion measures. We begin by showing that, with this definition, all manifolds which are generic with respect to an ergodic harmonic measure are recurrent and in fact have a unique diffusion measure.

Lemma 3.9. *Let μ be an ergodic harmonic measure on a space of manifolds with uniformly bounded geometry $\mathcal{M}$. Then μ is the unique diffusion measure for μ-almost every $(M, o, g) \in \mathcal{M}$.*

Proof. Consider the ergodic shift invariant lift $\widehat{\mu}$ of μ to $\widehat{\mathcal{M}}$ given by Theorem 3.5 and let $\{U_n\}$ be a countable basis of the topology on $\mathcal{M}$ by open sets whose boundaries have zero μ measure. By Birkhoff's ergodic theorem one has for $\widehat{\mu}$ almost every (M, ω) in $\widehat{\mathcal{M}}$ that

$$\mu(U_n) = \widehat{\mu}\left(\pi^{-1}(U_n)\right) = \lim_{t\to+\infty} \frac{1}{t}\int_0^t 1_{U_n}\left(\pi \circ \mathrm{shift}^s(M, \omega)\right) \mathrm{d}s$$

for all n.

This implies that

$$\mu = \lim_{t\to+\infty} \frac{1}{t}\int_0^t (\pi \circ \mathrm{shift}^s)_* P_{(M,\omega_0,g)} \mathrm{d}s$$

so μ is a diffusion measure for $\widehat{\mu}$ almost every (M, ω_0, g). This implies the claim since $\widehat{\mu}$ projects to μ. □

The following theorem shows, for example, that the plane with one handle cannot be the generic leaf of a compact foliation no matter what bounded geometry Riemannian metric we put on it.

Theorem 3.10. *Suppose that (M, o, g) is a non-compact bounded geometry manifold which belongs to the support of its only diffusion measure μ. Then for each $r > 0$ there are infinitely many disjoint diffeomorphic copies of $B_r(o)$ embedded in M.*

Proof. Given r we can choose a neighborhood U of (M, o, g) sufficiently small so that for all (M', o', g') in U there is a diffeomorphism $\varphi : B_r(o) \to M'$ with $\varphi(o) = o'$ and such that $\varphi(B_r(o))$ is contained in $B_{2r}(o')$.

Consider the set $A = \{x \in M : (M, x, g) \in U\}$ and notice that one has

$$\lim_{t\to+\infty} \frac{1}{t}\int_0^t 1_A(\omega_s)\mathrm{d}s = \mu(U) > 0$$

for P_o almost every Brownian path ω.

On the other hand, since M is non-compact and has bounded geometry it has infinite volume. This implies that the fraction of time spent by Brownian motion in any compact subset of M converges to 0 as $t \to +\infty$. It follows that A is unbounded and hence M contains infinitely many disjoint copies of $B_r(o)$. □

The main result of [Ghy95] can be stated by saying that the non-compact generic leaves with respect to any ergodic harmonic measure on a compact foliation by surfaces has either zero or infinite genus and either one, two, or infinitely many ends. The first part of this result follows from the theorem above.

Corollary 3.11 (Ghys)**.** *Any non-compact generic leaf with respect to an ergodic harmonic measure on a compact foliation by surfaces must have either zero or infinite genus.*

3.2 Busemann functions and linear drift

3.2.1 Busemann functions

The Busemann function $\xi_x : X \to \mathbb{R}$ associated to a point x in a pointed proper metric space (X, o) is defined by

$$\xi_x(y) = d(x, y) - d(x, o)$$

and the Busemann compactification of X is the closure of all such functions in the topology of local uniform convergence (this is equivalent to pointwise convergence since all Busemann functions are 1-Lipschitz). Given a set of pointed manifolds with uniformly bounded geometry $\mathcal{M}$ we will show in this section that the linear drift of a harmonic measure on $\mathcal{M}$ can be expressed in terms of the increment of a 'random Busemann function' on a fixed time interval of Brownian motion.

As a first step we must construct a space containing the Brownian paths $\widehat{\mathcal{M}}$ and Busemann functions. For this purpose we define $\widehat{\mathcal{M}}^1$ as the set of triplets (M, ω, ξ) where (M, ω) is in $\widehat{\mathcal{M}}$ and $\xi : M \to \mathbb{R}$ is 1-Lipschitz and satisfies $\xi(\omega_0) = 0$. Elements of $\widehat{\mathcal{M}}^1$ are considered up to the usual equivalence by isometries preserving ω and ξ.

The distance on $\widehat{\mathcal{M}}^1$ is defined similarly to that of $\widehat{\mathcal{M}}$ adding only one condition.

Definition 3.12. Define the distance between two elements (M, ω, ξ) and (M', ω', ξ') in $\widehat{\mathcal{M}}$ as either $1/2$ or, if such an ϵ exists, the infimum among all $\epsilon \in (0, 1/2)$ such that that there exist an admissible metric on the disjoint union $M \sqcup M'$ with the following properties:

1. $d(o, o') < \epsilon$.
2. $d(B_{1/\epsilon}(o), M') < \epsilon$ and $d(M, B_{1/\epsilon}(o')) < \epsilon$.
3. $d(\omega_t, \omega'_t) < \epsilon$ for all $t \in [-1/\epsilon, 1/\epsilon]$.
4. $|\xi(x) - \xi'(x')| < \epsilon$ whenever $d(x, x') < \epsilon$ and either $x \in B_{1/\epsilon}(o)$ or $x' \in B_{1/\epsilon}(o')$.

Usually the Busemann functions of a manifold M are defined as the locally uniform limits of the functions ξ_x (where $x \in M$). We define the generalized Busemann space $\widehat{\mathcal{M}}^b$ as the closure of the Busemann functions (M, ω, ξ_x) (where $x \in M$) in the above metric space. The elements of this space will be called generalized Busemann functions.

Lemma 3.13. *The metric space $\widehat{\mathcal{M}}^b$ is complete and separable, and the projection $\widehat{\pi}$: $\widehat{\mathcal{M}}^b \to \widehat{\mathcal{M}}$ is continuous surjective and proper (i.e. preimage of any compact set is compact).*

Proof. Completeness of $\widehat{\mathcal{M}}^b$ follows from an argument similar to that of Lemma 3.2. We consider a sequence (M^n, ω^n, ξ^n) with distance between consecutive elements less than 2^{-n}. We take an admissible metric d_n on $M_n \sqcup M_{n+1}$ satisfying the conditions of Definition 3.12 and use it to define a metric on the countable disjoint union $\bigsqcup M_n$ by setting $d(x, y)$ with $x \in M^n$ and $y \in M^{n+p}$ to be the infimum of $d_n(x, x_1) + \cdots + d_{n+p-1}(x_{p-1}, y)$ over all chains $x_1, \ldots, x_{p-1}$ with $x_i \in M_{n+i}$ for $i = 1, \ldots, p-1$. We take X to be the completion of this disjoint union and $M = X \setminus \bigsqcup M^n$. As in Lemma 3.2 one has that ω^n converge locally uniformly to a curve ω in M.

Take a point $x \in M$ and assume that $x = \lim x_n$ where $x_n \in M^n$ and $d(x_n, x_{n+1}) < 2^{-n}$ for each n. Then because $d(o_n, x_n) \to d(o, x)$ we will have for all n large enough that $|\xi^n(x_n) - \xi^{n+1}(x_{n+1})| < 2^{-n}$. This implies that $\lim \xi^n(x_n)$ exists. Furthermore, if y_n is another sequence with the same properties converging to x then $d(y_n, x_n) \to 0$ and since each ξ^n is 1-Lipschitz one has that $\lim \xi^n(y_n) = \lim \xi^n(x_n)$. We define $\xi(x) = \lim \xi^n(x_n)$. One verifies that ξ is 1-Lipschitz and defined on a dense subset of M, therefore it extends uniquely to a continuous function on M.

By definition the function F on X which coincides with ξ^n on each M^n and with ξ on M is continuous. Therefore it is uniformly continuous on compact sets. Hence, given $\epsilon > 0$ we may find $\delta > 0$ such that if $x \in B_{\epsilon+1/\epsilon}(o)$ and $d(x,y) < \delta$ then $|F(x) - F(y)| < \epsilon$. The admissible distance on $M_n \sqcup M$ which equals $d(x,y) + \epsilon - \delta$ whenever $x \in M_n$ and $y \in M$ and coincides with d otherwise, shows that the distance between (M_n, ω^n, ξ^n) and (M, ω, ξ) is less than ϵ for all n large enough. Hence $\widehat{\mathcal{M}}^b$ is complete.

Surjectivity of $\widehat{\pi}$ is immediate. Separability of $\widehat{\mathcal{M}}^b$ follows from that of $\widehat{\mathcal{M}}$ once we show that $\widehat{\pi}$ is proper (the preimage of each point in a dense countable subset of $\widehat{\mathcal{M}}$ is compact, hence has a countable dense set, the union of these sets is dense in $\widehat{\mathcal{M}}^b$).

To establish that $\widehat{\pi}$ is proper it suffices to show that if $(M^n, \omega^n, \xi_{x_n})$ is a sequence of Busemann functions in $\widehat{\mathcal{M}}^b$ such that (M^n, ω^n) converges to (M, ω) in $\widehat{\mathcal{M}}$, then there is a 1-Lipschitz function ξ on M such that a subsequence of $(M^n, \omega^n, \xi_{x_n})$ converges to (M, ω, ξ).

Repeating the construction above we may assume without loss of generality that there is an admissible metric on $\bigsqcup M_n \sqcup M$ such that the curves ω^n converge locally uniformly to ω and the metric restricted to $M_n \sqcup M_{n+1}$ satisfies properties 1,2, and 3 of Definition 3.12 for $\epsilon = 2^{-n}$. The same conditions are verified by the metric when restricted to $M_n \sqcup M$ for $\epsilon = 2^{-n+1}$.

Notice that each ξ_{x_n} is defined in terms of the distance on M^n. Hence it extends, using the admissible distance, to a 1-Lipschitz function F_n on all of $M \sqcup \bigsqcup M_n$. Since the sequence F_n is equicontinuous there is a locally uniformly convergent subsequence F_{n_k} which converges to a limit F. We define ξ as the restriction of F to M. An argument similar to the one above shows that $(M^{n_k}, \omega^{n_k}, \xi_{x_{n_k}})$ converges to (M, ω, ξ). □

3.2.2 A Furstenberg type formula for linear drift

We recall the idea of Furstenberg that the largest Lyapunov exponent

$$\chi = \lim \frac{1}{n} \log(|A_0 \cdots A_n|)$$

of a sequence $\ldots, A_{-1}, A_0, A_1, \ldots$ of independent identically distributed random 2×2 matrices with determinant equal to 1 can be expressed as

$$\chi = \mathbb{E}(\log(|A_0 v|))$$

where v is a random vector on the unit circle which is independent from A_0 (it belongs to the unstable direction of the sequence of matrices and depends only on $A_{-1}, A_{-2}, \ldots$).

The distribution of v above is, a priori, unknown but must satisfy a stationarity property, and even without full knowledge of v the formula can be used to establish that the exponent is positive under certain assumptions on the sequence of matrices. This idea has been generalized considerably and there are now several 'Furstenberg type formulas' which express asymptotic quantities of a random trajectory as integrals over a finite time segment of the trajectory involving some random element of a 'boundary space' whose distribution is unknown (see for example [KL11, Theorem 18]).

The purpose of this subsection is to establish a Furstenberg type formula for the linear drift of a harmonic measure (compare with [Led10, Proposition 1.1]). The Busemann function space $\widehat{\mathcal{M}}^b$ will play the role of the boundary (the circle in the above example) and our trajectory is the path of Brownian motion. We extend the shift maps shift^t to $\widehat{\mathcal{M}}^b$ so that (M,ω,ξ) goes to $(M,\omega_{t+\cdot},\xi-\xi(\omega_t))$ and notice that they are continuous.

By a Brownian motion on a stationary random manifold M we mean a random element (M,ω) of $\widehat{\mathcal{M}}$ whose distribution is the shift invariant lift $\widehat{\mu}$ of the harmonic distribution μ of M (see Theorem 3.5). In the following theorem the existence of the random element (M,ω,ξ) extending (M,ω) may depend on modifying the domain probability space of (M,ω) somewhat (without changing the distribution of (M,ω)).

Theorem 3.14 (Furstenberg type formula for linear drift)**.** *Let (M,ω) be a Brownian motion on an ergodic stationary random manifold M. Then (possibly modifying the domain of (M,ω)) there exists a random generalized Busemann function ξ such that the distribution of (M,ω,ξ) is shift invariant and one has*

$$\mathbb{E}\left(\xi(\omega_t)\right)=t\ell(M)$$

for all $t\in\mathbb{R}\setminus\{0\}$. Furthermore ξ can be chosen so that its distribution is ergodic for the shift maps and so that the conditional distribution of $\{\omega_t : t\geq 0\}$ given $M, o=\omega_0$ and ξ is P_o.

Proof. Let (M,ω) be a Brownian motion on an ergodic stationary random manifold M and let u be a uniform random variable in $[0,1]$. Because the projection $\widehat{\pi}$ from $\widehat{\mathcal{M}}^b$ to $\widehat{\mathcal{M}}$ is proper the distributions of the random variables $(M,\omega,\xi_{\omega_{-uT}})$ (where T ranges over $\mathbb{R}$) are tight (i.e. for each $\epsilon>0$ there is a compact subset of $\widehat{\mathcal{M}}^b$ with probability greater than $1-\epsilon$ for all the distributions). Hence there exists a sequence $T_n\to+\infty$ such that $(M,\omega,\xi_{\omega_{-uT_n}})$ converges in distribution to some random element (M,ω,ξ) in $\widehat{\mathcal{M}}^b$ (see [Bil99, Theorem 5.1], notice that this element projects to our original Brownian motion (M,ω)).

Since all functions in $\widehat{\mathcal{M}}^b$ are 1-Lipschitz we have that $|\xi(\omega_t)|\leq d(\omega_0,\omega_t)$. Notice also that, as shown in the previous chapter, the expectation of $d(\omega_0,\omega_t)$ is finite for any Brownian motion (M,ω) on a stationary random manifold. Also, by Skorohod's representation theorem (see [Bil99, Theorem 6.7]) there exist random elements with the distribution of $(M,\omega,\xi_{\omega_{-uT_n}})$ which converge pointwise to a random element with the distribution of (M,ω,ξ). Hence we may use dominated convergence to establish the first in the following chain of equalities (we assume for simplicity that $t>0$)

$$\begin{aligned}\mathbb{E}\left(\xi(\omega_t)\right)&=\lim_{n\to+\infty}\mathbb{E}\left(\xi_{\omega_{-uT_n}}(\omega_t)\right)=\lim_{n\to+\infty}\mathbb{E}\left(d(\omega_{-uT_n},\omega_t)-d(\omega_{-uT_n},\omega_0)\right)\\&=\lim_{n\to+\infty}\mathbb{E}\left(d(\omega_0,\omega_{uT_n+t})-d(\omega_0,\omega_{uT_n})\right).\end{aligned}$$

Notice that $u+t/T_n$ is uniformly distributed on $[t/T_n,t/T_n+1]$. Therefore, setting $v=u$ if $u>t/T_n$ and $v=1+u$ otherwise one obtains that v has the same distribution as $u+t/T_n$ from which it follows that

$$\mathbb{E}\left(\xi(\omega_t)\right)=\lim_{n\to+\infty}\mathbb{E}\left(d(\omega_0,\omega_{vT_n})-d(\omega_0,\omega_{uT_n})\right)$$

where the inner terms cancel except on a set of probability t/T_n where $u<t/T_n$.

Hence we have obtained

$$\mathbb{E}\left(\xi(\omega_t)\right)=t\lim_{n\to+\infty}\mathbb{E}\left(\frac{d(\omega_0,\omega_{s+T_n})}{T_n}\right)-t\lim_{n\to+\infty}\mathbb{E}\left(\frac{d(\omega_0,\omega_s)}{T_n}\right)$$

where s is independent from (M, ω) and uniformly distributed on $[0, t]$.

In second term we may bound the expected value of $d(\omega_0, \omega_s)$, using Ichihara's comparison result (see [Ich88]), by the expected diameter of a segment of length t of Brownian motion on a space of constant curvature lower than that of all possible values of M. Since $T_n \to +\infty$ it follows that the second limit is 0. For the first term we can bound the expected value of $|d(\omega_0, \omega_{s+T_n}) - d(\omega_0, \omega_{T_n})|$ by the same. This implies that one has

$$\mathbb{E}\left(\xi(\omega_t)\right) = t \lim_{n \to +\infty} \mathbb{E}\left(\frac{d(\omega_0, \omega_{T_n})}{T_n}\right) = t\ell(M)$$

where $\ell(M)$ is the linear drift of the stationary random manifold M. We omit the proof for negative t which is very similar.

To see that the distribution of (M, ω, ξ) is shift invariant take any $s > 0$ and notice that

$$\text{shift}^{-s}(M, \omega, \xi_{\omega_{uT_n}}) = (M, \omega_{\cdot - s}, \xi_{\omega_{uT_n}} - \xi_{\omega_{uT_n}}(\omega_{-s}))$$

has the same distribution as $(M, \omega, \xi_{\omega_{uT_n+s}})$. We may define v so that $uT_n + s$ has the same distribution as the vT_n by setting $v = u$ if $uT_n < s$ and $v = u + 1$ otherwise. Hence the $\text{shift}^{-s}(M, \omega, \xi_{\omega_{uT_n}})$ has the same distribution as (M, ω, ξ_{vT_n}) which coincides with (M, ω, ξ_{tT_n}) outside of a set of probability s/T_n. Taking limits when $n \to +\infty$ one obtains that $\text{shift}^{-s}(M, \omega, \xi)$ has the same distribution as (M, ω, ξ) for all $s > 0$, so the distribution of (M, ω, ξ) is shift invariant.

Since u is independent from (M, ω) and the conditional distribution of $\{\omega_t : t \geq 0\}$ relative to $M, o = \omega_0$ is P_o one obtains the same property for conditioning relative to $\xi_{\omega_{-uT_n}}$ and by taking limits also for ξ.

The possible shift invariant distributions of (M, ω, ξ) satisfying $\mathbb{E}(\xi(\omega_t)) = t\ell(M)$ for all $t \neq 0$ and such that the conditional distribution of $\{\omega_t : t \geq 0\}$ relative to $M, o = \omega_0, \xi$ are P_o, form a convex and weakly compact set. Any extremal element of this set (and such an element exists by the Krein-Milman theorem) is ergodic with respect to the shift maps. This shows that the distribution of ξ can be chosen to be ergodic. □

3.3 Entropy of reversed Brownian motion

3.3.1 Reversibility

The purpose of this section is to improve the inequality $\frac{1}{2}\ell(M)^2 \leq h(M)$ obtained in Theorem 2.15 in the case of Brownian motions on a stationary random Hadamard manifold to $2\ell(M)^2 \leq h(M)$ (recall that a Hadamard manifold is a manifold isometric to $\mathbb{R}^d$ endowed with a complete Riemannian metric of non-positive sectional curvature). This improvement was established by Kaimanovich and Ledrappier in the case of a single manifold with a compact quotient (no curvature assumption) and has strong rigidity consequences if for manifolds with negative curvature in this case (see [Led10] and the references therein [2]). We are able to prove the inequality under the assumption that Brownian motion is reversible; a technical hypothesis which is automatically satisfied in the case of a single manifold with compact quotient. We will discuss this assumption briefly in this subsection.

2. Since we use $q(t, x, y) = p(t/2, x, y)$ our definitions of ℓ and h differ from Ledrappier's by a factor of two. Hence then inequality $h \leq \ell v$ remains the same with both conventions but our claim that $2\ell^2 \leq h$ corresponds to $\ell^2 \leq h$ in Ledrappier's notation.

Consider a Brownian motion (M, ω) on a stationary random manifold M. The reverse process is defined by (M, ω') where $\omega'_t = \omega_{-t}$. We say the Brownian motion is reversible if (M, ω) and (M, ω') have the same distribution.

As an example consider a compact manifold M. The unique stationary measure for Brownian motion is the normalized volume measure hence the unique shift invariant measure on $C([0, +\infty), M)$ is $\int P_x \mathrm{d}x/\mathrm{vol}(M)$. We use this to define a unique shift invariant measure on the continuous paths $C(\mathbb{R}, M)$ defined on all of $\mathbb{R}$. Giving rise to a Brownian motion on a stationary random manifold (M, ω) where M is fixed and the basepoint ω_0 is uniformly distributed.

We claim that (M, ω) thus defined is reversible. For this purpose notice that given Borel sets A_0, A_1 in M and $t > 0$ one has

$$\mathbb{P}\left(\omega_0 \in A_0, \omega_t \in A_1\right) = \int\limits_{A_0 \times A_1} q(t, x, y)\mathrm{d}x\mathrm{d}y/\mathrm{vol}(M) = \mathbb{P}\left(\omega_0 \in A_1, \omega_t \in A_0\right)$$

because $q(t, x, y) = q(t, y, x)$. The claim follows from repeating this calculation for an arbitrary finite number of sets and times.

Given a compact foliation X and a harmonic measure μ one can define a measure on the space of paths $C([0, +\infty), X)$ corresponding to 'leafwise Brownian motion' with initial distribution μ. Similarly to the case discussed above where $X = M$ was a single compact leaf, the fact that μ is harmonic allows one to uniquely extend this probability in a shift invariant way to all of $C(\mathbb{R}, X)$. The results of Deroin and Klepsyn (see [DK07, Theorem B]) imply that for a minimal codimension one foliation without any transverse invariant measure the corresponding Brownian motion indexed on $\mathbb{R}$ is not reversible.

In the above example non-reversibility follows because holonomy is contracting in the forward time direction and hence expanding in the backward time direction. But transverse information such as holonomy is lost when one applies the leaf function to push forward the harmonic measure on X to a harmonic measure on the Gromov space. This poses the following question:

Question 3.1. Does there exist a non-reversible Brownian motion on a stationary random manifold?

3.3.2 Furstenberg type formula for Hadamard manifolds with pinched negative curvature

Recall that the Busemann functions of a manifold (M, o, g) are defined as local uniform limits of the function of the form $\xi_x(y) = d(x, y) - d(x, o)$. The Busemann functions which are not of the form ξ_x form the so-called Busemann boundary of M.

Lemma 3.15. *Let (M, o, g) be a Hadamard manifold with curvature bounded between two negative constants. Then for P_o almost every Brownian path ω the following limit exists and is an element of the Busemann boundary*

$$\xi = \lim_{t \to +\infty} \xi_{\omega_t}.$$

Proof. Convergence of ξ_{ω_t} to a boundary Busemann function follows if one shows that $d(\omega_0, \omega_t) \to +\infty$ and that the geodesic segment joining o to ω_t converges to a geodesic ray (see [Wan11, Proposition 2] and the references therein). Both of these properties of Brownian motion were established by Prat in the mid 70s (see [AT11, Theorem 3.2], [Pra71] and [Pra75]). ☐

The Martin boundary of a manifold M consists of the limits of functions of the form $y \mapsto G(x,y)/G(x,o)$ where G is a minimal Green's function (see [Wan11]). The elements of the Martin boundary are positive harmonic functions which we choose to consider up to multiplication by positive constants. In the context of Hadamard manifolds with pinched negative curvature the equivalence between Martin and Busemann boundaries was established by Anderson and Schoen.

Lemma 3.16. *Let (M,o,g) be a Hadamard manifold with curvature bounded between two negative constants. Then there is a natural homeomorphism $\xi \mapsto k_\xi$ between the Busemann and Martin boundaries. Furthermore the probability transition density of Brownian motion conditioned on the value of $\xi = \lim_{t\to+\infty} \xi_{\omega_t}$ is given by*

$$\frac{k_\xi(y)}{k_\xi(x)} q(t,x,y).$$

Proof. For the homeomorphism between the two boundaries see [AS85, Theorem 6.3]. The conditional distribution for Brownian motion is verified on page 36 of [Anc90] (this is a special case of the so-called h-transform due to Doob, see [Doo01]). □

We can now established a refined version of Theorem 3.14 for random Hadamard manifolds with pinched negative curvature.

Lemma 3.17. *Let (M,ω) be a reversible Brownian motion on a stationary random Hadamard manifold with sectional curvatures bounded between two negative constants. Then letting $\xi = \lim_{t\to+\infty} \xi_{\omega_{-t}}$ one has*

$$\mathbb{E}(\xi(\omega_t)) = t\ell(M)$$

for all $t \neq 0$.

Proof. By Lemma 3.15 the distribution of $(M,\omega,\xi_{\omega_{-uT}})$ where u is uniformly distributed in $[0,1]$ and independent from (M,ω), converges to that of (M,ω,ξ). Hence the result follows exactly as in the proof of Theorem 3.14. □

Boundary Busemann functions are at least two times continuously differentiable on any Hadamard manifold (see [HIH77]) this allows us to pass to the limit when $t \to 0$ in the formulas above (this idea for obtaining infinitesimal formulas for the linear drift goes back to [Kaĭ86]). Doing so along positive and negative t yields different formulas, in order to use Lemma 3.16 to obtain the distribution of reversed Brownian motion conditioned to ξ we must impose the hypothesis that our Brownian motion is reversible.

Lemma 3.18. *Let (M,ω) be a reversible Brownian motion on a stationary random Hadamard manifold with sectional curvatures bounded between two negative constants and let $\xi = \lim_{t\to+\infty} \xi_{\omega_{-t}}$. Then one has*

$$\ell(M) = \mathbb{E}\left(\frac{1}{2}\Delta\xi(o_M)\right) = -\mathbb{E}\left(\frac{1}{2}\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle\right).$$

Proof. By Lemma 3.17 on has

$$\ell(M) = \mathbb{E}\left(\frac{1}{t}\int q(t,o_M,x)\xi(x)\mathrm{d}x\right)$$

for all $t > 0$. Since the Busemann functions ξ are 1-Lipschitz and $\int q(t, o_M, x) d(o_M, x) \mathrm{d}x$ can be uniformly bounded on the support of the distribution of M, one can take limit when $t \to 0^+$ from which one obtains that

$$\ell(M) = \mathbb{E}\left(\frac{1}{2}\Delta\xi(o_M)\right).$$

Similarly, by our assumption of reversibility, ω_{-t} has distribution $q(t, o_M, x)$ on M. The conditional distribution with respect to ξ is given by Lemma 3.16 and one has

$$\ell(M) = -\mathbb{E}\left(\frac{1}{t}\int \frac{k_\xi(x)}{k_\xi(o_M)} q(t, o_M, x)\xi(x)\mathrm{d}x\right)$$

for all $t > 0$. Passing to the limit with $t \to 0^-$ one obtains [3]

$$\ell(M) = -\mathbb{E}\left(\frac{1}{2}\Delta\xi(o_M) + \langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle\right)$$

from which the second claimed formula for $\ell(M)$ follows using the first. □

As a toy example of the previous formulas for drift consider the hyperbolic half plane $M = \{(x, y) \in \mathbb{R}^2 : y > 0\}$ with the metric $\mathrm{d}s^2 = y^{-2}(\mathrm{d}x^2 + \mathrm{d}y^2)$ and base point $o_M = (0, 1)$. Notice that (M, o_M) is a stationary random manifold so one may indeed apply Lemma 3.18. All Busemann functions are obtained by applying isometries to $\xi(x, y) = -\log(y)$. Calculating the Laplacian of ξ at o_M one obtains 1 so $\ell(M) = 1/2$. On the other hand the positive harmonic function associated to ξ (the Poisson kernel function associated to the boundary point at infinity) is $k_\xi(x, y) = y$ so that in the second formula $\nabla \log k_\xi = -\nabla\xi$ and one obtains again $\ell(M) = 1/2$.

As stated before our objective in this section is to obtain the inequality $2\ell(M)^2 \leq h(M)$ for stationary random Hadamard manifolds with curvature bounded between two negative constants and reversible Brownian motion. So far we have the following.

Corollary 3.19. *Let (M, ω) be a reversible Brownian motion on a stationary random Hadamard manifold with sectional curvatures bounded between two negative constants and let $\xi = \lim_{t\to+\infty} \xi_{\omega_{-t}}$. Then one has*

$$2\ell(M)^2 \leq \mathbb{E}\left(\frac{1}{2}|\nabla \log k_\xi(o_M)|^2\right)$$

Proof. Using Lemma 3.18 followed by Jensen's inequality and the Cauchy-Schwarz inequality one obtains

$$2\ell(M)^2 \leq \mathbb{E}\left(\frac{1}{2}|\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle|^2\right) \leq \mathbb{E}\left(\frac{1}{2}|\nabla \log k_\xi(o_M)|^2\right)$$

where we have also used the fact that ξ is 1-Lipschitz. □

3. Notice that $k_\xi(x)/k_\xi(o_M)$ is bounded by $C\exp(Cd(o_M, x))$ for some C (see [ADT07, Corollary 4.5]) and ξ is Lipschitz so one has a uniform bound for the inner integral on all manifolds in the support of the distribution of M by virtue of the uniform upper heat kernel bounds.

3.3.3 Reverse entropy and entropy difference

The purpose of this section is to understand the right-hand side of the inequality in Corollary 3.19. We want to show that it is smaller than $h(M)$ in order to obtain the inequality $2\ell(M)^2 \le h(M)$. In short the proof consists in establishing that the right hand side $k(M)$ is the difference between the entropy $h(M)$ and an entropy associated to the reversed process (conditioned on $\xi = \lim\limits_{t\to+\infty} \xi_{\omega_{-t}}$), since this 'reversed entropy' is non-negative one obtains $k(M) \le h(M)$. As a first step we obtain an alternate formula for $k(M)$ as the expected increment of $\log k_\xi$ along a Brownian path.

Lemma 3.20. *Let (M,ω) be a reversible Brownian motion on a stationary random Hadamard manifold with curvature bounded between two negative constants, and let $\xi = \lim\limits_{t\to+\infty} \xi_{\omega_{-t}}$. Then setting*

$$k(M) = \mathbb{E}\left(\frac{1}{2}|\nabla \log k_\xi(o_M)|^2\right)$$

one has

$$k(M) = -\mathbb{E}\left(\frac{1}{t}\log\left(\frac{k_\xi(\omega_t)}{k_\xi(o_M)}\right)\right) = \mathbb{E}\left(\frac{1}{t}\log\left(\frac{k_\xi(\omega_{-t})}{k_\xi(o_M)}\right)\right)$$

for all $t > 0$.

Proof. Using that k_ξ is shift invariant one obtains for

$$K_t = -\mathbb{E}\left(\log\left(\frac{k_\xi(\omega_t)}{k_\xi(o_M)}\right)\right)$$

that $K_{t+s} = K_t + K_s$.

Furthermore since there is a uniform bound for $|\nabla \log k_\xi|$ over the entire support of (M,ω,ξ) (see [ADT07, Corollary 4.4]) one obtains that K_t is continuous with respect to t by dominated convergence. This implies $K_t = tK_1$ for all $t > 0$.

In particular one has for all $t > 0$ that

$$K_1 = -\mathbb{E}\left(\frac{1}{t}\int q(t,o_M,x)\log\left(\frac{k_\xi(x)}{k_\xi(o_M)}\right)\mathrm{d}x\right)$$

and taking limit when $t \to 0^+$ (which is justified by dominated convergence, again by the uniform bounds of [ADT07, Corollary 4.4]) one obtains

$$K_1 = \mathbb{E}\left(\frac{1}{2}\Delta \log k_\xi(o_M)\right) = \mathbb{E}\left(\frac{1}{2}|\nabla \log k_\xi(o_M)|^2\right)$$

as claimed.

The equality

$$-\mathbb{E}\left(\frac{1}{t}\log\left(\frac{k_\xi(\omega_t)}{k_\xi(o_M)}\right)\right) = \mathbb{E}\left(\frac{1}{t}\log\left(\frac{k_\xi(\omega_{-t})}{k_\xi(o_M)}\right)\right)$$

follows by shift invariance of k_ξ. □

Given a Hadamard manifold (M,o,g) with curvature bounded by two negative constants and a boundary Busemann function ξ let $P_{x,\xi}$ be the probability measure on $C([0,+\infty),M)$ which is the distribution of Brownian motion conditioned to exit at ξ

(the transition probability densities are given by Lemma 3.16). As in Chapter 1 we denote by $\mathcal{F}_t$ and $\mathcal{F}^T$ the σ-algebras generated by ω_s with $s \le t$ and $s \ge T$ respectively and by $\mathcal{F}^\infty$ the tail σ-algebra on $C([0,\infty), M)$.

In this context let $I_t^T(M,\xi)$ (where $0 < t < T \le \infty$) be the mutual information between $\mathcal{F}_t$ and $\mathcal{F}^T$ with respect to the probability $P_{o,\xi}$.

Lemma 3.21. *Let (M, o_M, g) be a Hadamard manifold with curvature strictly bounded between two negative constants and ξ be a boundary Busemann function. Then the following properties hold for all $0 < t < T < \infty$:*

1. $I_t^T(M,\xi) = \int \log\left(\frac{k_\xi(o_M)}{k_\xi(x)}\frac{q(T-t,x,y)}{q(T,o_M,y)}\right)\frac{k_\xi(y)}{k_\xi(o_M)} q(t,o,x)q(T-T,x,y)\mathrm{d}x\mathrm{d}y.$

2. *The function $T \mapsto I_t^T(M,\xi)$ is non-negative and non-increasing.*

Proof. The formula for $I_t^T(M,\xi)$ follows from Lemma 3.16 and the Gelfand-Yaglom-Peres theorem (see the proof of Theorem 1.10). Non-negativity and monotonicity follow directly form the definition of mutual information. $\square$

The following result identifies $k(M)$ as the difference between $h(M)$ and an entropy for the reverse process of (M,ω) conditioned to its limit Busemann function ξ.

Lemma 3.22. *Let (M,ω) be a reversible Brownian motion on a stationary random Hadamard manifold with curvature bounded between two negative constants, and let $\xi = \lim\limits_{t\to+\infty} \xi_{\omega_{-t}}$. Then for all $t > 0$ one has*

$$0 \le \lim_{T\to+\infty} \mathbb{E}\left(I_t^T(M,\xi)\right) = t(h(M) - k(M)).$$

In particular $k(M) \le h(M)$.

Proof. We calculate using Lemma 3.21, Lemma 3.16, and the fact that (M,ω,ξ) is shift invariant to obtain

$$\begin{aligned}
0 \le \mathbb{E}\left(I_t^T(M,\xi)\right) &= \mathbb{E}\left(\log\left(\frac{k_\xi(\omega_0)}{k_\xi(\omega_{-t})}\right) + \log(q(T-t,\omega_{-t},\omega_{-T}) - \log(q(T,\omega_0,\omega_{-T})\right)\\
&= \mathbb{E}\left(\log\left(\frac{k_\xi(\omega_t)}{k_\xi(\omega_0)}\right) + \log(q(T-t,\omega_0,\omega_{T-t})) - \log(q(T,\omega_0,\omega_T))\right)\\
&= tk(M) - H_{T-t} + H_T
\end{aligned}$$

where $H_t = \mathbb{E}(\int \log(q(t,o_M,x))q(t,o_M,x)\mathrm{d}x)$.

When $T \to +\infty$ one has that $H_T - H_{T-t}$ converges to $th(M)$ (see the proof of Theorem 2.11) and therefore one has

$$0 \le \mathbb{E}\left(I_t^T(M,\xi)\right) = t(h(M) - k(M))$$

as claimed. $\square$

To conclude we combine the previous results to obtain a sharp lower bound for the entropy $h(M)$ of a stationary random Hadamard manifold in terms of its linear drift $\ell(M)$.

Theorem 3.23. *Let (M, ω) be a reversible Brownian motion on a stationary random Hadamard manifold with curvature bounded between two negative constants, and let $\xi = \lim\limits_{t \to +\infty} \xi_{\omega_{-t}}$. Then one has $2\ell(M)^2 \le h(M)$ with equality if and only if*

$$\nabla \log k_\xi(o_M) = -2\ell(M)\nabla\xi(o_M)$$

almost surely.

Proof. By Lemma 3.18 one has

$$\ell(M) = -\mathbb{E}\left(\frac{1}{2}\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle\right).$$

Squaring and applying Jensen's inequality one obtains

$$2\ell(M)^2 \le \mathbb{E}\left(\frac{1}{2}|\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle|^2\right).$$

Notice at this point that if the equality $\mathbb{E}(X)^2 = \mathbb{E}(X^2)$ holds for some random variable X then X is almost surely constant. Hence if equality holds in the last inequality above one obtains that $\ell(M) = -\frac{1}{2}\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle$ almost surely.

Next we apply the Cauchy-Schwartz inequality and the fact that ξ is 1-Lipschitz to obtain

$$\mathbb{E}\left(\frac{1}{2}|\langle \nabla \log k_\xi(o_M), \nabla\xi(o_M)\rangle|^2\right) \le k(M)$$

and by Lemma 3.22 conclude that $2\ell(M)^2 \le k(M) \le h(M)$.

At this step we notice that equality would imply that $\nabla \log k_\xi(o_M)$ and $\nabla\xi(o_M)$ are collinear almost surely. Combining this with the previous observation about equality in the Jensen's inequality one sees that if $2\ell(M)^2 = h(M)$ then one would have

$$\nabla \log k_\xi(o_M) = -2\ell(M)\nabla\xi(o_M)$$

almost surely as claimed. □

Part II

Gromov-Hausdorff convergence of leaves of compact foliations

Chapter 4

The leaf function of compact foliations

4.1 Introduction

In this chapter we study the regularity of the leaf function of a compact foliation X, i.e. the function associating to each $x \in X$ the leaf L_x of x considered as an element of the Gromov space.

The proof of measurability of the leaf function (which we needed in order to push-forward harmonic measures of any foliation to the Gromov space, see Lemma 2.8) depends on the results of this chapter. However, we are more interested in continuity properties and how they relate to Reeb type stability results.

Recall that the Reeb local stability theorem [Ree47, Theorem 2] states that if the fundamental group of a compact leaf in a foliation is finite then all nearby leaves are finite covers of it. In the special case when a leaf is compact and has trivial holonomy one can strengthen the conclusion to yield that all nearby leaves are diffeomorphic to the given leaf (this is the case for example for simply connected leaves such as spheres).

In [ÁC03] Álvarez and Candel introduced the leaf function as part of a program for studying the geometry (e.g. quasi-isometry invariants) of generic leaves in foliations. One result in this program is that the leaf function of any compact foliation is continuous on the set of leaves (compact or otherwise) without holonomy (see [ÁC03, Theorem 2]).

In general a sequence of manifolds can converge in the Gromov-Hausdorff sense to a compact manifold without any element of the sequence being homeomorphic to the limit (for example one can shrink the handle on a sphere with one handle to obtain a sequence converging to a sphere, see [BBI01, Figure 7.4]).

Our first result, Theorem 4.1, shows that this type of sequences do not exist within a compact foliation. As a consequence one can conclude that on compact foliations Álvarez and Candel's continuity theorem implies Reeb stability of compact leaves with trivial holonomy as a special case.

Our second and main result, Theorem 4.3, is that the leaf function of a compact foliation is semicontinuos in the sense that the limit of any sequence of leaves is a Riemannian covering of the limiting leaf. An upper bound (the so-called holonomy covering) is provided for the coverings obtainable this way and allows us to obtain Reeb's local stability theorem and Ávarez and Candel's continuity theorem as special cases.

4.2 Examples of leaf functions

Recall that by a d-dimensional foliation we mean a metric space X partitioned into disjoint subsets called leaves. Each leaf is assumed to be a continuously and injectively immersed d-dimensional connected complete Riemannian manifold. We further assume that each $x \in X$ belongs to an open set U such that there exists a Polish space T and a homeomorphism $h : \mathbb{R}^d \times T \to U$ with the following properties:

1. For each $t \in T$ the map $x \mapsto h(x,t)$ is a smooth injective immersion of $\mathbb{R}^d$ into a single leaf.

2. For each $t \in T$ let g_t be the metric on $\mathbb{R}^d$ obtained by pullback under $x \mapsto h(x,t)$ of the corresponding leaf's metric. If a sequence t_n converges to $t \in T$ then the Riemannian metrics g_{t_n} converge smoothly on compact sets to g_t.

Given a point x in a foliation X we denote by (L_x, x, g_{L_x}) the leaf of x considered as a pointed Riemannian manifold with basepoint x. We sometimes write only L_x and leave the basepoint x and metric g_{L_x} implicit. Homeomorphisms satisfying the conditions of h above are called foliated parametrizations and their inverses are foliated charts.

We recall that in any metric space (X,d) there is a natural distance between subsets, Hausdorff distance, which is defined by

$$d_H(A,B) = \inf\{\epsilon > 0 : d(a,B) < \epsilon \text{ and } d(A,b) < \epsilon \text{ for all } a \in A \text{ and } b \in B\}.$$

In what follows we use $B_r(x)$ to denote the open ball centered at a point x in a metric space and $\overline{B_r}(x)$ to denote its closure. A metric space is said to be proper if all closed balls are compact.

It will be convenient for this chapter to work with the distance on Gromov-space defined between two pointed proper metric spaces (X_i, x_i, d_i) where $i = 1, 2$ as

$$\mathrm{d}_{\mathcal{GS}}(X_1, X_2) = \sum_{n=1}^{+\infty} 2^{-n} \min(1, d_n)(X_1, X_2)$$

where

$$d_n(X_1, X_2) = \inf\left\{d(x_1, x_2) + d_H(\overline{B_n}(x_1), \overline{B_n}(x_2))\right\}$$

the infimum being taken over all distances d on the disjoint union $\overline{B_n}(x_1) \sqcup \overline{B_n}(x_2)$ which coincide with d_i when restricted to $\overline{B_n}(x_i)$ for $i = 1, 2$.

The notion of convergence induced by $\mathrm{d}_{\mathcal{GS}}$ is that of Hausdorff convergence of all integer radius balls. This coincides with Gromov-Hausdorff convergence only on the closed subset formed by length spaces (see [BBI01, Exercise 8.1.3, Theorem 8.1.9]). Since we will be working exclusively with Riemannian manifolds we will use the above distance instead of the one defined in Section 2.1.1 without further comment.

The leaf function of a foliation X is the function from X to $\mathcal{GS}$ is defined by

$$x \mapsto L_x$$

where the leaf L_x is considered up to pointed isometry.

We begin our study of the regularity of this function with a series of examples.

4.2.1 Example: the vinyl record foliation

Consider a foliation of the closed annulus $\{(x, y) \in \mathbb{R}^2 : 1 \leq x^2 + y^2 \leq 2\}$ such that the two boundary circles are leaves and all other leaves are spirals which accumulate on both boundary components. The leaf function of such a foliation is clearly not continuous since there are leaves which are isometric to $\mathbb{R}$ accumulating on a leaf isometric to an Euclidean circle.

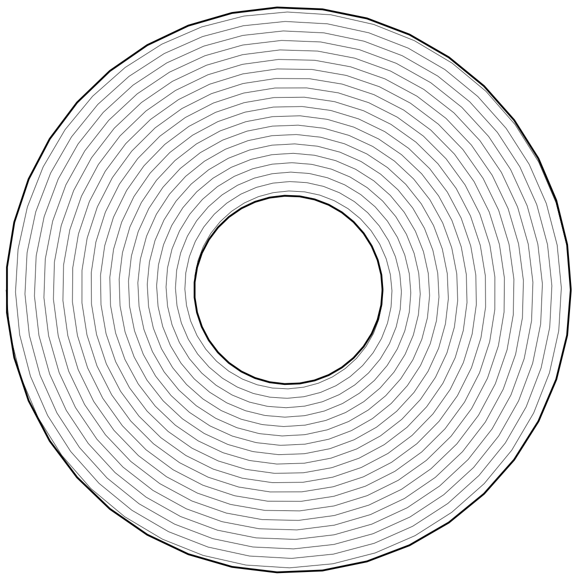

Figure 4.1: The vinyl record foliation.

4.2.2 Example: the Reeb cylinder

Consider the foliation of the solid cylinder $C = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 \leq \pi/2\}$ where the boundary cylinder is a leaf and all other leaves are of the form $\{(x, y, z) \in \mathbb{R}^d : z = t - \tan(x^2 + y^2)^2\}$ for $t \in \mathbb{R}$.

In this example the leaf function is continuous but there are simply connected leaves accumulating on a non-simply connected leaf. Hence the function

$$p \mapsto \widetilde{L}_p$$

associating to each point in C the universal covering of its leaf, is not continuous.

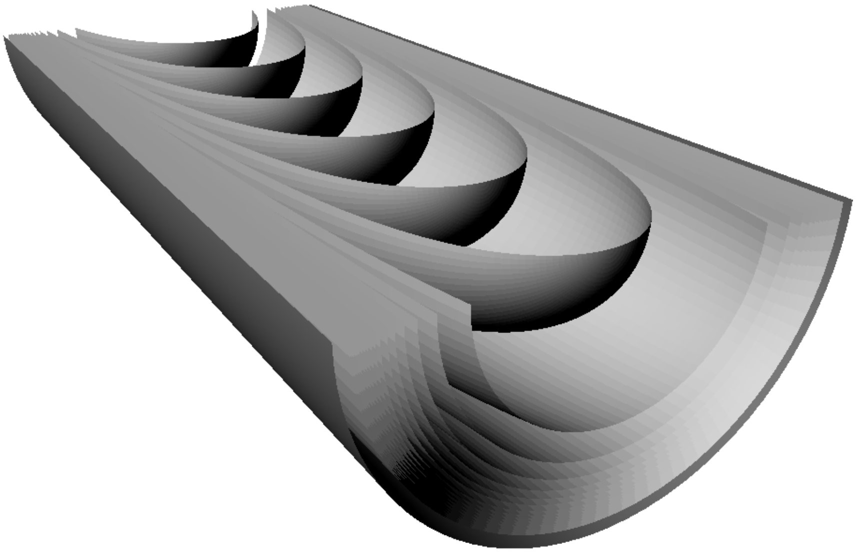

Figure 4.2: A section of the Reeb cylinder.

4.2.3 Example: the Reeb component

One may take the quotient space of a Reeb cylinder by a translation along the axis to obtain a foliation of the solid torus normally called a Reeb component.

The leaf function of a Reeb component is not continuous since for any sequence x_n of interior points converging to a boundary point x one has that the sequence of leaves L_{x_n} converges to a cylinder M while the leaf L_x is a torus.

We notice that the cylinder M is a covering space of the torus leaf L_x. Furthermore one can choose a covering map from M to L_x in such a way that the image of the fundamental group of M is exactly the set of curves in L_x without holonomy.

Hence one sees that in this example the function

$$x \mapsto \widetilde{L_x}^{\mathrm{hol}}$$

associating to each point the holonomy covering of its leaf (see Section 4.9), is continuous.

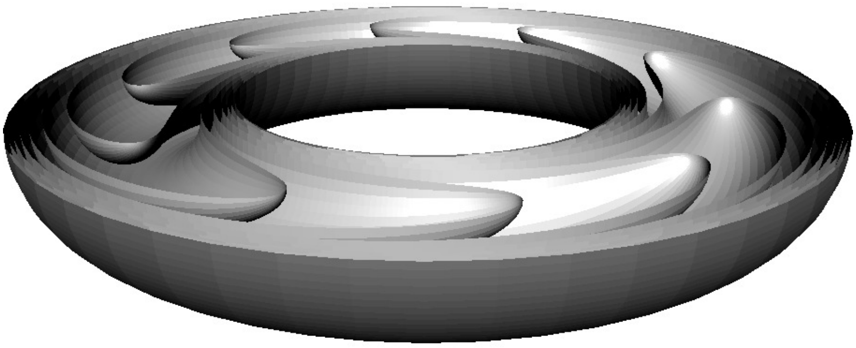

Figure 4.3: Half a Reeb component.

4.2.4 Example: the broken record foliation

Consider a foliation of the closed annulus which is obtained by pasting a copy of the vinyl record foliation with a trivially foliated annulus (i.e. foliated by parallel circles).

The holonomy of leaves in the trivially foliated annulus is trivial and hence they coincide with their holonomy covers. However a sequence of such leaves can be chosen to converge to the single circular leaf separating the two components. This leaf has non-trivial holonomy and hence its holonomy cover is isometric to $\mathbb{R}$. Hence in this example on sees that the function

$$x \mapsto \widetilde{L_x}^{\text{hol}}$$

is not continuous.

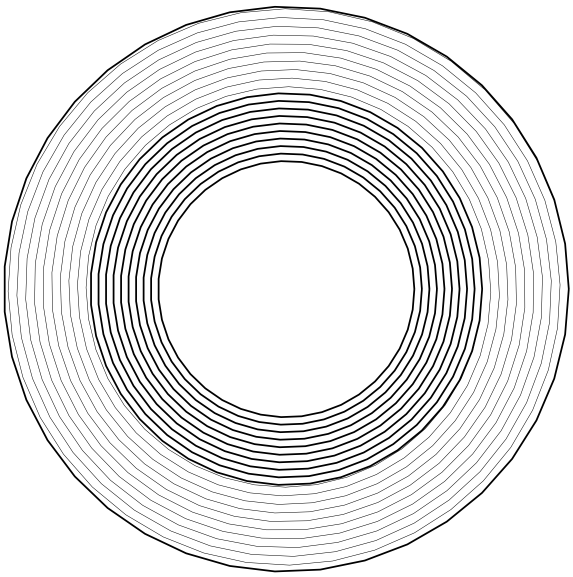

Figure 4.4: A broken record foliation. Circular leaves with trivial holonomy accumulate on a circular leaf with non-trivial holonomy.

4.2.5 Example: the Reeb transition

The following example was introduced by Reeb in [Ree48].

Consider the product Riemannian manifold $S^2 \times S^1 \times S^1$ where $S^2 = \{(x, y, z) \in \mathbb{R}^2 : x^2 + y^2 + z^2 = 1\}$ is the standard two-dimensional sphere, and $S^1 = \{z \in \mathbb{C} : |z| = 1\}$ the standard circle. We consider the coordinates $((x, y, z), e^{is}, e^{it})$ and the one forms

$$\begin{cases} \omega_1 &= \mathrm{d}t \\ \omega_2 &= \left((1 - \sin(t))^2 + x^2\right) \mathrm{d}s + \sin(t)\mathrm{d}x \end{cases}$$

The conditions of Frobenius' integrability theorem (see [CLN85, Theorem 2, pg. 185]) are satisfied and hence there is a unique two-dimensional foliation such that the tangent space of each leaf is contained in the kernel of ω_1 and ω_2. The equation $\omega_1 = 0$ for vectors tangent to the foliation implies that each leaf is contained in a set of the form $S^2 \times S^1 \times \{\text{constant}\}$ and hence we may consider the foliation as a family of foliations on $S^2 \times S^1$ parametrized by t.

When $\sin(t) = 0$ it is easy to verify that one obtains the foliation of $S^2 \times S^1$ by leaves of the form $S^2 \times \{\text{constant}\}$. However, when $\sin(t) = 1$ one has

$$\omega_2 = x^2 \mathrm{d}s + \mathrm{d}x$$

so that the torus in $S^2 \times S^1$ defined by $x = 0$ is a leaf, while the other leaves are planes parametrized by functions of the form

$$(x, y, z) \mapsto ((x, y, z), e^{i(c+1/x)})$$

on the hemispheres $x < 0$ and $x > 0$, for different values of the constant c.

Whenever $\sin(t) \neq 0$ one obtains a foliation of $S^2 \times S^1$ by spheres such that all leaves are obtained by applying a rotation to the S^1 components of a single leaf (i.e. they are all graphs of functions from S^2 to S^1 which in fact can be written explicitly).

Hence the set of spherical leaves is given by $\{\sin(t) \neq 1\}$, and the set of non-compact leaves is defined by $\{\sin(t) = 1, x \neq 0\}$.

One can explain this example geometrically. By pasting two copies of the partition of the solid torus $D \times S^1$ into closed disks $D \times \{\text{constant}\}$ one can obtain the trivial foliation of $S^2 \times S^1$ by leaves of the form $S^2 \times \{\text{constant}\}$. Pushing each disk at its center in the direction of the central circle of the solid torus one deforms the foliation but all leaves are still copies of S^2. This is done in such a way that the number of turns each disk does around the solid torus diverges, at which point the boundary torus becomes a leaf and we obtain a foliation of $S^2 \times S^1$ by two Reeb components. We call this process a Reeb transition.

Reeb noticed that in any such example there must be spherical leaves with arbitrarily large volume. We will show that this is a consequence of the regularity properties of the leaf function.

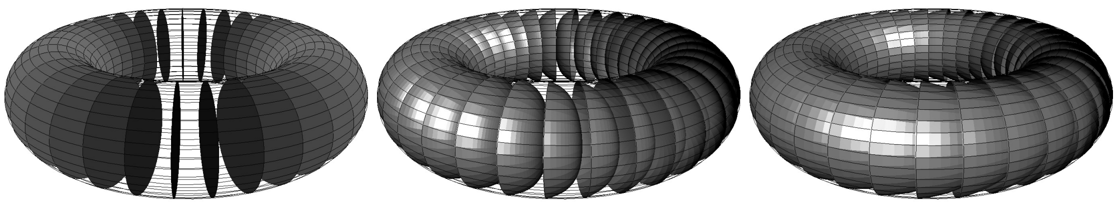

Figure 4.5: A Reeb transition: the trivial partition of a solid torus into disks is deformed into a Reeb component.

4.3 Regularity of leaf functions

In this section we state and prove our two main results after which we discuss applications to Reeb-type stability results and the Reeb transition example of the previous section.

4.3.1 Regularity theorems

A sequence of pointed complete connected Riemannian manifolds of the same dimension (M_n, o_n, g_n) is said to smoothly converge to a pointed complete Riemannian manifold (M, o, g) if there exists for each $r > 0$ a sequence of pointed smooth embeddings $f_n : B_r(o) \to M_n$ of the open ball of radius r centered at o into M_n defined for n large enough with the property that the pullback Riemannian metrics $f_n^* g_n$ converge smoothly to g on all compact subsets of $B_r(o)$ (see [Pet06, Chapter 10.3.2] and Section 4.7).

In principle smooth convergence of a sequence of manifolds is much stronger than $\mathcal{GS}$-convergence of the same sequence. However we can use compactness results from Riemannian geometry to obtain the following results.

Theorem 4.1 (Precompactness of the leaf function)**.** *Let X be a compact d-dimensional foliation. Then the leaf function of X takes values in a compact subset $\mathcal{M}$ of $\mathcal{GS}$ which contains only complete Riemannian manifolds of dimension d. Furthermore smooth and $\mathcal{GS}$-convergence are equivalent on $\mathcal{M}$.*

Proof. We establish in Section 4.8 that there exists $r > 0$ and a sequence C_k such that the injectivity radius of all leaves is at least r and the tensor norm of k-th derivative of the curvature tensor of any leaf is at most C_k.

Hence all leaves belong to the set $\mathcal{M}$ of (isometry classes of) pointed complete d-dimensional Riemannian manifolds with geometry bounded by $(r, \{C_k\})$ (see Section 4.4).

We establish in Theorem 4.11 that $\mathcal{M}$ is $\mathcal{GS}$-compact and that a sequence in $\mathcal{M}$ converges smoothly if and only if it $\mathcal{GS}$-converges. □

Corollary 4.2. *If x_n is a sequence converging to a point x in a compact foliation X and the sequence of leaves L_{x_n} $\mathcal{GS}$-converges to a pointed metric space M then, in fact, M is a smooth complete Riemannian manifold and L_{x_n} converges smoothly to M. In particular if M is compact then L_{x_n} is diffeomorphic to M for all n large enough.*

By a Riemannian covering we mean a pointed local isometry $f : M \to N$ between complete pointed Riemannian manifolds. If such a covering exists we say that M is a Riemannian covering (or just a covering) of N and that N is covered by M. See Section 4.9 for the definition of the holonomy covering of a leaf.

Theorem 4.3 (Semicontinuity of the leaf function)**.** *Let X be a compact foliation and x_n be a sequence converging to a point $x \in X$. If the sequence of leaves L_{x_n} $\mathcal{GS}$-converges to a pointed Riemannian manifold M then M is a Riemannian covering space of L_x and is covered by $\widetilde{L_x}^{hol}$.*

Proof. By Theorem 4.1 the leaf function takes values in a compact subspace of $\mathcal{GS}$ where Gromov-Hausdorff and smooth convergence are equivalent. Hence M is a complete Riemannian manifold and the sequence converges smoothly to M.

By smooth convergence (see Section 4.7), for each $r > 0$ there is a sequence of pointed embedding $f_{n,r} : B_r(o_M) \to L_{x_n}$ (defined for n large enough) such that $|f_{n,r}^* g_{L_{x_n}} - g_M|_{g_M}$ converges uniformly to 0 on $B_r(o_M)$. We show in Lemma 4.34 that this implies that the maps $f_{n,r}$ have a subsequence which converges locally uniformly to a local isometry $f_r : B_r(o_M) \to L_x$.

Now consider the family of functions $f_r : B_r(o_M) \to L_x$ when $r \to +\infty$. Since all these functions are local isometries one obtains a local isometry $f : M \to L_x$ as a the uniform limit on compact subsets f_{r_k} for some subsequence $r_k \to +\infty$. Hence M is a Riemannian covering of L_x via the covering map f.

Suppose that for some pair of distinct points $x, y \in M$ one has $f(x) = f(y)$ and let $\alpha : [0,1] \to M$ be a curve joining x and y. Take $r > 0$ large enough so that $B_r(o_M)$ contains $\alpha([0,1])$ and let $f_{n,r} : B_r(o_M) \to L_{x_n}$ be a sequence of embeddings as above which converges locally uniformly to f on $B_r(o_M)$.

Since each $f_{n,r}$ is injective and the pullback metrics converge to g_M the leafwise distance between $f_{n,r}(x)$ and $f_{n,r}(y)$ is bounded below by a positive constant for n large enough.

However since $f_{n,r} \circ \alpha$ converges uniformly to $f \circ \alpha$ we obtain that the holonomy along the closed curve $f \circ \alpha$ is non-trivial (see Corollary 4.31).

We have established that any closed curve in L_x having a lift under f which is not closed has non-trivial holonomy. In particular the lift of any curve with trivial holonomy in L_x is closed in M and hence the image of the fundamental group of M under f contains the subgroup of curves with trivial holonomy. By the classification of covering spaces (see Lemma 4.27) $\widetilde{L_x}^{\mathrm{hol}}$ is a Riemannian cover of M. □

4.3.2 Applications to continuity and Reeb stability

The main result of [EMT77] is that in any foliation the set of leaves without holonomy is residual. Combined with Theorem 4.3 we obtain that the leaf function is continuous on a residual set. Potential applications of this result to the study of quasi-isometry invariants of leaves are discussed by Álvarez and Candel in [ÁC03, Section 2].

Corollary 4.4 (Álvarez-Candel continuity theorem)**.** *The leaf function of any compact foliation is continuous on the set of leaves without holonomy. In particular the set of continuity points contains a residual set.*

Smooth convergence of a sequence to a compact manifold implies that the sequence elements are eventually diffeomorphic to the limit. Combined with Theorem 4.3 one obtains Reeb's local stability theorem (see [Ree47, Theorem 2]).

Corollary 4.5 (Reeb's local stability theorem)**.** *Let X be a compact foliation and $x \in X$ be such that $\widetilde{L_x}$ is compact. Then there exists a neighborhood U of x such that $\widetilde{L_y}$ is diffeomorphic to $\widetilde{L_x}$ for all $y \in U$.*

The same argument gives the usual generalization of Reeb's stability theorem to compact leaves with trivial or finite holonomy (see for example [CLN85, pg. 70]).

Corollary 4.6 (Stability of compact leaves with finite holonomy)**.** *Let X be a compact foliation and $x \in X$ be such that $\widetilde{L_x}^{hol}$ is compact. Then there is a neighborhood U of x such that for each $y \in U$ the leaf L_y is compact and diffeomorphic to a covering space of L_x.*

We say X is a foliation by compact leaves if all leaves are compact. The volume function of such a foliation is the function

$$x \mapsto \mathrm{vol}\,(L_x)$$

associating to each leaf its volume (which is finite). Since a Riemannian covering has larger volume then the space it covers one obtains the following.

Corollary 4.7 (Volume function semicontinuity)**.** *Let X be a compact foliation by compact leaves. Then the volume function of X is lower semicontinuous.*

Notice that since any sequence of leaves has a smoothly convergent subsequence we obtain the following part of Epstein's structure theorem (see [Eps76, Theorem 4.3]).

Corollary 4.8 (Epstein)**.** *Let X be a compact foliation by compact leaves whose volume function is bounded. Then every point $x \in X$ has a neighborhood U such that for all $y \in U$ the leaf L_y is diffeomorphic to a finite covering of L_x.*

We say a foliation X has codimension k if it admits an atlas by foliated charts $\{h_i : U_i \to \mathbb{R}^d \times T_i, i \in I\}$ with $T_i = \mathbb{R}^k$ for all i. Under this hypothesis X is automatically a topological manifold.

Notice that any holonomy transformations of a codimension one foliation will be a homeomorphism between two open subsets of $\mathbb{R}$. We say a codimension one foliation is transversally orientable if every holonomy transformation associated to a closed chain of compatible charts is increasing.

The following elementary lemma implies that in a transversally oriented codimension one foliation by compact leaves all leaves have trivial holonomy (here we denote by $f^n(x) = f(f(\cdots f(x)\cdots))$ the n-th iterate of the point x under the function f and notice that in order for it to be well defined $f^k(x)$ must belong to the domain of f for all $k = 0, \ldots, n-1$):

Lemma 4.9. *Let $h : U \to V \subset \mathbb{R}$ be an increasing homeomorphisms between two neighborhoods of $0 \in \mathbb{R}$ such that $h(0) = 0$. Then either h is the identity map or there exists $x \in U$ and $f = h^{\pm 1}$ such that the set $\{f^n(x) : n \geq 0\}$ is well defined and infinite.*

Epstein established in [Eps72] that a flow on a 3-manifold for which all orbits are periodic has the property that the periods are bounded. This was later generalized to state that compact codimension two foliations by compact leaves have bounded volume functions (see [EMS77]). Notice that these results are very subtle since they are false for foliations of codimension 3 or more (see [EV78]). The codimension one case follows directly from our results and the above elementary lemma.

Corollary 4.10. *Let X be a connected compact transversally oriented codimension one foliation. Then the leaf function of X is continuous. In particular all leaves are diffeomorphic and the volume function is continuous.*

For tranversally oriented codimension one foliations of connected manifolds Reeb's local stability combines with properties of one dimensional dynamics in the spirit of the lemma above to yield Reeb's global stability theorem (see [Ree47, Theorem 3] and [CLN85, pg. 72]) which states that if a leaf has a compact universal cover than all leaves are diffeomorphic.

In view of these results one might conjecture that the set of leaves with compact universal cover, besides being open, is always closed. However this is false as shown by the Reeb transition example given in Section 4.2.5. We will now discuss some aspects of this example.

The fact that in the Reeb transition there must be spheres with arbitrarily large volume follows from Corollary 4.4 and the smooth convergence given by Theorem 4.1. To see this consider a sequence of points x_n belonging to compact leaves which converge to a point x whose leaf is non-compact and notice that the sequence of manifolds L_{x_n} smoothly converge to L_x.

Consider now in the same example a sequence x_n on spherical leaves which converges to a point x on the single torus leaf. By Theorem 4.3 any smooth limit point of the sequence L_{x_n} must either be a finite covering of the torus L_x or the cylinder $\widetilde{L_x}^{\mathrm{hol}}$. The first case is impossible because convergence to a compact limit would imply that the manifolds in the sequence L_{x_n} are eventually diffeomorphic to the limit manifold which would have to be a torus. Hence the sequence of spheres L_{x_n} converges smoothly to the cylinder $\widetilde{L_x}^{\mathrm{hol}}$.

4.4 Uniformly bounded geometry

In this section we prove that certain subsets of $\mathcal{GS}$ consisting of manifolds with 'uniformly bounded geometry' are compact and that furthermore smooth and $\mathcal{GS}$-convergence coincide on them. This result was used in the proof of Theorem 4.1.

4.4.1 Spaces of manifolds with uniformly bounded geometry

We say a complete d-dimensional Riemannian manifold has geometry bounded by $r > 0$ and a sequence C_k if the injectivity radius of M is at least r at all points and the curvature tensor of M satisfies

$$|\nabla^k R| \le C_k$$

for all k, where ∇ denotes the covariant derivative and we are using the tensor norms induced by the Riemannian metric.

We use $\mathcal{M}(d, r, \{C_k\})$ to denote the subset of $\mathcal{GS}$ consisting of all isometry classes of d-dimensional complete pointed Riemannian manifolds with geometry bounded by r and the sequence C_k.

An element of $\mathcal{M}(d, r, \{C_k\})$ is represented by a triplet (M, o_M, g_M) and two triplets represent the same element if there is a pointed isometry between them. We will sometimes write $M \in \mathcal{M}(d, r, \{C_k\})$ in which case it is implied that the basepoint will be denoted by o_M and the Riemannian metric by g_M.

4.4.2 A smooth compactness theorem

Usually $\mathcal{GS}$-convergence of a sequence of manifolds is much weaker than smooth convergence. However we will show they are equivalent on sets of manifolds with uniform bounded geometry.

To understand this it might be helpful to consider the following fact: Let $\mathcal{F}$ be a C^1 compact family of functions from the interval $[0, 1]$ to $\mathbb{R}$. Then if a sequence f_n in $\mathcal{F}$ converges uniformly to a limit f, in fact $f \in \mathcal{F}$ and the derivatives f'_n converge uniformly to f'.

The proof can also be thought of as an application of the fact that a continuous bijective function whose domain is compact has a continuous inverse (in the setting of the previous paragraph the domain would be $\mathcal{F}$ with the C^1-topology the codomain would be the same set with the C^0-topology and function would be the identity). The difficulty in our case is in establishing compactness of the domain plus a subtle technical point which is discussed immediately after the proof.

We will now state the main result of this section.

Theorem 4.11. *Let $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ for some choice of dimension d, radius r, and sequence C_k. Then $\mathcal{M}$ is a compact subset of $\mathcal{GS}$ on which $\mathcal{GS}$-convergence and smooth convergence are equivalent.*

Proof with gap. The proof rests on the following facts

1. The set $\mathcal{M}$ is precompact with respect to smooth convergence.

2. The set $\mathcal{M}$ is closed under smooth convergence.

3. Smooth convergence implies pointed Gromov-Hausdorff convergence.

We will establish facts 1 and 2 in sections 4.5 and 4.6 respectively.

Fact 3 is generally accepted (e.g. see [Pet06, Section 10.3.2] and [BBI01, Section 7.4.1]) but we include a proof in the next subsection for completeness.

Using these facts the proof proceeds as follows.

Given a sequence M_n in $\mathcal{M}$ we may, using smooth precompactness, extract a smoothly convergent subsequence M_{n_k} with limit M. Since $\mathcal{M}$ is closed under smooth convergence we have $M \in \mathcal{M}$. Finally, since smooth convergence implies pointed Gromov-Hausdorff convergence one has

$$\lim_{n \to +\infty} \mathrm{d}_{\mathcal{GS}}(M_{n_k}, M) = 0.$$

This establishes that $\mathcal{M}$ is a compact subset of $\mathcal{GS}$.

Suppose now that some sequence M_n in $\mathcal{M}$ converges in the pointed Gromov-Hausdorff sense to $M \in \mathcal{M}$. Since any subsequence of M_n will have a further subsequence which converges smoothly and any smooth limit must in fact coincide with M we obtain that the original sequence M_n converges smoothly to M. $\square$

There is a gap in the above proof which is illustrated by the following example (see Figure 4.6).

Consider the sequence of functions indexed on finite strings of zeros and ones defined by

$$f_{a_1 \ldots a_r} : [0, 1] \to \mathbb{R}$$

$$f_{a_1 \ldots a_r}(x) = \begin{cases} 1 & \text{if } \sum_{k=1}^{r} a_k 2^{-k} < x < 2^{-r} + \sum_{k=1}^{r} a_k 2^{-k}. \\ 0 & \text{otherwise.} \end{cases}$$

The sequence does not converge Lebesgue almost surely to any function. However any subsequence has a further subsequence which converges almost surely to 0. In particular the arguments in our proof above would imply that L^2 convergence and almost sure convergence coincide on the set of functions $\{0\} \cup \{f_{a_1 \ldots a_r}\}$ but this conclusion is false.

To exclude this type of behavior it suffices to show that smooth convergence comes from a topology. We do this in Section 4.7.

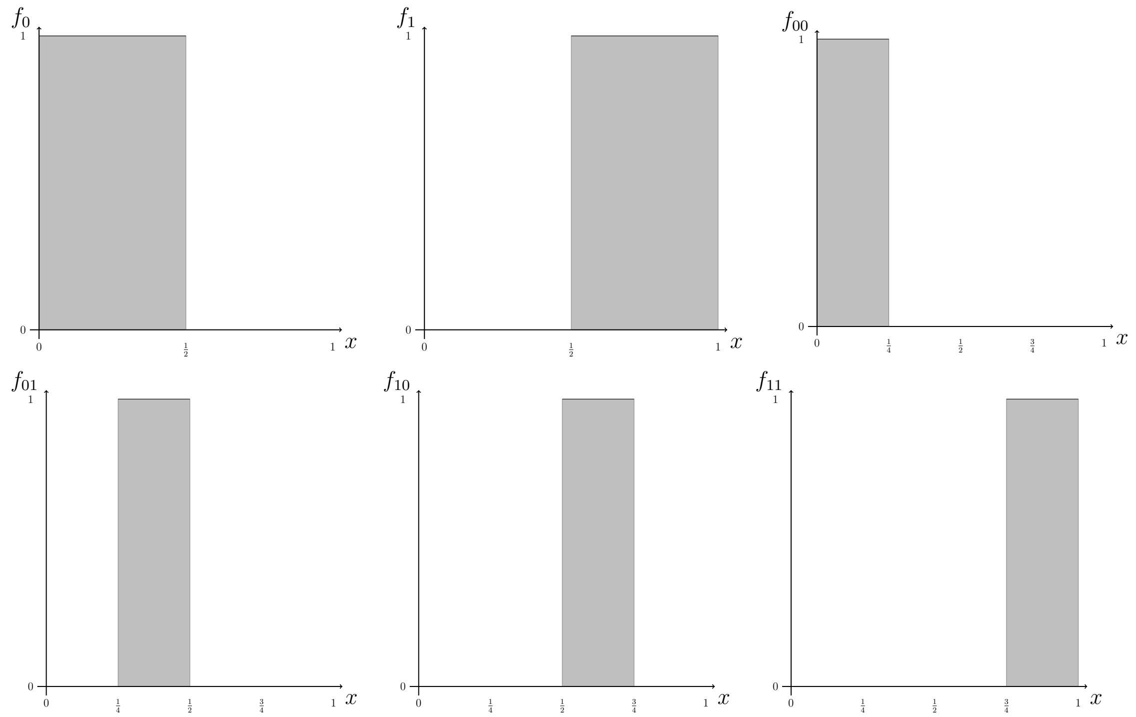

Figure 4.6: Six elements of a sequence of functions which does not converge almost surely to 0 but has no other limit points.

4.4.3 Smooth vs Gromov-Hausdorff convergence

For the readers convenience we present a proof of the fact that smooth convergence is stronger than $\mathcal{GS}$-convergence. The key ideas are contained in the proof of part 2 of [BBI01, Theorem 7.3.25] and the indications given in Section 7.4.1 of the same reference.

Lemma 4.12. *If a sequence (M_n, o_n, g_{M_n}) converges smoothly to (M, o, g) then it also $\mathcal{GS}$-converges to the same limit.*

Proof. We must show that for each $r > 0$ the sequence of pointed compact metric spaces $\overline{B_r}(o_n)$ (where the metric is inherited from M_n) converges in the Gromov-Hausdorff sense to $\overline{B_r}(o)$.

By smooth convergence (see Section 4.7) given $r > 0$ there exists a sequence of smooth pointed embeddings $f_n : B_{3r}(o) \to M_n$ with the property that the pullback metrics $g_n = f_n^* g_{M_n}$ satisfy

$$a_n = \sup\{|g_n(x) - g(x)|_g : x \in B_{3r}(o)\} \to 0$$

when $n \to +\infty$.

Notice that whenever $a_n = 0$ one has that $\overline{B_r}(o)$ is isometric to $\overline{B_r}(o_n)$ via f_n so that there is nothing to prove. Hence we may assume without loss of generality in what follows that $a_n \neq 0$. Also, since we are only interested in behavior when $n \to +\infty$ we may assume that $a_n < 1$.

Let d be the Riemannian distance of M and d_n be the pullback under f_n of the Riemannian distance on M_n. Since the shortest curve between $f_n(x)$ and $f_n(y)$ might in principle exit $f_n(B_{3r}(o))$ it is not necessarily true that d_n equals the distance on $B_{3r}(o)$ induced by the metric g_n.

However notice that if v is a tangent vector in $B_{3r}(o)$ of unit norm for g then

$$|g_n(v,v) - 1| \leq a_n$$

so that the g_n norm of v is between $(1-a_n)^{1/2}$ and $(1+a_n)^{1/2}$. This implies that the g_n-length of any curve in $B_{3r}(o)$ is within a multiplicative factor $b_n^{\pm 1}$ of its g-length where $b_n = \max\left((1-a_n)^{-1/2}, (1+a_n)^{1/2}\right)$. In particular, for n large enough, the Riemannian distance induced by g_n on $B_{3r}(o)$ coincides with d_n when restricted to $\overline{B_r}(o)$.

The previous comparison of lengths of curves also implies for n large that

$$|d_n(x,y) - d(x,y)| \leq (b_n - 1)d(x,y) \leq 2r(b_n - 1)$$

for all $x, y \in \overline{B_r}(o)$ (the first inequality relies on the fact that $1 - b_n^{-1} \leq b_n - 1$ which is true since $b_n \geq 1$).

Following the proof of part 2 of [BBI01, Theorem 7.3.25] we consider for each n the distance $\tilde{d}_n$ on the disjoint union $\overline{B_r}(o) \sqcup \overline{B_r}(o)$ which coincides with d on the left-hand copy, with d_n on the right-hand copy and for x, z in different copies is defined by

$$\tilde{d}_n(x,z) = \inf\left\{d(x,y) + 2r(b_n - 1) + d_n(y,z) : y \in \overline{B_r}(o)\right\}.$$

The Hausdorff distance between the two copies of $\overline{B_r}(o)$ with the above defined distance is less than $3r(b_n - 1)$ and therefore goes to 0 when $n \to +\infty$. This shows that the Gromov-Hausdorff distance between $\overline{B_r}(o)$ and $f_n(\overline{B_r}(o))$ (the later inheriting its metric from M_n) converges to 0 when $n \to +\infty$.

To conclude it suffices to establish that the Hausdorff distance (with respect to the Riemannian distance on M_n) between $f_n(\overline{B_r}(o))$ and $\overline{B_r}(o_n)$ goes to 0 when $n \to +\infty$. This follows from our comparison of d and d_n since $f_n(\overline{B_r}(o))$ contains the ball of radius $b_n^{-1}r$ and is contained in the ball of radius $b_n r$ centered at o_n. $\square$

4.5 Smooth precompactness

In this section we prove that sets of manifolds with uniformly bounded geometry are precompact with respect to smooth convergence. This was used in the proof of Theorem 4.11.

We recall (see [Pet06, Chapter 10] and Section 4.7) that, in similar fashion to the definition of smooth convergence, a sequence of complete Riemannian manifolds (M_n, o_n, g_n) is said to converge C^k to (M, o, g) if for each $r > 0$ there exists a sequence of smooth pointed embeddings $f_n : B_r(o) \to M_n$ (defined for large enough n) such that the pullback metrics $f_n^* g_n$ converge C^k to g on compact subsets of $B_r(o)$.

Lemma 4.13. *All subsets of $\mathcal{GS}$ of the form $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ are sequentially precompact with respect to smooth convergence.*

Proof. For each $M \in \mathcal{M}$ we consider the atlas $\mathcal{A}$ by normal coordinates on the balls of radius r' given by Lemma 4.16 below.

A theorem of Eichhorn (see Lemma 4.15 below) shows that there exists a sequence C_{nor}^k such that all the metrics on $B_{r'}$ obtained from such coordinates have coefficients which satisfy

$$|\partial_{i_1} \cdots \partial_{i_k} g_{ij}| \leq C_{\text{nor}}^k$$

for all choices of indices $i_1, \ldots, i_k$.

Furthermore we establish in Lemma 4.16 that there is a sequence C^k_{tran} bounding the k-th order partial derivatives of the transition maps of any such atlas $\mathcal{A}$ and that the Euclidean and Riemannian norms on $B_{r'}$ differ at most by a multiplicative factor of $2^{\pm 1/2}$.

This shows that for each k there exists Q such that all manifold in $\mathcal{M}$ have C^k norm less than or equal to Q on a scale of r in the sense of Petersen (see the definition in subsection 4.5.1 below).

Applying Petersen's compactness theorem (see Theorem 4.14 below) one obtains that $\mathcal{M}$ is C^k precompact for all k, and hence smoothly precompact as claimed. □

4.5.1 Norms and sequential compactness

Following [Pet06, Chapter 10.3.1] (taking, for simplicity, $\alpha = 1$ in his notation) we say that a manifold M has C^k-norm less than or equal to Q on a scale of r if there exists an atlas $\mathcal{A}$ of M which satisfies the following properties:

1. Every ball of radius $e^{-Q}r/10$ is contained in the domain of some chart in $\mathcal{A}$.
2. For each chart $\varphi \in \mathcal{A}$ one has $|D\varphi| \le e^Q$ and $|D\varphi^{-1}| \le e^Q$, where $D\varphi$ is the tangent map to the chart and one uses the operator norm between the tangent space of M with the Riemannian metric and Euclidean space with the usual Euclidean metric.
3. For each chart $\varphi \in \mathcal{A}$ and each $0 \le i \le k$ the partial derivatives of order i of the coefficients of $\varphi_* g_M$ are $Q/(r^{i+1})$-Lipschitz.
4. For each $\varphi_1, \varphi_2 \in \mathcal{A}$ the C^{k+2}-norm (i.e. sum of suprema of absolute values of all partial derivatives up to order $k+2$) of the transition map $\varphi_2 \circ \varphi_1^{-1}$ is less than or equal to $(10 + r)e^Q$.

We now restate Petersen's [Pet06, Theorem 72] as we will use it.

Theorem 4.14 (Petersen). *For any positive constants r and Q the class of pointed, complete, d-dimensional Riemannian manifolds with C^k-norm less than or equal to Q on a scale of r is sequentially compact with respect to C^k convergence.*

4.5.2 Normal coordinates

We recall that a normal parametrization of a manifold $M \in \mathcal{M}(d, r, \{C_k\})$ at a point p is a function $\psi : \mathbb{R}^d \to M$ satisfying

$$\psi(x) = \exp_{\psi(0)} \circ f(x)$$

where $\exp : T_{\psi(0)}M \to M$ is the Riemannian exponential map and $f : \mathbb{R}^d \to T_{\psi(0)}M$ is a linear isometry between $\mathbb{R}^d$ and the tangent space $T_{\psi(0)}M$ at $\psi(0)$.

If $M \in \mathcal{M}(d, r, \{C_k\})$ then any normal parametrization ψ is a diffeomorphism when restricted to the ball B_r of radius r centered at $0 \in \mathbb{R}^d$. Hence the pullback $g = \psi^* g_M$ of the Riemannian metric of M to B_r is also a Riemannian metric (i.e. non-degenerate).

We recall that the coefficients of a metric g defined on some open subset of $\mathbb{R}^d$ are the functions

$$x \mapsto g(x)(e_i, e_j) = g_{ij}(x)$$

where $e_1, \ldots, e_d$ is the canonical basis of $\mathbb{R}^d$.

The coefficients obtained in this manner from manifolds in $\mathcal{M}(d, r, \{C_k\})$ are uniformly C^k bounded as is shown by the following lemma (see [Eic91, Corollary 2.6]).

Lemma 4.15 (Eichhorn). *Given $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ for each $k \geq 0$ there exists a constant C^k_{nor} such that if $g = \psi^* g_M$ is a metric on B_r obtained by pulling back the metric of some manifold $M \in \mathcal{M}$ via a normal parametrization ψ then one has:*

$$|\partial_{i_1} \cdots \partial_{i_k} g_{ij}| \leq C^k_{nor}$$

for all indices $i, j, i_1, \ldots, i_k$.

4.5.3 Transition maps

This subsection is devoted to establishing the following uniform estimate for the derivatives of transition maps between normal coordinates.

Lemma 4.16. *Given $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ there exists $r' < r$ and for each $k \geq 0$ a constant C^k_{tran} such that the k-th order partial derivatives of any transition map between normal coordinates on balls of radius r' in any manifold $M \in \mathcal{M}$ are bounded in absolute value by C^k_{tran}.*

For partial derivatives of order one and two the above result can be compared to Lemma 3.4 and Lemma 4.3 of [Che70].

The first derivative of the change of coordinates between maximal normal coordinates based at the north and south pole on the standard two dimensional sphere is not bounded. This shows that it's indeed necessary to take $r' < r$ in the above lemma.

Our proof proceeds in three steps. First we bound the k-th order covariant derivative of any curve of the form $t \mapsto x + tv$ for any metric on the Euclidean ball of radius r' in $\mathbb{R}^d$ obtained by pullback from a normal parametrization of a manifold in $\mathcal{M}$. Second, we bound the the actual (Euclidean) k-th order derivative of any curve whose covariant derivatives satisfy the previously obtained bounds (the point here being that covariant derivatives are invariant under the transition maps). Finally, combining the preceding result one obtains a bound for the k-th derivative of any transition map along any straight line which implies the same bound is satisfied for the partial derivatives of order k (this amounts to the statement that a symmetric k-linear function attains its maximum norm on the diagonal, see [Wat90] for a proof).

To begin we recall that the Christoffel symbols of a metric on an open subset of $\mathbb{R}^d$ with coefficients g_{ij} are given by

$$\Gamma^k_{ij} = \frac{1}{2} g^{kl} \left(\partial_j g_{il} + \partial_i g_{lj} - \partial_l g_{ij} \right)$$

where g^{ij} are the coefficients of the inverse of the matrix (g_{ij}) and summation is implied over the repeated indices of each term.

In what follows we use B_s for the open Euclidean ball of radius s centered at $0 \in \mathbb{R}^d$.

Lemma 4.17. *Given $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ there exists $r' < r$ and for each $k \geq 0$ a constant C'_k such that for any metric g on $B_{r'}$ obtained by pullback from a normal parametrization of a manifold $M \in \mathcal{M}$ one has:*

1. *The k-th order partial derivatives of the metric coefficients g_{ij}, the coefficients of the inverse matrix g^{ij}, and the Christoffel symbols Γ^l_{ij}, are bounded in absolute value by C'_k for all k.*

2. *For all $v \in \mathbb{R}^d$ and $x \in B_{r'}$ one has $2^{-1}|v| \leq |v|_{g(x)} \leq 2|v|$ where $|v|$ is the Euclidean norm of v and $|v|_{g(x)}$ its norm with respect to the inner product $g(x)$.*

Proof. Notice that for any of the coefficients g_{ij} under consideration one has $(g_{ij}(0)) = (\delta_{ij})$ where the right-hand side is the $d \times d$ identity matrix. Let K be a compact neighborhood of the identity matrix such that any inner product whose matrix of coefficients (i.e. the matrix whose entry in the i-th row and j-th column is the inner product between the i-th and j-th vectors of the canonical basis of $\mathbb{R}^d$) is in K satisfies property 2 above.

Since one has a uniform bound C^1_{nor} (given by Lemma 4.15) for the first order derivatives of g_{ij} on B_r there exits $r' < r$ (depending only on this C^1_{nor}) such that for all the metrics under consideration $(g_{ij}(x))$ belongs to K for all $x \in B_{r'}$.

By Lemma 4.15 one has uniform bounds on the partial derivatives of the metric coefficients g_{ij} on B_r (and in particular on $B_{r'}$). Combining this with the fact that matrix inversion is smooth on K one obtains uniform bounds for the partial derivatives of all orders of the inverse matrix (g^{ij}) on $B_{r'}$. From this one can bound the partial derivatives of the Christoffel symbols as well. □

The covariant derivative of a vector field $v(t)$ over a curve $x(t)$ in $\mathbb{R}^d$ with respect to a metric with Christoffel symbols Γ^k_{ij} is given by

$$\nabla_{x'} v = v' + \Gamma^k_{ij}(x^i)' v^j e_k \tag{4.1}$$

where a superscript i denotes the i-th coordinate and $'$ denotes derivative with respect to t. We convene that $\nabla^0_{x'} v(t) = v(t)$ and define inductively $\nabla^{k+1}_{x'} v(t) = \nabla_{x'} \nabla^k_{x'} v(t)$.

Lemma 4.18. *Fix $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ and let C'_k and r' be given by Lemma 4.17. There exists a sequence C''_k such that for any metric g on $B_{r'}$ obtained by pullback from a normal parametrization of a manifold $M \in \mathcal{M}$ and any curve of the form*

$$x(t) = x_0 + tv$$

where $x_0 \in B_{r'}$ and $|v| = 1$ one has

$$|\nabla^k_{x'} x'|_g \leq C''_k$$

for all $k \geq 0$.

Proof. From Lemma 4.17 the Riemannian norm of v is bounded by 2 at all points in $B_{r'}$. This shows that one can take $C''_0 = 2$.

In order to bound the higher order covariant derivatives define inductively

$$v_2(t) = \nabla_{x'} v = v^i v^j \Gamma^k_{ij} e_k$$

and

$$v_{n+1}(t) = \nabla_{x'} v_n(t) = v'_n + v^i v^j_n \Gamma^k_{ij} e_k.$$

Since the coordinates v^i of v are constants of absolute value less than or equal to 1 the Euclidean norm of v_{n+1} can be bounded in terms of that of v_n and the derivatives of the Christofell symbols. This is possible and is equivalent to bounding the Riemannian norm due to Lemma 4.17. □

We denote by $x^{(k)}(t)$ denote the k-th (Euclidean) derivative of a curve in $\mathbb{R}^d$.

Lemma 4.19. *Fix $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ and let C_k'' and r' be given by Lemma 4.17. There exists a sequence C_k''' such that for any metric g on $B_{r'}$ obtained by pullback from a normal parametrization of a manifold $M \in \mathcal{M}$ and any curve $x(t)$ satisfying*

$$|\nabla_{x'}^k x'|_g \leq C_k''$$

for all $k \geq 0$ one has

$$|x^{(k)}(t)| \leq C_k'''$$

for all $k \geq 0$.

Proof. By Lemma 4.17 the Euclidean and Riemannian norms differ at most by a factor of $2^{\pm 1/2}$.

In particular one can take $C_0''' = 2C_0''$ and the Euclidean norm of

$$\nabla_{x'} x' = x'' + \Gamma_{ij}^k (x^i)'(x^j)' e_k$$

is bounded by $2C_1''$.

Since one has $|x'| \leq 2C_0''$ one obtains from the last equation a bound for $|x''|$.

The higher order case follows by induction since there is a single term in $\nabla_{x'}^k x'$ which is equal to $x^{(k+1)}$ and the rest can be bounded in terms of lower order derivatives of x and the derivatives of the Christoffel symbols. □

We now complete the final step for the proof of Lemma 4.16.

Lemma 4.20. *Let $f : U \subset \mathbb{R}^d \to \mathbb{R}^d$ be a smooth function satisfying*

$$|g^{(k)}(0)| \leq C_k''$$

for all $k \geq 0$ and g of the form $g(t) = f(x + tv)$ with $|v| = 1$ and $x \in U$. Then for all $x \in U$ one has

$$|\partial_{i_1} \cdots \partial_{i_k} f(x)| \leq C_k''$$

for all $k \geq 0$ and $i_1, \ldots, i_k \in \{1, \ldots, d\}$.

Proof. Define inductively

$$D_x f(v) = \lim_{h \to 0} \frac{f(x + hv) - f(x)}{h}$$

$$D_x^2 f(v_1, v_2) = \lim_{h \to 0} \frac{D_{x+hv_1} f(v_2) - D_x f(v_2)}{h}$$

$$D_x^{k+1} f(v_1, \ldots, v_{k+1}) = \lim_{h \to 0} \frac{D_{x+hv_1}^k f(v_2, \ldots, v_{k+1}) - D_x^k f(v_2, \ldots, v_{k+1})}{h}.$$

Letting $P_x^k f(v) = D_x^k f(v, \ldots, v)$ we have by hypothesis and multilinearity that $|P_x^k f(v)| \leq C_k'' |v|^k$.

Since partial derivatives commute the multilinear function $D_x^k f : (\mathbb{R}^d)^k \to \mathbb{R}^d$ is symmetric and $P_x^k f$ determines $D_x^k f$ by polarization. This implies a bound for the mixed partial derivatives, and in fact one has $|D_x^k f(v_1, \ldots, v_k)| \leq C_k'' |v_1| \cdots |v_k|$ as shown in [Wat90]. □

4.6 Curvature and injectivity radius

In this section we prove that sets of manifolds with uniform bounded geometry are closed with respect to smooth convergence. This was used in the proof of Theorem 4.11.

Lemma 4.21. *Suppose $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ for some value of the parameters. If (M_n, o_n, g_n) is a sequence in $\mathcal{M}$ converging smoothly to (M, o, g) then $M \in \mathcal{M}$.*

Proof. The fact that the injectivity radius of M is larger than or equal to r follows because the injectivity radius is upper semicontinuous with respect to smooth convergence as we will show in the next subsection (see Lemma 4.22).

We will now establish that M satisfies the curvature bounds

$$|\nabla^k R|_g \le C_k.$$

Let (g_{ij}) be the matrix of coefficients of a metric g on an open subset of $\mathbb{R}^d$ and (g^{ij}) the inverse matrix. The g norm of a (p, q) tensor field [1]

$$T = a_{i_1,\ldots,i_q}^{i_{q+1},\ldots,i_{p+q}} e^{i_1} \otimes \cdots \otimes e^{i_q} \otimes e_{i_{q+1}} \otimes \cdots e_{i_{p+q}}$$

(where we denote by e_i the canonical basis and e^i the dual basis of $\mathbb{R}^d$) is given by

$$|T|_g^2 = a_{i_1,\ldots,i_q}^{i_{q+1},\ldots,i_{p+q}} a_{j_1,\ldots,j_q}^{j_{q+1},\ldots,j_{p+q}} g^{i_1 j_1} \cdots g^{i_q j_q} g_{i_{q+1} j_{q+1}} \cdots g_{i_{p+q} j_{p+q}}.$$

The curvature tensor of g is the $(1, 3)$-tensor field $R = R_{ijk}^l e^i \otimes e^j \otimes e^k \otimes e_l$ given by (e.g. see [Bre10, Section 5])

$$R_{ijk}^l = \partial_j \Gamma_{ki}^l - \partial_k \Gamma_{jk}^l + \Gamma_{jm}^k \Gamma_{ki}^m - \Gamma_{km}^l \Gamma_{ji}^m$$

where the Christoffel symbols Γ_{ij}^k are defined by

$$\Gamma_{ij}^k = \frac{1}{2} g^{kl} (\partial_i g_{il} + \partial_j g_{jl} - \partial_l g_{ij}).$$

Since matrix inversion is smooth the two formulas above prove that if a sequence of metrics g_n converges uniformly on compact sets to g then the norm of their curvature tensors converge pointwise to that of g.

Similarly, for each k the covariant derivative $\nabla^k R$ is a $(1, 3+k)$-tensor field whose coefficients are smooth functions of the partial derivatives of the coefficients g_{ij} and g^{ij}. This shows that the bound $|\nabla^k R| \le C_k$ passes to the limit when a sequence of manifolds converges C^∞ to another. Hence one has that the limit manifold M of the the sequence M_n also satisfies these bounds. □

4.6.1 Semicontinuity of the injectivity radius

Continuity of the injectivity radius with respect to a varying family of metrics on a single compact manifold was established in [Ehr74] and [Sak83].

The injectivity radius is not continuous under smooth convergence of pointed manifolds as can be seen by considering a metric g on $\mathbb{R}^2$ which has finite injectivity radius but is flat outside of a compact set. In this setting the sequence of pointed manifolds $(\mathbb{R}^2, x_n, g)$

1. In all tensor calculations we use the convention that summation is implied over indices which are repeated in a term

will smoothly converge to $\mathbb{R}^2$ endowed with the Euclidean metric if $x_n \to \infty$ when $n \to +\infty$. Hence we have a sequence of manifolds with finite and constant injectivity radius converging to a manifold whose injectivity radius is infinite.

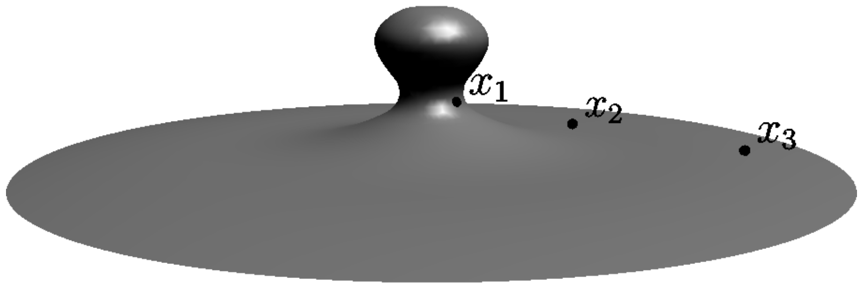

Figure 4.7: An asymptotically flat surface with finite injectivity radius. Changing the basepoint gives an sequence converging to a limit whose injectivity radius is infinite.

However, upper semicontinuity still holds as we will now show.

Lemma 4.22 (Semicontinuity of the injectivity radius)**.** *The injectivity radius is upper semicontinuous with respect to smooth convergence.*

Proof. Suppose for the sake of contradiction that there is a sequence (M_n, o_n, g_n) with the injectivity radius of each term larger than or equal to some $r > 0$ which converges smoothly to a manifold (M, o, g) whose injectivity radius is strictly less than r.

By Proposition 19 and Lemma 14 of [Pet06, pg. 139-142], there exists a geodesic $\alpha : [0,1] \to M$ of length $L < r$ and some other smooth curve $\beta : [0,1] \to M$ with the same endpoints with length $L' < L$.

By the definition of smooth convergence there is an open set Ω containing $\alpha([0,1])$ and $\beta([0,1])$ and a sequence of pointed embeddings $f_n : \Omega \to M_n$ such that $f_n^* g_n$ converges C^∞ to g on compact subsets of Ω.

Consider $\alpha_n : [0,1] \to M$ the geodesic for the metric $f_n^* g_n$ with initial condition $\alpha'(0)$. We claim that $\alpha_n(1) \to \alpha(1)$ and that the $f_n^* g_n$ length of α_n converges to L when $n \to +\infty$. By covering $\alpha([0,1])$ with a finite number of charts and noticing that in each chart the coefficients of $f_n^* g_n$ converge C^∞ on compact sets to those of g, this follows from continuity of solutions to ordinary differential equations with respect to the vector field (see [DK00, Theorem B3, pg. 333]). We omit further details.

Smooth convergence of $f_n^* g_n$ to g implies that the $f_n^* g_n$ length of β converges to L' and the $f_n^* g_n$ distance between $\beta(1)$ and $\alpha_n(1)$ converges to 0.

Hence for n large enough the manifold M_n contains a geodesic of length strictly less than r which is not the shortest curve between its endpoints. By the Hopf-Rinow theorem we will find two geodesics of length strictly less than r joining the same endpoints in M_n contradicting the fact that the injectivity radius of M_n is larger than or equal to r. $\square$

4.7 Smooth convergence and tensor norms

In this section we discuss in detail the definition of C^k and smooth convergence of pointed Riemannian manifolds. In particular we provide a coordinate free definition of convergence in terms of tensor norms and prove that it is equivalent to definition given in [Pet06, Chapter 10.3.2].

We also establish that C^k and smooth convergence on certain subsets of $\mathcal{GS}$ comes from a topology, a fact that was used in the proof of Theorem 4.11.

4.7.1 Coordinate free definition of convergence

Following [Pet06, 10.3.2] a sequence (M_n, o_n, g_n) of pointed connected complete Riemannian manifolds is said to converge C^k to (M, o, g) if for every $r > 0$ there exists a domain Ω containing $B_r(o)$ and (for n large enough) a sequence of pointed embeddings $f_n : \Omega \to M_n$ such that $f_n(\Omega) \supset B_r(o_n)$ and $f_n^* g_n$ converges C^k to g on compact subsets of Ω. Smooth convergence is by definition C^k convergence for all k.

Recall that the coefficients of a Riemannian metric g defined on an open subset U of $\mathbb{R}^d$ are the functions

$$x \mapsto g(x)(e_i, e_j) = g_{ij}(x)$$

where $e_1, \ldots, e_d$ is the canonical basis of $\mathbb{R}^d$.

By C^k convergence of $f_n^* g_n$ to g on compact subsets of Ω we mean that for any smooth parametrization $h : U \to V \subset \Omega$ the coefficients of the metrics $h^* f_n^* g_n$ converge to those of $h^* g$ in the C^k topology on every compact subset of U.

To see that the restriction to compact subsets of U is necessary consider the sequence of Riemannian metrics g_n on the open interval $(0, 1)$ defined by

$$g_n(x)(v, w) = e^{x/n} vw.$$

The sequence of coefficients $x \mapsto e^{x/n}$ in this example converges uniformly to the coefficient of the metric g on $(0, 1)$ given by

$$g(x)(v, w) = vw$$

however taking pullback under the diffeomorphism $h : (0, 1) \to (0, 1)$ defined by $h(x) = x^\alpha$ one obtains

$$h^* g_n(x)(v, w) = e^{x^\alpha/n} \alpha^2 x^{2(\alpha-1)} vw$$

so that taking for example $\alpha = 1/2$ one sees that uniform convergence of the sequence of coefficients no longer holds.

We now present a coordinate free definition of C^k convergence.

For this purpose we recall that a (p, q) tensor on a vector space V is an element of $(V^*)^{\otimes q} \otimes V^{\otimes p}$. If g is an inner product on V then g induces an inner product and norm on the space of (p, q) tensors. This inner product can be defined by taking any g-orthonormal basis $v_1, \ldots, v_d$ of V, considering the dual basis $v^1, \ldots, v^d$, and declaring that the tensors of the form $v^{i_1} \otimes \cdots v^{i_q} \otimes v_{i_{1+q}} \otimes \cdots \otimes v_{i_{p+q}}$ are orthonormal.

In particular given a Riemannian manifold (M, g) and a (p, q) tensor field T one can consider the tensor norm $|T(x)|_g$ of the tensor $T(x)$ over the tangent space $T_x M$ with respect to the inner product $g(x)$.

Lemma 4.23 (Characterization of convergence). *A sequence (M_n, o_n, g_n) of pointed connected complete Riemannian manifolds converges C^k to (M, o, g) if and only if for each*

$r > 0$ there exists a sequence of pointed embeddings (defined for n large enough) $f_n : B_r(o) \to M_n$ such that

$$\lim_{n\to+\infty} \sup\{|\nabla^i(f_n^* g_n - g)(x)|_g : x \in B_r(o), i = 0, \dots, k\} \to 0$$

where ∇ denotes the covariant derivative corresponding to the Riemannian metric g (in particular for $i \neq 0$ one has $\nabla^i g = 0$).

Proof. Assume first that a sequence (M_n, o_n, g_n) in $\mathcal{M}$ converges C^k to (M, o, g) and fix $r > 0$.

By definition of C^k convergence there exists a domain $\Omega \supset B_{2r}(o)$ sequence of pointed embeddings $f_n : \Omega \to M_n$ such that $f_n^* g_n$ converges C^k on compact sets of Ω to g. This means that in any local chart the coefficients of $f_n^* g_n$ will converge C^k on compact sets to those of g. By Lemma 4.25 below this implies $|\nabla^i(f_n^* g_n - g)| \to 0$ uniformly on compact subset of $B_{2r}(o)$ for $i = 0, \dots, k$. In particular since $\overline{B_r}(o)$ is compact one has

$$\lim_{n\to+\infty} \sup\{|\nabla^i g_n(x) - \nabla^i g(x)|_g : x \in B_r(o), i = 0, \dots, k\} \to 0$$

as claimed.

We will now prove the converse claim.

Given r we must obtain a sequence of embeddings f_n of an open set $\Omega \supset B_r(o)$ into M_n such that $f_n(\Omega) \supset B_r(o_n)$ and $f_n^* g_n$ converges C^k to g on compact sets. We will show that one can take $\Omega = B_{2r}(o)$.

By hypothesis there exists a sequence of pointed embeddings $f_n : B_{2r}(o) \to M_n$ such that $|\nabla^i(f_n^* g_n - g)| \to 0$ uniformly for $i = 0, \dots, k$. By Lemma 4.25 below this implies that $f_n^* g_n$ converges C^k to g on compact subsets of $B_{2r}(o)$.

We must now establish that $f_n(B_{2r}(o)) \supset B_r(o_n)$ for all n large enough.

To see this let $v_1, \dots, v_d$ be a g-orthonormal basis of the tangent space $T_x M$ at a point $x \in B_{2r}(o)$ and $v^1, \dots, v^d$ the dual basis. One has $f_n^* g_n(x) = a_{ij} v^i \otimes v^j$ and

$$|(f_n^* g_n - g)(x)|_g^2 = \sum_{i,j} (a_{ij} - \delta_{ij})^2$$

where (δ_{ij}) is the identity matrix.

For all n large enough the left-hand side above will be small enough to guarantee that $a_{11} = f_n^* g_n(x)(v_1, v_1) = |v_1|^2_{f_n^* g_n} > 1/4$. And, since one can choose any g-orthonormal basis to calculate the norm above, this implies

$$\frac{1}{2}|v|_g < |v|_{f_n^* g_n}$$

for all $v \in T_x M$ and all $x \in B_{2r}(o)$.

In particular for n large enough the $f_n^* g_n$ length of any curve joining o and the boundary of $B_{2r}(o)$ will be at least r. So that $f_n(B_{2r}) \supset B_r(o_n)$ as claimed. $\square$

The following consequence was used in the proof of Theorem 4.11.

Lemma 4.24. *On any subset of $\mathcal{GS}$ of the form $\mathcal{M} = \mathcal{M}(d, r, \{C_k\})$ smooth convergence is topologizable.*

Proof. We define the k-th order (r,ϵ)-neighborhood of a manifold $M \in \mathcal{M}$ as the set of $N \in \mathcal{M}$ such that there exists a pointed embedding $f : B_r(o_M) \to N$ satisfying

$$\sup\left\{|\nabla^i(g_M - f^*g_N)(x)|_g : x \in B_r(o_M), i = 0,\dots,k\right\} < \epsilon.$$

By Lemma 4.23 convergence with respect to the topology on $\mathcal{M}$ generated by all k-th order (r,ϵ)-neighborhoods (for all $k \in \mathbb{N}, r > 0$ and $\epsilon > 0$) coincides with smooth convergence. $\square$

4.7.2 Convergence of tensor fields

We will now complete the calculations in local coordinates needed for the proof of Lemma 4.23.

Recall that the coefficients of a (p,q) tensor field T on an open set of $\mathbb{R}^d$ are the functions

$$x \mapsto T(x)(e_{i_1},\dots,e_{i_q},e^{i_{q+1}},\dots,e^{i_{p+q}}).$$

In what follows we use $|T|$ to denote the Euclidean tensor norm and $|T|_g$ to denote the tensor norm coming from a metric g.

The following result characterizes C^k convergence of the coefficients of such a tensor on a compact set in a coordinate invariant manner.

Lemma 4.25 (Convergence of tensor fields)**.** *Let U be an open subset of $\mathbb{R}^d$, g a Riemannian metric on U, K a compact subset of U, and T_n a sequence of (p,q)-tensor fields on U. Then the following two statements are equivalent for all $k \geq 0$:*

1. *The coefficients of T_n and their partial derivatives up to order k converge to 0 uniformly on K.*
2. *For each $i = 0,1,\dots,k$ one has*

$$\lim_{n\to+\infty} \max\left\{|\nabla^i T_n(x)|_g : x \in K\right\} = 0.$$

Proof. Since K is compact and the metric coefficients and Christoffel symbols are smooth there exist constants $C \geq 1$ and $\Gamma > 0$ such that

1. The absolute value of the derivatives of the Christoffel symbols up to order k are bounded by Γ on K.
2. For any tensor field T of type (p,q') with $q \leq q' \leq q+k$ one has

$$C^{-1}|T(x)| \leq |T(x)|_g \leq C|T(x)|$$

for all $x \in K$.

Notice that, by the existence of the constant C above, if T_n is a sequence of (p,q')-tensor fields with $q \leq q' \leq q+k$ then $|T_n(x)|$ converges to 0 uniformly on K if and only if $|T_n(x)|_g$ does. On the other hand $|T_n(x)|$ is the square root of the sum of squares of the coefficients of T_n which implies that both the previous statements are equivalent to the uniform convergence of the coefficients to 0 on K.

In particular, this establishes the case $k = 0$ of the lemma. We will prove the lemma by induction on k but first we must establish some basic properties of the coefficients of $\nabla^i T_n$.

For this purpose, assuming that T is a (p, q')-tensor field, observe that the coefficients of ∇T are obtained from the equation

$$\nabla T(Y, X_1, \ldots, X_{p+q'}) = \nabla_Y T(X_1, \ldots, X_{p+q'}) - \sum_{i=1}^{p+q'} T(X_1, \ldots, \nabla_Y X_i, \ldots, X_{p+q'})$$

by substituting elements of the canonical basis for $Y, X_1, \ldots, X_{q'}$ and elements of the dual basis for $X_{q'+1}, \ldots, X_{p+q'}$.

The first term above is simply the derivative in the direction of the basis vector Y of a coefficient of T while the other terms are products of the coefficients of T with Christoffel symbols.

By induction one can establish that for each i one has

1. Each coefficient of $\nabla^i T$ is the sum of one i-th order partial derivative of a coefficient of T with products of lower order partial derivatives coefficients of T with partial derivatives of the Christoffel symbols of order less than or equal to i.

2. Every partial derivative of order i of each coefficient of T appears in at least one of the aforementioned sums.

Now assume that our lemma is true for $k - 1$.

If $|\nabla^i T_n|_g$ converges to 0 uniformly on K for each $i \leq k$ then by the induction hypothesis the partial derivatives of the coefficients of T_n up to order $k - 1$ converge uniformly to 0 on K. Using the properties of $\nabla^k T_n$ established above and the bound Γ on partial derivatives of the Christoffel symbols it follows that the k-th order partial derivatives of the coefficients of T_n converge to 0 uniformly on K as well.

Similarly if the partial derivatives up to order k of the coefficients of T_n converge to 0 uniformly on K then by the properties of $\nabla^i T_n$ established above and the bounds on the Christoffel factors one obtains that the coefficients of $\nabla^i T_n$ converge to 0 uniformly for each $i \leq k$ on K. This implies (using the constant C defined above) the claim on $|\nabla^i T_n|_g$. □

4.8 Bounded geometry of leaves

We now verify that the leaves of a compact foliation have uniformly bounded geometry. This was used in the proof of Theorem 4.1.

Lemma 4.26. *If X is a compact d-dimensional foliation then there exists $r > 0$ and a sequence $\{C_k : k \geq 0\}$ such that all leaves belong to the space $\mathcal{M}(d, r, \{C_k\})$.*

Proof. We have shown in Section 4.6 that the norm of the k-th covariant derivative of the curvature tensor is a continuous function of the metric coefficients, the coefficients of the inverse matrix, and a finite number of their partial derivatives. This implies (by looking at the leaf metrics in a foliated chart) that this norm is continuous on X and hence has a global maximum C_k.

Let $h : \mathbb{R}^d \times T \to U \subset X$ be a foliated parametrization and for each $t \in T$ let g_t be the Riemannian metric on $\mathbb{R}^d$ obtained by pullback under $x \mapsto h(x, t)$.

Let $\exp_{x,t}$ denote the exponential map of the metric g_t at x (i.e. $\exp_{x,t} : \mathbb{R}^d \to \mathbb{R}^d$ with $\exp_t(v) = \alpha(1)$ where α is the g_t-geodesic satisfying $\alpha(0) = x$ and $\alpha'(0) = v$).

By continuity of the solution to an ordinary differential equation with respect to the vector field (see [DK00, Theorem B3, pg. 333]) one has that $(x, t) \mapsto \exp_{x,t}$ is continuous when the codomain is endowed with the topology of C^k convergence on compact subsets of $\mathbb{R}^d$.

In particular each $(x,t) \in \mathbb{R}^d \times T$ has a neighborhood U on which there is a radius $r > 0$ such that the operator norm of the difference between the differential of $\exp_{y,s}$ and the identity is less than $\frac{1}{2}$ at all points in $B_{2r}(0)$ for all $(y,s) \in U$. This implies that $\exp_{y,s} - z$ is a contraction mapping $B_{2r}(0)$ into itself for all $z \in B_r(y)$. In particular the injectivity radius of g_s at y is at least r.

It follows that each point $x \in X$ has a neighborhood on which there is a uniform positive lower bound for the leafwise injectivity radius at each point. Covering X by a finite number of these open sets one obtains that there is a global positive lower bound for the injectivity radius of all leaves. □

4.9 Covering spaces and holonomy

In this section we recall some basic facts on Riemannian coverings and provide the definitions and results on holonomy which are relevant for Theorem 4.3.

4.9.1 Riemannian coverings

By a Riemannian covering of a pointed complete connected Riemannian manifold (M, o, g) we mean a pointed local isometry $f : N \to M$ from some pointed complete connected Riemannian manifold (N, o_N, g_N) to M. We sometimes omit the function f and simply say that N is a Riemannian covering of M (meaning there exists at least one suitable f).

Any covering space N in the sense of [Hat02, Chapter 1.3] can be made a Riemannian covering by constructing local charts which are compositions of the covering map and the local charts of the covered manifold M. Reciprocally any Riemannian covering satisfies the 'pile of disks' property for the preimage of balls of radius smaller than half the injectivity radius of M.

We recall that the fundamental group $\pi_1(M)$ of (M, o, g) is the group of (endpoint fixing) homotopy classes of closed curves starting and ending at the basepoint o. Any covering $f : N \to M$ induces a morphism f_* from $\pi_1(N)$ to $\pi_1(M)$.

With these observations we restate [Hat02, Theorem 1.38] and the comment immediately after about ordering covering spaces as we shall use them.

Lemma 4.27 (Classification of covering spaces)**.** *Let (M, o, g) be a pointed complete Riemannian manifold. For each subgroup H of $\pi_1(M)$ there exists a Riemannian covering $f : N \to M$ with $f_*(\pi_1(N)) = H$ and this covering space is unique up to pointed isometries. If two Riemannian coverings N and N' correspond to subgroups $H \subset H'$ then N is a Riemannian covering of N'.*

The Riemannian covering associated to the trivial subgroup of $\pi_1(M)$ is the universal covering which we denote by $\widetilde{M}$.

4.9.2 Holonomy covering

Let X be a foliation and $h_i : U_i \to \mathbb{R}^d \times T_i$ (where $i = 1, 2$) be foliated charts.

The charts h_1, h_2 are said to be compatible if there exists a homeomorphism ψ : $h_1(U_1 \cap U_2) \to h_2(U_1 \cap U_2)$ such that

$$h_2 \circ h_1^{-1}(x,t) = (\varphi(x,t), \psi(t))$$

for a certain (automatically smooth with respect to x) function φ and all (x, t) with $h_1(x, t) \in U_1 \cap U_2$. The map ψ is called the holonomy from h_1 to h_2.

Notice that the Vinyl record foliation of Section 4.2.1 cannot be covered by only two compatible charts.

By a chain of compatible foliated charts we mean a finite sequence $h_i : U_i \to \mathbb{R}^d \times T_i$ indexed on $i = 0, 1, \ldots, r$ of foliated charts such that U_i intersects U_{i+1} and h_i is compatible with h_{i+1} for all $i = 0, \ldots, r-1$. The chain is closed if $h_r = h_0$. The holonomy of the chain is the map $\psi_{r-1,r} \circ \cdots \circ \psi_{0,1}$ where $\psi_{i,i+1}$ is the holonomy from h_i to h_{i+1} and we assume the maximal possible domain for the composition.

A leafwise curve is a continuous function $\alpha : [0, 1] \to X$ whose image is contained in a single leaf. We say α is covered by a compatible chain of foliated charts $\{h_i, i = 0, \ldots, r\}$ if there exists a finite sequence $t_0 = 0 < \cdots < t_r = 1$ such that $\alpha([t_i, t_{i+1}]) \subset U_i$ for all $i = 0, \ldots, r-1$ where U_i is the domain of h_i.

A closed leafwise curve $\alpha : [0, 1] \to X$ is said to have trivial holonomy if there exists a closed chain of compatible charts $\{h_i : i = 0, \ldots, r\}$ covering α whose holonomy map is the identity on a neighborhood of $t \in T$ where t is the second coordinate of $h_0(\alpha(0))$.

The holonomy covering $\widetilde{L_x}^{\text{hol}}$ of a leaf L_x is defined as the Riemannian covering corresponding (via Lemma 4.27) to the subgroup H of homotopy classes of closed curves based at x in L_x which have trivial holonomy.

To show that this is well defined we must prove that:

1. Each closed leafwise curve admits a covering by a compatible closed chain of foliated charts.
2. The property of having trivial holonomy does not depend on the choice of covering.
3. The property of having trivial holonomy is invariant under homotopy.

A leaf L_x is said to have trivial holonomy if L_x is isometric to its holonomy cover (equivalently all leafwise closed curves based at x have trivial holonomy). The fact that this property does not depend on the basepoint x follows from Lemma 4.32 below, this lemma also covers item 3 in the above list and shows that the holonomy cover is a normal covering space (although we will not use this fact).

4.9.3 Trivial holonomy

In this subsection we verify the claims necessary for the definition of the holonomy covering of a leaf. We also provide a characterization of trivial holonomy which was used in the proof of Theorem 4.3.

Recall that an atlas of a foliation is simply a collection of foliated charts whose domains cover the foliation. We say one atlas refines another if the domain of each chart of the former is contained in the domain of some chart of the later. We call an atlas consisting of pairwise compatible charts 'admissible'.

Lemma 4.28. *Every atlas of a compact foliation has an admissible refinement.*

Proof. Let A be an atlas of a compact foliation X. Since X is compact we may take a finite subatlas B of A.

Let $h : U \to \mathbb{R}^d \times T$ be a chart in B. Given an open ball $D \subset \mathbb{R}^d$ and an open subset $S \subset T$ we may construct a chart $g : h^{-1}(D \times S) \to \mathbb{R}^d \times S$ by letting $g(x) = f(h(x))$ where f acts as the identity on the second coordinate and a fixed diffeomorphism between

D and $\mathbb{R}^d$ on the second. We call any such chart a restriction of h and note that any two restrictions of the same chart are compatible.

Now let $h_i : U_i \to \mathbb{R}^d \times T_i$ be charts in B for $i = 1, 2$. Even if these charts are not compatible the fact that they are foliated charts implies that each point x in $U_1 \cap U_2$ has a neighborhood $V = h_1^{-1}(D \times S)$ where $D \subset \mathbb{R}^d$ is an Euclidean open ball and S is an open subset of T_1, such that

$$h_2 \circ h_1^{-1}(y) = (\varphi(y, t), \psi(t))$$

on $h_1(V)$, where ψ is a homeomorphism between certain open sets in T_1 and T_2. This implies that restricting h_1 to V one obtains a chart which is compatible with any restriction of h_1 or h_2.

Since there are only finitely many charts we may choose for each point x a neighborhood, and a restriction of a certain chart in B to this neighborhood which will be compatible with (the restrictions of) all charts in B. The collection of such charts gives a compatible refinement C of A. □

From the above result it follows that any closed curve has a covering by a compatible chain of foliated charts.

We now establish the fact that having trivial holonomy does not depend on the choice of covering. We recall that the plaques of a foliated chart $h : U \to \mathbb{R}^d \times T$ are the sets of the form $h^{-1}(\mathbb{R}^d \times \{t\})$.

Lemma 4.29. *If α is a leafwise closed curve in a compact foliation X. Then α has trivial holonomy if and only if for each sequence α_n of (possibly non-closed) leafwise curves which converge uniformly to α and any foliated chart $h : U \to \mathbb{R}^d \times T$ where $\alpha(0) \in U$ one has that $\alpha_n(0)$ and $\alpha_n(1)$ belong to the same plaque for n large enough.*

Proof. Observe that if $h_i : U_i \to \mathbb{R}^d \times T_i$ (where $i = 1, 2$) are compatible foliated charts and β is a leafwise curve whose image is contained in $U_1 \cup U_2$ then there exists $t_i \in T_i$ $(i = 1, 2)$ such that β is in the plaque $h_i^{-1}(\mathbb{R}^d \times \{t_i\})$ whenever it is in U_i. Furthermore $t_2 = \psi(t_1)$ where ψ is the holonomy between the charts.

By definition α is covered by a closed chain of compatible charts $h_i : U_i \to \mathbb{R}^d \times T_i$ where $i = 0, \ldots, r$ and there exist $t_0 = 0 < \cdots < t_r = 1$ such that $\alpha([t_i, t_{i+1}]) \subset U_i$ for all $i = 0, \ldots, r-1$. Furthermore the holonomy ψ of the chain is the identity on a neighborhood of the second coordinate of $h_0(\alpha(0))$.

Take $\epsilon > 0$ such that

1. Any leafwise curve $\beta : [0, 1] \to X$ at distance less than ϵ from α one has $\beta([t_i, t_{i+1}]) \subset U_i$ for all $i = 0, \ldots, r-1$.

2. Any point p at distance less than ϵ from $\alpha(0)$ is of the form $h_0^{-1}(x, t)$ where $\psi(t) = t$.

It follows that any leafwise curve at uniform distance less than ϵ from α will start and end in the same plaque. □

The first observation in the above proof yields the following.

Corollary 4.30. *If α is a leafwise closed curve with trivial holonomy in a compact foliation X. Then any closed chain of compatible charts $h_i : U_i \to \mathbb{R}^d \times T_i$ (where $i = 0, \ldots, r$) which covers α has trivial holonomy in a neighborhood of the second coordinate of $h_0(\alpha(0))$.*

By the leafwise distance between to points on the same leaf L_x we mean the distance with respect to the Riemannian metric g_{L_x}.

The following corollary amounts to the observation that if x_n, y_n are two sequences of points converging to the same limit x which belong to the same sequence of plaques with respect to some chart h covering x, then the leafwise distance between x_n and y_n converges to 0.

Corollary 4.31. *If a sequence of leafwise curves $\alpha_n : [0,1] \to X$ converges uniformly to a closed leafwise curve α and the leafwise distance between $\alpha_n(0)$ and $\alpha_n(1)$ does not converge to* 0*, then α has non-trivial holonomy.*

We say two leafwise curves $\alpha, \beta : [0,1] \to X$ are leafwise freely homotopic if they belong to the same leaf L and there exists a continuous function $h : [0,1] \times [0,1] \to X$ such that $t \mapsto h(s,t) = h_s(t)$ is a leafwise closed curve in L for all s, $h_0 = \alpha$ and $h_1 = \beta$.

Lemma 4.32. *Let X be a compact foliation and $\alpha : [0,1] \to X$ a closed leafwise curve with trivial holonomy. Then any closed leafwise curve which is leafwise freely homotopic to α also has trivial holonomy.*

Proof. Take an admissible finite atlas A of X and let $\epsilon > 0$ be the Lebesgue number of the associated open covering. It follows from Corollary 4.30 that if two closed curves belong to the same leaf and are at uniform distance less than ϵ and one of them has trivial holonomy then they both do.

Letting h_s be a homotopy between α and β one can find times $s_0 = 0 < \ldots < s_r = 1$ such that h_{s_i} is at uniform distance less than ϵ from $h_{s_{i+1}}$ for all $i = 0, \ldots, r-1$ from which the lemma follows. □

4.10 Convergence of leafwise functions

In this section we justify the claims on convergence of immersions into foliations which were used in the proof of Theorem 4.3.

4.10.1 Adapted distances

Let X be a compact foliation. Denote by d_L the leafwise distance in X which is defined by

$$d_L(x,y) = \begin{cases} d_{L_x}(x,y) & \text{if } y \in L_x. \\ +\infty & \text{otherwise.} \end{cases}$$

where d_{L_x} is the Riemannian distance on the leaf L_x.

We call a distance d on X adapted if it metricizes the topology of X and satisfies $d(x,y) \le d_L(x,y)$.

Consider a smooth Riemannian manifold X foliated by smoothly immersed leaves each of which inherits the ambient Riemannian metric (e.g. any example in Section 4.2). The Riemannian distance between two points $x, y \in X$ is the infimum of the lengths of arbitrary curves connecting them while the leafwise distance is the infimum among leafwise curves. Hence one clearly has that the Riemannian distance is adapted to the foliation. This makes the following result plausible.

Lemma 4.33. *Every compact foliation has an adapted distance.*

Proof. Let X be a compact foliation. We will construct an adapted distance by averaging pseudodistances obtained by a local construction.

Let $h : V \to \mathbb{R}^d \times T$ be a foliated chart and let g_t be the family of Riemannian metrics on $\mathbb{R}^d$ parametrized by $t \in T$ obtained by pushforward of the leafwise metrics under h. Fix a complete distance d_T on T and metrizice $\mathbb{R}^d \times T$ by defining

$$\rho_1((x,t),(x',t')) = |x - x'| + d_T(t,t')$$

for all $(x,t),(x',t') \in \mathbb{R}^d \times T$.

We observe that because X is compact any point in $\mathbb{R}^d \times T$ has a compact neighborhood, and it follows that T is locally compact.

Fix $(x,t) \in \mathbb{R}^d \times T$ and consider a precompact open neighborhood $S \subset T$ of t. The family of Riemannian metrics on the Euclidean ball $B_1(x)$ of the form g_s for $s \in S$ is smoothly precompact. Hence there exists a constant $C > 0$ such that the g_s-length of any curve between points $y, y' \in \overline{B_1}(x)$ is at least $C|y - y'|$ for all $s \in S$.

For this constant C we choose a continuous function $\varphi : \mathbb{R}^d \times T \to [0, C]$ which is strictly positive exactly on the set $B_1(x) \times S$ and zero outside of it and define for all $(y,s),(y',s') \in \mathbb{R}^d \times T$

$$\rho_2((y,s),(y',s')) = \inf\left\{\sum_{i=0}^{k-1} \frac{\varphi(y_i,s_i) + \varphi(y_{i+1},s_{i+1})}{2} \cdot \rho_1((y_i,s_i),(y_{i+1},s_{i+1}))\right\}$$

where the infimum is among all $k \in \mathbb{N}$ and all finite chains in $\mathbb{R}^d \times T$ with $(y_0,s_0) = (y,s)$ and $(y_k,s_k) = (y',s')$.

Because one can reverse a chain and concatenate two of them one obtains that ρ_2 is symmetric and satisfies the triangle inequality. Notice also that ρ_2 is zero for any pair of points not in $B_1(x) \times S$.

Now consider $(y,s) \in B_1(x) \times S$ and the function $f(y',s') = \rho_2((y,s),(y',s'))$. By the triangle inequality f is constant outside of $B_1(x) \times S$. Given $(y',s') \neq (y,s)$ one may choose $r > 0$ such that the ρ_1-ball $B_{\rho_1,r}(y,s)$ of radius r centered at (y,s) does not contain (y',s') and the values of φ on this ball are bounded from below by a positive constant ϵ. For any finite chain (y_i,s_i) joining (y,s) and (y',s') one may take the first k such that $(y_k,s_k) \notin B_{\rho_1,r}(y,s)$ and since ρ_1 is a distance one has

$$\sum_{i=0}^{k-1} \frac{\varphi(y_i,s_i) + \varphi(y_{i+1},s_{i+1})}{2} \cdot \rho_1((y_i,s_i),(y_{i+1},s_{i+1})) \geq \frac{1}{2}\epsilon r.$$

Hence f is zero only at (y,s). Combined with the inequality $\rho_2 \leq C\rho_1$ one obtains that ρ_2 is a continuous bounded pseudodistance on $\mathbb{R}^d \times T$ which is an actual distance when restricted to $B_1(x) \times S$ and which is zero on pairs of points not belonging to $B_1(x) \times S$.

Hence the pullback of ρ_2 to V via h can be extended to a bounded continuous pseudodistance $\rho : X \times X \to [0, +\infty)$ which is an actual distance when restricted to the open set $U = h^{-1}(B_1(x) \times S)$ and which is zero on pairs of points not belonging to this set.

We will now establish that $\rho(p,p') \leq d_L(p,p')$ whenever p and p' are on the same leaf. The only interesting case (i.e. $\rho \neq 0$) is if either p or p' belong to U. Suppose $p \in U$ and let $\alpha : [0,1] \to X$ be the leafwise geodesic of length $d_L(p,p')$ joining p and p'. There are two cases to consider: either α leaves U or it does not.

If $\alpha([0,1]) \subset U$ then taking $\beta = h \circ \alpha$ and setting $(y,s) = \beta(0)$ and $(y',s') = \beta(1)$ one obtains that $s = s'$. Since $\rho(p,p') = \rho_2((y,s),(y',s)) \leq C|y - y'|$ which is a lower bound for the g_s length of any curve joining y and y' one obtains that $\rho(p,p') \le d_L(p,p')$ as claimed.

Now suppose that α leaves U and take $T \in [0,1)$ so that $\beta = h \circ \alpha$ is well defined on $[0,T]$ and $\beta(T) \notin B_1(x) \times S$. Setting $(y,s) = \beta(0)$ and $(y',s') = \beta(T)$ one obtains once again that $s = s'$ and that the g_s-length of β is at least $C|y-y'|$ which is larger than $\rho(\alpha(0),\alpha(T)) = \rho_2((y,s),(y',s'))$. If $p' \notin U$ we are done since $\rho(\alpha(T),p') = 0$. Otherwise we take $T < T_2 < 1$ so that $\beta_2 = h \circ \alpha$ is well defined on $[T_2,1]$ and repeat the preceeding argument to obtain that $\rho(p,p') \le \rho(p,\alpha(T)) + \rho(\alpha(T_2),p')$ is less than the length of α as claimed.

We have succeeded in constructing for each $p \in X$ an open neighborhood U and a continuous bounded pseudodistance ρ which is an actual distance when restricted to $U \times U$ and which satisfies $\rho(q,q') \le d_L(q,q')$. Covering X with a finite number of such neighborhoods $U_1,\ldots,U_n$ with associated pseudodistances $\rho_1,\ldots,\rho_n$ and setting $d(p,q) = \frac{1}{n}\sum\limits_{i=1}^{n} \rho_i(p,q)$ one obtains an adapted distance for the foliation X. □

4.10.2 Convergence of leafwise immersions

We conclude the section with the following result which was used to proved Theorem 4.3 (recall that a function into a foliation is said to be leafwise if its image is contained in a single leaf).

Lemma 4.34. *Let X be a compact foliation and (M,o,g) be a complete pointed Riemannian manifold. If $f_n : B_r(o) \to X$ is a sequence of leafwise functions such that $|f_n^* g_{L_{f_n(o)}} - g|_g$ converges to 0 uniformly then there exists a subsequence converging locally uniformly to a leafwise local isometry $f : B_r(o) \to X$.*

Proof. The hypothesis implies that the f_n are locally uniformly Lipschitz with respect to any adapted distance on X. By the Arzelà-Ascoli Theorem there exists a subsequence f_{n_k} which converges locally uniformly to a limit f.

Given $x \in B_r(o)$ we may consider a foliated parametrization $h : \mathbb{R}^d \times T \to U \subset X$ of a neighborhood of $f(x)$ such that $f(x) = h(0,t)$. For each $s \in T$ let g_s be the Riemannian metric on $\mathbb{R}^d$ obtained by pullback under $z \mapsto h(z,s)$.

Let $\epsilon > 0$ be such that the g_t-ball centered at 0 of radius 2ϵ is bounded with respect to the Euclidean metric on $\mathbb{R}^d$, the ball of radius 2ϵ centered at x is contained in $B_r(o)$, and $f(B_{2\epsilon}(x))$ is contained in the open set parametrized by h.

We will show that $\pi_1 \circ h^{-1} \circ f$ is an isometry from $B_\epsilon(x)$ into $\mathbb{R}^d$ with the Riemannian metric g_t where $\pi_1 : \mathbb{R}^d \times T \to \mathbb{R}^d$ is the projection onto the first coordinate.

For this purpose consider $y \in B_\epsilon(x)$. The sequences $p_n = \pi_1 \circ h \circ f_n(x)$ and $q_n = \pi_1 \circ h \circ f_n(y)$ are eventually well defined and converge to $p = \pi_1 \circ h \circ f(x)$ and $q = \pi_1 \circ h \circ f(y)$ respectively. Furthermore letting t_n be the common coordinate in T of $h \circ f_n(x)$ and $h \circ f_n(y)$ one has that the g_{t_n}-distance between p_n and q_n converges to $d_M(x,y)$ (d_M being the Riemannian distance on M) [2]. Since g_{t_n} converges smoothly on compact sets to g_t one has that the g_t-distance between p and q equals $d_M(x,y)$ as claimed. □

2. Here we use the fact that, if n is large enough, the g_{t_n}-ball of radius ϵ is bounded in $\mathbb{R}^d$.

Bibliography

[ÁC03] J. A. Álvarez López and A. Candel. Generic aspects of the geometry of leaves of foliations. `http://www.csun.edu/~ac53971/research/kyoto_GF2003.pdf`, 2003.

[ADT07] Marc Arnaudon, Bruce K. Driver, and Anton Thalmaier. Gradient estimates for positive harmonic functions by stochastic analysis. *Stochastic Process. Appl.*, 117(2):202–220, 2007.

[Anc90] A. Ancona. Théorie du potentiel sur les graphes et les variétés. In *École d'été de Probabilités de Saint-Flour XVIII—1988*, volume 1427 of *Lecture Notes in Math.*, pages 1–112. Springer, Berlin, 1990.

[AS85] Michael T. Anderson and Richard Schoen. Positive harmonic functions on complete manifolds of negative curvature. *Ann. of Math. (2)*, 121(3):429–461, 1985.

[AT11] Marc Arnaudon and Anton Thalmaier. Brownian motion and negative curvature. In *Random walks, boundaries and spectra*, volume 64 of *Progr. Probab.*, pages 143–161. Birkhäuser/Springer Basel AG, Basel, 2011.

[Ave74] André Avez. Théorème de Choquet-Deny pour les groupes à croissance non exponentielle. *C. R. Acad. Sci. Paris Sér. A*, 279:25–28, 1974.

[Ave76] A. Avez. Harmonic functions on groups. In *Differential geometry and relativity*, pages 27–32. Mathematical Phys. and Appl. Math., Vol. 3. Reidel, Dordrecht, 1976.

[BBI01] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.

[BC12] Itai Benjamini and Nicolas Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17:no. 93, 1–20, 2012.

[Bil99] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.

[Bre10] Leo Brewin. A brief introduction to cadabra: A tool for tensor computations in general relativity. *Computer Physics Communications*, 181(3):489 – 498, 2010.

[BSW12] Sara Brofferio, Maura Salvatori, and Wolfgang Woess. Brownian motion and harmonic functions on $\mathrm{Sol}(p, q)$. *Int. Math. Res. Not. IMRN*, (22):5182–5218, 2012.

[Can03] Alberto Candel. The harmonic measures of Lucy Garnett. *Adv. Math.*, 176(2):187–247, 2003.

[Cha84] Isaac Chavel. *Eigenvalues in Riemannian geometry*, volume 115 of *Pure and Applied Mathematics*. Academic Press Inc., Orlando, FL, 1984. Including a chapter by Burton Randol, With an appendix by Jozef Dodziuk.

[Che70] Jeff Cheeger. Finiteness theorems for Riemannian manifolds. *Amer. J. Math.*, 92:61–74, 1970.

[CLN85] César Camacho and Alcides Lins Neto. *Geometric theory of foliations*. Birkhäuser Boston Inc., Boston, MA, 1985. Translated from the Portuguese by Sue E. Goodman.

[Cri08] Jan Cristina. Gromov-hausdorff convergence of metric spaces. `http://www.helsinki.fi/~cristina/pdfs/gromovHausdorff.pdf`, 2008.

[Der76] Yves Derriennic. Lois "zéro ou deux" pour les processus de Markov. Applications aux marches aléatoires. *Ann. Inst. H. Poincaré Sect. B (N.S.)*, 12(2):111–129, 1976.

[Der85] Y. Derriennic. *Entropie, théorèmes limite et marches aléatoires*. Publications de l'Institut de Recherche Mathématique de Rennes. [Publications of the Rennes Institute of Mathematical Research]. Université de Rennes I Institut de Recherche Mathématique de Rennes, Rennes, 1985.

[DGM77] A. Debiard, B. Gaveau, and E. Mazet. Théorèmes de comparaison en géométrie riemannienne. *Publ. Res. Inst. Math. Sci.*, 12(2):391–425, 1976/77.

[dJ77] Andrés del Junco. On the decomposition of a subadditive stochastic process. *Ann. Probability*, 5(2):298–302, 1977.

[DK00] J. J. Duistermaat and J. A. C. Kolk. *Lie groups*. Universitext. Springer-Verlag, Berlin, 2000.

[DK07] Bertrand Deroin and Victor Kleptsyn. Random conformal dynamical systems. *Geom. Funct. Anal.*, 17(4):1043–1105, 2007.

[dlR93] Thierry de la Rue. Espaces de Lebesgue. In *Séminaire de Probabilités, XXVII*, volume 1557 of *Lecture Notes in Math.*, pages 15–21. Springer, Berlin, 1993.

[DM88] E. B. Davies and N. Mandouvalos. Heat kernel bounds on hyperbolic space and Kleinian groups. *Proc. London Math. Soc. (3)*, 57(1):182–208, 1988.

[Doo01] Joseph L. Doob. *Classical potential theory and its probabilistic counterpart*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1984 edition.

[Ehr74] Paul E. Ehrlich. Continuity properties of the injectivity radius function. *Compositio Math.*, 29:151–178, 1974.

[Eic91] Jürgen Eichhorn. The boundedness of connection coefficients and their derivatives. *Math. Nachr.*, 152:145–158, 1991.

[EMS77] Robert Edwards, Kenneth Millett, and Dennis Sullivan. Foliations with all leaves compact. *Topology*, 16(1):13–32, 1977.

[EMT77] D. B. A. Epstein, K. C. Millett, and D. Tischler. Leaves without holonomy. *J. London Math. Soc. (2)*, 16(3):548–552, 1977.

[Eps72] D. B. A. Epstein. Periodic flows on three-manifolds. *Ann. of Math. (2)*, 95:66–82, 1972.

[Eps76] D. B. A. Epstein. Foliations with all leaves compact. *Ann. Inst. Fourier (Grenoble)*, 26(1):viii, 265–282, 1976.

[EV78] D. B. A. Epstein and E. Vogt. A counterexample to the periodic orbit conjecture in codimension 3. *Ann. of Math. (2)*, 108(3):539–552, 1978.

[EW11] Manfred Einsiedler and Thomas Ward. *Ergodic theory with a view towards number theory*, volume 259 of *Graduate Texts in Mathematics*. Springer-Verlag London Ltd., London, 2011.

[Fur63] Harry Furstenberg. Noncommuting random products. *Trans. Amer. Math. Soc.*, 108:377–428, 1963.

[FV10] Peter K. Friz and Nicolas B. Victoir. *Multidimensional stochastic processes as rough paths*, volume 120 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010. Theory and applications.

[Gar81] Lucy Jane Garnett. *Functions and measures harmonic along the leaves of a foliation and the ergodic theorem*. ProQuest LLC, Ann Arbor, MI, 1981. Thesis (Ph.D.)–Dartmouth College.

[Gar83] Lucy Garnett. Foliations, the ergodic theorem and Brownian motion. *J. Funct. Anal.*, 51(3):285–311, 1983.

[Ghy95] Étienne Ghys. Topologie des feuilles génériques. *Ann. of Math. (2)*, 141(2):387–422, 1995.

[GMM12] S. Gouëzel, F. Mathéus, and F. Maucourant. Sharp lower bounds for the asymptotic entropy of symmetric random walks. *ArXiv e-prints*, September 2012.

[Gra11] Robert M. Gray. *Entropy and information theory. 2nd ed.* New York, NY: Springer, 2nd ed. edition, 2011.

[Gri09] Alexander Grigor'yan. *Heat kernel and analysis on manifolds*, volume 47 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI, 2009.

[Gro81] Mikhael Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, (53):53–73, 1981.

[Gui80] Y. Guivarc'h. Sur la loi des grands nombres et le rayon spectral d'une marche aléatoire. In *Conference on Random Walks (Kleebach, 1979) (French)*, volume 74 of *Astérisque*, pages 47–98, 3. Soc. Math. France, Paris, 1980.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[HIH77] Ernst Heintze and Hans-Christoph Im Hof. Geometry of horospheres. *J. Differential Geom.*, 12(4):481–491 (1978), 1977.

[Hsu02] Elton P. Hsu. *Stochastic analysis on manifolds*, volume 38 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.

[Ich88] Kanji Ichihara. Comparison theorems for Brownian motions on Riemannian manifolds and their applications. *J. Multivariate Anal.*, 24(2):177–188, 1988.

[Kaĭ86] V. A. Kaĭmanovich. Brownian motion and harmonic functions on covering manifolds. An entropic approach. *Dokl. Akad. Nauk SSSR*, 288(5):1045–1049, 1986.

[Kaĭ88] V. A. Kaĭmanovich. Brownian motion on foliations: entropy, invariant measures, mixing. *Funktsional. Anal. i Prilozhen.*, 22(4):82–83, 1988.

[Kai92] Vadim A. Kaimanovich. Measure-theoretic boundaries of Markov chains, 0-2 laws and entropy. In *Harmonic analysis and discrete potential theory (Frascati, 1991)*, pages 145–180. Plenum, New York, 1992.

[Kec95] Alexander S. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

[Kin68] J. F. C. Kingman. The ergodic theory of subadditive stochastic processes. *J. Roy. Statist. Soc. Ser. B*, 30:499–510, 1968.

[Kin73] J. F. C. Kingman. Subadditive ergodic theory. *Ann. Probability*, 1:883–909, 1973. With discussion by D. L. Burkholder, Daryl Daley, H. Kesten, P. Ney, Frank Spitzer and J. M. Hammersley, and a reply by the author.

[KL07] Anders Karlsson and François Ledrappier. Propriété de Liouville et vitesse de fuite du mouvement brownien. *C. R. Math. Acad. Sci. Paris*, 344(11):685–690, 2007.

[KL11] Anders Karlsson and François Ledrappier. Noncommutative ergodic theorems. In *Geometry, rigidity, and group actions*, Chicago Lectures in Math., pages 396–418. Univ. Chicago Press, Chicago, IL, 2011.

[KV83] V. A. Kaĭmanovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy. *Ann. Probab.*, 11(3):457–490, 1983.

[KW02] Vadim A. Kaimanovich and Wolfgang Woess. Boundary and entropy of space homogeneous Markov chains. *Ann. Probab.*, 30(1):323–363, 2002.

[Led84] F. Ledrappier. Quelques propriétés des exposants caractéristiques. In *École d'été de probabilités de Saint-Flour, XII—1982*, volume 1097 of *Lecture Notes in Math.*, pages 305–396. Springer, Berlin, 1984.

[Led96] F. Ledrappier. Profil d'entropie dans le cas continu. *Astérisque*, (236):189–198, 1996. Hommage à P. A. Meyer et J. Neveu.

[Led10] François Ledrappier. Linear drift and entropy for regular covers. *Geom. Funct. Anal.*, 20(3):710–725, 2010.

[LS84] Terry Lyons and Dennis Sullivan. Function theory, random paths and covering spaces. *J. Differential Geom.*, 19(2):299–323, 1984.

[LS12] F. Ledrappier and L. Shu. Entropy rigidity of symmetric spaces without focal points. *ArXiv e-prints*, May 2012.

[Lu12] Peng Lu. Convergence of fundamental solutions of linear parabolic equations under Cheeger-Gromov convergence. *Math. Ann.*, 353(1):193–217, 2012.

[Pet06] Peter Petersen. *Riemannian geometry*, volume 171 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2006.

[Pin64] M.S. Pinsker. *Information and information stability of random variables and processes.* Holden-Day Series in Time Series Analysis. San Francisco-London-Amsterdam: Holden-Day, Inc. XII. 243 p. , 1964.

[Pra71] Jean-Jacques Prat. Étude asymptotique du mouvement brownien sur une variété riemannienne à courbure négative. *C. R. Acad. Sci. Paris Sér. A-B*, 272:A1586–A1589, 1971.

[Pra75] Jean-Jacques Prat. Étude asymptotique et convergence angulaire du mouvement brownien sur une variété à courbure négative. *C. R. Acad. Sci. Paris Sér. A-B*, 280(22):Aiii, A1539–A1542, 1975.

[Ree47] Georges Reeb. Variétés feuilletées, feuilles voisines. *C. R. Acad. Sci. Paris*, 224:1613–1614, 1947.

[Ree48] Georges Reeb. Remarque sur les variétés feuilletées contenant une feuille compacte à groupe de Poincaré fini. *C. R. Acad. Sci. Paris*, 226:1337–1339, 1948.

[Roh52] V. A. Rohlin. On the fundamental ideas of measure theory. *Amer. Math. Soc. Translation*, 1952(71):55, 1952.

[Sak83] Takashi Sakai. On continuity of injectivity radius function. *Math. J. Okayama Univ.*, 25(1):91–97, 1983.

[Sol70] Robert M. Solovay. A model of set-theory in which every set of reals is Lebesgue measurable. *Ann. of Math. (2)*, 92:1–56, 1970.

[Ste89] J. Michael Steele. Kingman's subadditive ergodic theorem. *Ann. Inst. H. Poincaré Probab. Statist.*, 25(1):93–98, 1989.

[Var86] Nicholas Th. Varopoulos. Information theory and harmonic functions. *Bull. Sci. Math. (2)*, 110(4):347–389, 1986.

[Wan11] Xiaodong Wang. Compactifications of complete Riemannian manifolds and their applications. In *Surveys in geometric analysis and relativity*, volume 20 of *Adv. Lect. Math. (ALM)*, pages 517–529. Int. Press, Somerville, MA, 2011.

[Wat90] William C. Waterhouse. The absolute-value estimate for symmetric multilinear forms. *Linear Algebra Appl.*, 128:97–105, 1990.