



HAL
open science

Intrinsic random walks and sub-Laplacians in sub-Riemannian geometry

Ugo Boscain, Robert Neel, Luca Rizzi

► **To cite this version:**

Ugo Boscain, Robert Neel, Luca Rizzi. Intrinsic random walks and sub-Laplacians in sub-Riemannian geometry. *Advances in Mathematics*, 2017, 10.1016/j.aim.2017.04.024 . hal-01122735v2

HAL Id: hal-01122735

<https://hal.science/hal-01122735v2>

Submitted on 4 Nov 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Intrinsic random walks and sub-Laplacians in sub-Riemannian geometry

UGO BOSCAIN¹, ROBERT NEEL², AND LUCA RIZZI¹

¹CNRS, CMAP École Polytechnique, and Équipe INRIA GECO Saclay Île-de-France, Paris, France

²Department of Mathematics, Lehigh University, 14 East Packer Avenue, Bethlehem, PA 18015, USA

March 20, 2015

Abstract

On a sub-Riemannian manifold we define two type of Laplacians. The *macroscopic Laplacian* Δ_ω , as the divergence of the horizontal gradient, once a volume ω is fixed, and the *microscopic Laplacian*, as the operator associated with a sequence of geodesic random walks. We consider a general class of random walks, where *all* sub-Riemannian geodesics are taken in account. This operator depends only on the choice of a complement \mathbf{c} to the sub-Riemannian distribution, and is denoted $L^\mathbf{c}$.

We address the problem of equivalence of the two operators. This problem is interesting since, on equiregular sub-Riemannian manifolds, there is always an intrinsic volume (e.g. Popp's one \mathcal{P}) but not a canonical choice of complement. The result depends heavily on the type of structure under investigation:

- On contact structures, for every volume ω , there exists a unique complement \mathbf{c} such that $\Delta_\omega = L^\mathbf{c}$.
- On Carnot groups, if H is the Haar volume, then there always exists a complement \mathbf{c} such that $\Delta_H = L^\mathbf{c}$. However this complement is *not unique* in general.
- For quasi-contact structures, in general, $\Delta_\mathcal{P} \neq L^\mathbf{c}$ for any choice of \mathbf{c} . In particular, $L^\mathbf{c}$ is not symmetric with respect to Popp's measure. This is surprising especially in dimension 4 where, in a suitable sense, $\Delta_\mathcal{P}$ is the unique intrinsic macroscopic Laplacian.

A crucial notion that we introduce here is the *N-intrinsic volume*, i.e. a volume that depends only on the set of parameters of the nilpotent approximation. When the nilpotent approximation does not depend on the point, a N-intrinsic volume is unique up to a scaling by a constant and the corresponding N-intrinsic sub-Laplacian is unique. This is what happens for dimension less than or equal to 4, and in particular in the 4-dimensional quasi-contact structure mentioned above.

Finally, we prove a general theorem on the convergence of a sequence of random walks to a diffusion, that gives, as a particular case, the convergence of the random walks mentioned above to the diffusion generated by $L^\mathbf{c}$.

Contents

1	Introduction	2
1.1	The Riemannian setting	2
1.2	The sub-Riemannian setting	4
1.2.1	The sub-Riemannian macroscopic Laplacian	5
1.2.2	On intrinsic volumes	5
1.2.3	The sub-Riemannian microscopic Laplacian	6
1.3	The equivalence problem	7
1.3.1	Contact structures	8
1.3.2	Carnot groups	8
1.3.3	Quasi-contact structures	9
1.3.4	Convergence of random walks	9
1.4	Comparison with recent literature	9
2	Preliminaries	11
2.1	Computations with frames	12
2.2	A Taylor expansion with frames	13

3	On volume forms in (sub)-Riemannian geometry	14
3.1	The nilpotent approximation	14
3.2	Popp's volume	15
3.2.1	Behavior under isometries	16
3.3	Intrinsic volumes	16
3.3.1	Equi-nilpotentizable structures	17
3.3.2	Homogeneous (sub)-Riemannian structures	18
4	The macroscopic Laplacian	18
5	The microscopic Laplacian	19
5.1	A class of adapted measures	19
5.2	Horizontal divergence	22
6	The equivalence problem	22
7	Carnot groups	23
8	Corank 1 structures	25
8.1	On natural Riemannian extensions and Popp's volume	26
9	Contact structures	26
9.1	Integrability conditions	26
10	Quasi-contact structures	27
10.1	The quasi-Reeb vector field	27
10.2	An example of non-existence	29
11	Convergence of random walks	30
11.1	A general class of random walks	30
11.2	Background on diffusions and pathspace	30
11.3	Control of the interpolating paths	31
11.4	Convergence results	32

1 Introduction

1.1 The Riemannian setting

Let M be a (smooth, connected, orientable, complete) n -dimensional Riemannian manifold. In Riemannian geometry the infinitesimal conservation condition for a smooth scalar quantity ϕ , a function of the point q and of the time t (for example, the temperature, the concentration of a diffusing substance, the noise, the probability density of a randomly-moving particle, etc.), that flows via a flux F (which says how much of ϕ is, infinitesimally, crossing a unit area of surface and in a unit of time) is expressed via the “continuity” equation $\partial_t \phi + \operatorname{div}(F) = 0$, where $\operatorname{div}(\cdot)$ is the divergence computed with respect to the Riemannian volume. If one postulates that the flux is proportional to minus the Riemannian gradient (the constant of proportionality being fixed to 1 for simplicity), that is $F = -\operatorname{grad}(\phi)$, one obtains the Riemannian heat equation

$$\partial_t \phi = \Delta \phi, \tag{1}$$

where $\Delta = \operatorname{div} \circ \operatorname{grad}$ is the Laplace-Beltrami operator. Since equation (1) has been obtained from a continuity equation that views ϕ as a fluid without a microscopic structure (the fluid is modeled as a continuous substance), in the following we refer to it as to the *macroscopic heat equation* and to the corresponding operator Δ as the *macroscopic Laplacian*. It is useful to write Δ in terms of an orthonormal frame. If X_1, \dots, X_n is a local orthonormal frame for the Riemannian structure we have the formulas

$$\operatorname{grad}(\phi) = \sum_{i=1}^n X_i(\phi) X_i \quad \text{and} \quad \Delta(\phi) = \sum_{i=1}^n (X_i^2 + \operatorname{div}(X_i) X_i)(\phi).$$

It is well known that the heat equation (1) also admits a stochastic interpretation as the evolution equation associated to a diffusion process. Actually, we need to be more precise here, since there are two evolution equations associated to a diffusion process. In particular, there is the evolution of the expectation of a test function $f \in C_0^\infty(M)$ evaluated along the paths of the diffusion, called the forward Kolmogorov equation, and

there is the evolution equation for the transition density of the diffusion (relative to a smooth volume), called the backward Kolmogorov equation. The second-order operators in these two equations are adjoints (with respect to the same smooth volume as the transition density). However, thanks to the geodesic completeness assumption, the Laplace-Beltrami operator is essentially self-adjoint (with respect to the Riemannian volume) so that both operators, and thus both parabolic PDEs, coincide.

Moreover, the diffusion process associated to the (Riemannian) heat equation can be obtained as a (parabolic) scaling limit of a random walk. We think of the random walk as giving a “microscopic” viewpoint on the evolution of the quantity measured by ϕ , since, at least in an idealized sense, it models the motion of an individual particle. In taking the scaling limit, we pass from the small-scale behavior of an individual particle to the large-scale average behavior of the particle, or equivalently, to the aggregate behavior of a large number of such particles, and thus pass from the microscopic to the macroscopic viewpoint. We now need to explain these random walks and the associated ideas in more detail.

In the Riemannian case currently under discussion, we are interested in an isotropic random walk. In particular, starting from a point q , the unit sphere \mathbb{S}^{n-1} in the tangent space T_qM has the standard $(n-1)$ -volume induced by the metric on T_qM . Normalizing this to a probability measure μ_q (by dividing by the total volume $2\pi^{n/2}\Gamma(n/2)$), we can then choose a direction $\theta \in \mathbb{S}^{n-1}$ at q randomly, in a way which is obviously isotropic with respect to the Riemannian structure. The particle then travels along the geodesic tangent to θ for a distance ε in time δt (at constant speed, determined by these conditions). We let X_t^ε be the (random) position of the particle at time $t \in [0, \delta t]$. We can continue this process by next choosing a direction, at the point $X_{\delta t}^\varepsilon$, at random via the measure $\mu_{X_{\delta t}^\varepsilon}$ on the unit sphere $\mathbb{S}^{n-1} \subset T_{X_{\delta t}^\varepsilon}M$, and independently from the previous choice of direction q . Then the particle follows the geodesic from $X_{\delta t}^\varepsilon$ in this direction for distance ε in time δt . We can continue this process indefinitely, since after i steps, the particle has traveled along a piecewise geodesic for a total distance of $i\varepsilon$, and since we are assuming that M is geodesically complete, a next step can always be made (equivalently, this process cannot explode by exiting every compact set in finite time). The result is a random (piecewise geodesic) path X_t^ε (for $t \in [0, \infty)$) starting from $q = X_0^\varepsilon$. Further, the positions of the particle at the times when it randomly changes direction, namely, the sequence $X_0^\varepsilon, X_{\delta t}^\varepsilon, X_{2\delta t}^\varepsilon, \dots$ is a Markov chain on M , and for $t \in (i\delta t, (i+1)\delta t)$, X_t^ε interpolates between $X_{i\delta t}$ and $X_{(i+1)\delta t}$ along a geodesic between them (for $i = 0, 1, 2, \dots$).

As mentioned, we are interested in a scaling limit of such random walks (and this is why we have already indexed our walk X_t^ε by the step-size ε). We will take $\delta t = \varepsilon^2/\alpha$, where α is a normalizing constant to be chosen later. Then $\delta t \rightarrow 0$ as $\varepsilon \rightarrow 0$, and moreover, this is the parabolic scaling (which we note is in a sense responsible for the infinite velocity of propagation of the heat).

We are interested in the behavior of a single step under this parabolic scaling, which by the Markov property and homogeneity in time may as well be the first step. If we consider the change in the expectation of a function $\phi \in C_0^\infty(M)$ sampled after the step, normalized by dividing by δt , we obtain an operator which we denote by L_ε . More concretely, we have

$$\begin{aligned} (L_\varepsilon\phi)(q) &= \frac{\mathbb{E}[\phi(X_{\delta t}^\varepsilon) | X_0^\varepsilon = q] - \phi(q)}{\delta t} \\ &= \frac{\alpha}{\varepsilon^2} \left(\mathbb{E} \left[\phi \left(X_{\varepsilon^2/\alpha}^\varepsilon \right) \middle| X_0^\varepsilon = q \right] - \phi(q) \right) \\ &= \frac{\alpha}{\varepsilon^2} \left(\int_{\mathbb{S}^{n-1}} \phi(\exp_q(\varepsilon, \theta)) \mu_q(\theta) - \phi(q) \right), \end{aligned}$$

where $\exp_q(\varepsilon, \theta)$ is the point obtained by following the arc-length parametrized geodesic from q in the direction of θ for distance ε .

It turns out that the limiting behavior of the sequence of random walks, as $\varepsilon \rightarrow 0$, is governed by the limiting behavior of these operators. Indeed, Section 11 is devoted to a discussion of this issue and to the proof of theorems on the convergence of a sequence of random walks to a diffusion in the Riemannian or sub-Riemannian context (see Theorem 70 for a general convergence result and Theorem 71 for the case of random walks of the type just described). At any rate, expanding the exponential map in normal coordinates around q shows that L_ε converges to an operator L as $\varepsilon \rightarrow 0$ (in the sense that $L_\varepsilon\phi \rightarrow L\phi$ uniformly on compacts for $\phi \in C_0^\infty(M)$), and that $L = \frac{\alpha}{2n}\Delta$ (the computation is a special case of the (sub)-Riemannian computations below, so we don't reproduce it). It follows that the sequence of random walks converges to the diffusion generated by L (or the diffusion that solves the martingale problem for L), which we denote X_t^0 . Thus, if ϕ is a smooth scalar quantity depending on a point q and a time t such that its evolution is governed infinitesimally by the dynamics of the random walk (in the sense that $\partial_t\phi(q, t) = \lim_{\delta \rightarrow 0} \frac{1}{\delta} \mathbb{E}[\phi(X_\delta^0, t) - \phi(q) | X_0^0 = q]$), we obtain the equation

$$\partial_t\phi = L\phi,$$

where

$$L\phi(q, t) = \lim_{\varepsilon \rightarrow 0} \frac{\alpha}{\varepsilon^2} \left(\int_{\mathbb{S}^{n-1}} \phi(\exp_q(\varepsilon, \theta), t) \mu_q(\theta) - \phi(q, t) \right).$$

In the following we refer to the equation (2) as to the “microscopic heat equation” and to L as to the “microscopic Laplacian”. As we’ve already seen, if we take $\alpha = 2n$, then we have

$$\Delta = L. \tag{4}$$

(As an aside, the need for the normalizing constant α to grow with the dimension n is a manifestation of the concentration of measure on the sphere.) Hence the macroscopic diffusion equation and the microscopic one coincide.

Further, define the heat kernel $p_t(q, q_0)$ as the density of the random variable $X_t^0 | X_0^0 = q_0$ with respect to the Riemannian volume \mathcal{R} . Said more analytically, p_t is the fundamental solution to Equation (2), so that

$$\phi(q_0, t) = \int_M \phi(q) p_t(q, q_0) \mathcal{R}(q) = \mathbb{E} [\phi(X_t^0) - \phi(q_0) | X_0^0 = q_0]$$

solves the Cauchy problem for Equation (2) with initial condition $\phi(q_0, 0) = \phi(q_0)$. Then because L is essentially self-adjoint with respect to the Riemannian volume (it’s just a scalar multiple of Δ), $p_t(\cdot, q_0)$ also satisfies the microscopic heat equation 2. Thus p_t , which measures the probability density of the random paths themselves, rather than a quantity that is sampled along them, can be understood in terms of the same equations.

All of this can be viewed in the following way. On one side, the microscopic perspective is a good interpretation of the macroscopic heat equation. On the other side, the microscopic Laplacian L is a good operator because it is essentially self-adjoint with respect to a volume (the Riemannian one). This is due to the fact that it is symmetric since it can be written in divergence form thanks to Equation (4) and to the geodesic completeness of the manifold.

The essential self-adjointness of L with respect to some volume ω , even if not necessary to define the corresponding process, is important, both because it means that the heat kernel satisfies the same equation and because it permits one to study the evolution equation (2) in $L^2(M, \omega)$. See Theorem 45 below.

Remark 1. The Riemannian volume can be defined equivalently as the volume \mathcal{R} such that $\mathcal{R}(X_1, \dots, X_n) = 1$ for any oriented local orthonormal frame X_1, \dots, X_n , or as the n -dimensional Hausdorff or spherical-Hausdorff volume (up to some constant). The above construction gives an alternative way of characterizing the Riemannian volume. It is the unique volume (up to constant rescaling) such that the microscopic Laplacian can be written in divergence form. These fact are much less trivial in the sub-Riemannian context. Indeed in sub-Riemannian geometry there are several notions of intrinsic volumes and the definition of the microscopic Laplacian requires additional structure.

1.2 The sub-Riemannian setting

In this paper, a sub-Riemannian structure is a triple $(M, \blacktriangle, \mathbf{g})$, where M is a n -dimensional differentiable manifold, \blacktriangle is a smooth distribution of constant rank $k < n$ satisfying the Hörmander condition and \mathbf{g} is a Riemannian metric on \blacktriangle . Locally the structure can be assigned by an orthonormal frame $X_1, \dots, X_k \in \Gamma(\blacktriangle)$. Here, $\Gamma(E)$ denotes the $C^\infty(M)$ -module of smooth sections of any vector bundle E over M .

With the term (sub)-Riemannian, we mean structures that can be also Riemannian, i.e. defined as above but with $k \leq n$. Riemannian manifolds are by definition equiregular. All the information about the structure is contained in the Hamiltonian function $H : T^*M \rightarrow \mathbb{R}$ defined by

$$H(\lambda) = \frac{1}{2} \sum_{i=1}^k \langle \lambda, X_i \rangle^2, \quad \lambda \in T^*M,$$

where $\langle \lambda, \cdot \rangle$ denotes the action of covectors on vectors. It is a well known fact in (sub)-Riemannian geometry that integral lines $\lambda(\cdot)$ of the Hamiltonian flow defined by H , laying on the level set $H = 1/2$, project to smooth curves $\gamma(t) := \pi(\lambda(t))$ on M that are arc-length parametrized geodesics, i.e. $\|\dot{\gamma}(t)\| = 1$ and, for every sufficient small interval $[t_1, t_2]$, the restriction $\gamma|_{[t_1, t_2]}$ is a minimizer of the *sub-Riemannian length* $\ell(\gamma) = \int_{t_1}^{t_2} \|\dot{\gamma}(t)\| dt$ (with fixed endpoints). Here $\|\cdot\|$ denotes the norm induced by \mathbf{g} on \blacktriangle . In Riemannian geometry these curves are precisely all the arc-length parameterized geodesics of the structure. In the sub-Riemannian case these are called *normal geodesics*. There is also another class of geodesics called *abnormal geodesics*, that may not follow the above dynamic.

1.2.1 The sub-Riemannian macroscopic Laplacian

As we will see later, in sub-Riemannian geometry the definition of an intrinsic volume, playing the role of the Riemannian volume is a subtle question. For the moment, let us assume that a volume ω is fixed on the sub-Riemannian manifold. In this case one can write the macroscopic heat equation similarly to the Riemannian case. The only difference is that one should postulate that the flux is proportional to the horizontal gradient. The horizontal gradient $\text{grad}_H(\cdot)$ of a C^∞ function ϕ is defined similarly to the Riemannian gradient but is a vector field belonging to the distribution (see, for instance, [6]):

$$\mathbf{g}_q(v, \text{grad}_H(\phi)_q) = d_q\phi(v), \quad \forall v \in \mathbf{A}_q.$$

We have then for the macroscopic heat equation

$$\partial_t \phi = \Delta_\omega \phi,$$

where

$$\Delta_\omega = \text{div}_\omega \circ \text{grad}_H$$

is the macroscopic Laplacian. In terms of a local orthonormal frame of the sub-Riemannian manifold we have the same formulas as in the Riemannian case, but summing up to the rank of the distribution:

$$\text{grad}_H(\phi) = \sum_{i=1}^k X_i(\phi)X_i, \quad \Delta_\omega(\phi) = \sum_{i=1}^k (X_i^2 + \text{div}_\omega(X_i)X_i)(\phi).$$

In the following we always suppress the H from grad_H (the horizontal gradient coincides with the standard one in the Riemannian case).

1.2.2 On intrinsic volumes

The difficulty concerning the definition of a macroscopic Laplacian is related to the choice of ω . What is the “correct” volume in the sub-Riemannian case, the analogue to the Riemannian one? One needs some algorithm to assign, with any sub-Riemannian structure on M , a volume form ω_M . Moreover, the construction of ω_M should depend, loosely speaking, only on the metric invariants of the structure.

Definition 1. An *intrinsic definition of volume* is a map that associates, with any (oriented) (sub)-Riemannian structure $(M, \mathbf{A}, \mathbf{g})$ a volume form ω_M on M such that if $\phi : M \rightarrow N$ is a sub-Riemannian isometry between $(M, \mathbf{A}_M, \mathbf{g}_M)$ and $(N, \mathbf{A}_N, \mathbf{g}_N)$, then $\phi^* \mu_N = \mu_M$.

Remark 2. In order to avoid this verbose terminology, with the term “intrinsic volume” we mean either the actual volume form ω_M or the definition of volume given by a map $M \mapsto \omega_M$.

Even in the Riemannian case, there are many intrinsic volumes. The classical Riemannian one is the unique volume form \mathcal{R} such that $\mathcal{R}(X_1, \dots, X_n) = 1$ for any (oriented) orthonormal frame. But what about the volume form defined by $\mathcal{R}'(X_1, \dots, X_n) = 1 + \kappa^2$, where κ is the scalar curvature? Both are perfectly fine definitions of volume, according to our definition. The first, loosely speaking, is more “intrinsic” than the second. In fact, it is true that both depend only on the metric invariant of the structure, but \mathcal{R}' involves second order information about the structure. To rule out \mathcal{R}' and select the classical Riemannian volume we need a more precise definition.

Roughly speaking we say that an intrinsic definition of volume is *N-intrinsic* if its value of ω_M at q depends only on the metric invariants of the nilpotent approximation of the structure at q (the metric tangent space to the structure). For the precise definition see Section 3. In the Riemannian case, there is only one nilpotent approximation, which is the flat \mathbb{R}^n , and it has no non-trivial metric invariants. As a consequence, there is a unique N-intrinsic volume, the Riemannian one. As it is well known, in the sub-Riemannian case, nilpotent approximations may be different at different points.

Definition 2. A (sub)-Riemannian manifold is *equi-nilpotentizable* if nilpotent approximations at different points are isometric.

As just said above, Riemannian manifolds are equi-nilpotentizable. Equi-nilpotentizable structures are equiregular (see Section 2). By definition, all Carnot groups are equi-nilpotentizable. It is well known that all equiregular sub-Riemannian structures in dimension less than or equal to four are equi-nilpotentizable (see [4]). In particular this is true for 3D sub-Riemannian manifolds (the nilpotent approximation is given by the Heisenberg group \mathbb{H}_3 at each point). The simplest non-equi-nilpotentizable sub-Riemannian structure is given by a generic structure of rank 4 in dimension 5. Other examples of non-equi-nilpotentizable sub-Riemannian

manifolds are given by generic contact manifolds of dimension greater than or equal to 5 and by generic quasi-contact sub-Riemannian manifolds in dimension greater than or equal to 6.

A general N-intrinsic volume, the Popp's volume, well defined for any equiregular sub-Riemannian structure was given by Montgomery in his book [25] and later studied in [6, 8, 4]. This definition was indeed pioneered by Brockett for contact structures in [12]. In the following we call this volume \mathcal{P} . It turns out that, in the Riemannian case, Popp's construction recovers the Riemannian volume. Moreover, in [4], it has been proved that, for non-equinilpotentizable structures, Popp's volume does not coincide with the (spherical) Hausdorff measure. Finally, in Section 3 we prove the following results.

Proposition 3. *Let $(M, \mathbf{\Delta}, \mathbf{g})$ be a equi-nilpotentizable (sub)-Riemannian manifold. Then Popp's volume \mathcal{P} is the unique N-intrinsic definition of volume, up to a multiplicative constant.*

In this sense Popp's volume generalizes the Riemannian one. When ω is an N-intrinsic volume, we say that Δ_ω is an *N-intrinsic macroscopic Laplacian*. Since the divergence of a vector field does not change if the volume is multiplied by a constant, we have the following.

Proposition 4. *Let $(M, \mathbf{\Delta}, \mathbf{g})$ be a equi-nilpotentizable (sub)-Riemannian manifold. Then there exists a unique N-intrinsic macroscopic Laplacian, $\Delta_{\mathcal{P}}$ that is built with Popp's volume.*

1.2.3 The sub-Riemannian microscopic Laplacian

Another problem in the sub-Riemannian case is the definition of microscopic Laplacian, as limit process coming from a random walk along geodesics starting from a given point. Indeed, while in the Riemannian context the geodesics starting from a given point can always be parametrized by the direction of the initial velocity, i.e., by the points of a $(n - 1)$ -dimensional sphere, in the sub-Riemannian context the geodesics starting from a given point are always parameterized by a non-compact set, namely by the points of a cylinder $\square_q = H^{-1}(1/2) \cap T_q^*M$ having the topology of $\mathbb{S}^{k-1} \times \mathbb{R}^{n-k}$. How to define an intrinsic finite volume on \square_q , and thus a probability measure, is a non-trivial question.

Moreover, while in Riemannian geometry, for sufficiently small ε the Riemannian metric sphere of radius ε centered at q coincides with the endpoints of the geodesics starting from q and having length ε , in sub-Riemannian geometry this is not true: for a fixed a point q , there are geodesics starting from q that lose optimality arbitrarily close to q . Hence one has to decide whether to average only on the sub-Riemannian sphere (i.e., on geodesics that are optimal up to length ε) or on the sub-Riemannian front (i.e., on all geodesics of length ε). Another approach is to choose a k -dimensional linear subspace \mathbf{h}_q in T_q^*M transverse to the non compact directions of the cylinder (playing the role of the "most horizontal" geodesics) and to average on $\mathbf{h}_q \cap \square_q$. This last approach is essentially the one followed in [15, 16, 17, 18]. Certainly the problem of finding a canonical subspace \mathbf{h}_q in T_q^*M is a non-trivial question and in general it is believed that it is not possible.

In any case, all these problems are indeed encompassed in the specification of a measure μ_q (possibly singular) on \square_q for all $q \in M$. Once such a collection of measures $\mu = \{\mu_q\}_{q \in M}$ is fixed, we define the *microscopic Laplacian* as

$$(L_\mu \phi)(q) := \lim_{\varepsilon \rightarrow 0} \frac{\alpha}{\varepsilon^2} \left(\int_{\square_q} \phi(\exp_q(\varepsilon, \lambda)) \mu_q(\lambda) - \phi(q) \right), \quad (5)$$

assuming the limit exists. Here $t \mapsto \exp_q(t, \lambda)$ is the arc-length parametrized normal geodesics starting from q with initial covector $\lambda \in \square_q$. Moreover α is a constant of proportionality that will be fixed later.

Remark 3. In this paper, the microscopic Laplacian is built only with the geodesics of the exponential map that, by definition, are normal. In sub-Riemannian geometry there is also another type of geodesic, called (strict) *abnormal*, that are not described by the sub-Riemannian exponential map (i.e. they are not the projection of the Hamiltonian flow described above). The "size" of the set of points reached by abnormal geodesics (starting from a fixed origin) is a hard open problem known as the *Sard conjecture in sub-Riemannian geometry*, see for instance [2, 27]). It is believed that only a set of measure zero is reached by abnormal geodesics. What is known in general is that the set of points reached by optimal normal geodesics is open and dense (see [1, 5] for precise statements). We notice that, on contact structures, there are no nontrivial abnormal geodesics.

If one would like to include also abnormal geodesics for the construction of the microscopic Laplacian, one should decide which measure give to them and in principle one could get a different operator. This research direction is beyond the purpose of the present paper.

In view of Theorem 5, we restrict the class of possible measures, and we consider only those induced by a complement as follows. For all $q \in M$, consider a complement \mathbf{c}_q such that $T_q M = \mathbf{\Delta}_q \oplus \mathbf{c}_q$. By duality $T_q^* M = \mathbf{v}_q \oplus \mathbf{h}_q$, where $\mathbf{v}_q := \mathbf{\Delta}_q^\perp$ (resp. $\mathbf{h}_q := \mathbf{c}_q^\perp$) denote the annihilators of $\mathbf{\Delta}_q$ (resp. \mathbf{c}_q). We can see \mathbf{v}_q as the space of "vertical" covectors, and \mathbf{h}_q the space of "horizontal" ones. Now we can define an Euclidean structure on \mathbf{h}_q by identifying it with $\mathbf{\Delta}_q$ (this is equivalent to considering the restriction $2H|_{\mathbf{h}_q}$, which is a

positive definite quadratic form). The intersection $\mathbb{S}_q^{k-1} = \square_q \cap \mathbf{h}_q$ (See Figure 1 in Section 5) is precisely the Euclidean sphere in this space. Then the cylinder of initial covectors splits as

$$\square_q = \mathbb{S}_q^{k-1} \times \mathbf{v}_q.$$

We stress that this identification depends on the choice of \mathbf{c} . We restrict to the class of product measures induced by the choice of a complement on \square_q of the form

$$\mu_{\mathbf{c}_q} = \mu_{\mathbb{S}_q^{k-1}} \times \mu_{\mathbf{v}_q}, \quad (6)$$

where $\mu_{\mathbb{S}_q^{k-1}}$ is the Euclidean probability measure on $\mathbb{S}_q^{k-1} = \square_q \cap \mathbf{h}_q$ and $\mu_{\mathbf{v}_q}$ is any probability measure on \mathbf{v}_q . Moreover, we assume that $\mu_{\mathbf{v}_q}$ is sufficiently regular as a function of q , and 2-decreasing, namely that any linear function on \mathbf{v}_q is $L^2(\mathbf{v}_q, \mu_{\mathbf{v}_q})$ (see Definition 50). Any such a measure is called *adapted*.

Repeating this construction at each point we recover a differential operator, as in Eq. (5), that we call $L_{\mu_{\mathbf{c}}}$. It turns out that the latter does not depend on the choice of adapted measure, but only on the complement \mathbf{c} . In Section 5 we prove the following main result.

Theorem 5. *Let $TM = \mathbf{A} \oplus \mathbf{c}$ and let $\mu_{\mathbf{c}}$ any adapted measure. Then $L_{\mu_{\mathbf{c}}}$ depends only on \mathbf{c} . Moreover, let X_1, \dots, X_k a local orthonormal frame for \mathbf{A} , and X_{k+1}, \dots, X_n a local frame for \mathbf{c} . Then*

$$L_{\mu_{\mathbf{c}}} = \sum_{i=1}^k X_i^2 + \sum_{i,j=1}^k c_{ji}^j X_i,$$

where the structural functions $c_{ij}^\ell \in C^\infty(M)$ are defined by $[X_i, X_j] = \sum_{\ell=1}^n c_{ij}^\ell X_\ell$ for all $i, j = 1, \dots, n$. Finally, the convergence of Eq. (5) is uniform on compact sets.

Observe that the equiregularity assumption is not needed here. Thanks to Theorem 5, in the following we use the notation $L^{\mathbf{c}} := L_{\mu_{\mathbf{c}}}$.

Remark 4. One of the byproducts of Theorem 5 is the following. Fixing \mathbf{c} is equivalent to assign a subspace of “horizontal covectors” \mathbf{h}_q in T_q^*M . The expression of $L^{\mathbf{c}}$ at $q \in M$ is the same averaging only on the horizontal geodesics (i.e. with a measure $\mu_{\mathbf{c}_q} = \mu_{\mathbb{S}_q^{k-1}} \times \delta_{\mathbf{v}_q}$, where δ is the Dirac delta) or averaging on all possible geodesics with a measure of the type (6). The particular choice $\mu_{\mathbf{c}_q} = \delta_{\mathbf{v}_q}$ recovers the construction of [17, 18, 15, 16]), where the authors choose a Riemannian extension and use this to define the space of horizontal covectors.

Remark 5. We can rewrite the operator $L^{\mathbf{c}}$ in a more elegant way by introducing the concept of *horizontal divergence*. As we will make precise in Section 5.2, the horizontal divergence of a vector field X computes the infinitesimal change of volume, under the flow of a vector field X , of the standard parallelotope of \mathbf{A} . To do this, we need a well defined projection $\pi : TM \rightarrow \mathbf{A}$ that, indeed, requires the choice of a complement \mathbf{c} . We denote with $\text{div}^{\mathbf{c}}(X)$ the horizontal divergence of X , and we have

$$L^{\mathbf{c}} = \text{div}^{\mathbf{c}} \circ \text{grad}.$$

1.3 The equivalence problem

Once a volume ω on M is chosen (hence a macroscopic Laplacian is defined) and a complement \mathbf{c} is fixed (hence a microscopic Laplacian is defined), it is natural to ask:

Q1: Under which conditions on ω and \mathbf{c} do we have $\Delta_\omega = L^{\mathbf{c}}$?

In other words, we would like to know when a macroscopic Laplacian admits a microscopic interpretation and when a microscopic Laplacian can be written in divergence form (and hence is symmetric) w.r.t. some volume on the manifold. Moreover:

Q2: Given a volume ω , is it possible to find a complement \mathbf{c} such that $\Delta_\omega = L^{\mathbf{c}}$? If so, is it unique?

This question is interesting since on any sub-Riemannian structure there is a smooth, intrinsic volume, Popp’s one. Then an answer to **Q2** gives a way to assign an intrinsic complement \mathbf{c} . A counting argument suggests that such a question has an affirmative answer. In fact, a volume form is given by a non-zero function, while a complement is given by $(n - k)k$ functions. However the answer is more complicated because some integrability conditions must be taken in account. A more specific question is the following:

Q3: Let \mathbf{c} a complement, and let $\mathbf{g}_{\mathbf{c}}$ a smooth scalar product on \mathbf{c} . Then the orthogonal direct sum $\mathbf{g} \oplus \mathbf{g}_{\mathbf{c}}$ is a *Riemannian extension* of the sub-Riemannian structure (also called taming metric). Let $\omega_{\mathbf{c}}$ be the corresponding Riemannian volume. Is it true that for the sub-Riemannian diffusion operator $\Delta_{\omega_{\mathbf{c}}} = L^{\mathbf{c}}$?

This last question is even more interesting when it is possible to find an intrinsic Riemannian extension, i.e.

some choice of \mathbf{c} and $\mathbf{g}_{\mathbf{c}}$ that depends only on the sub-Riemannian structure $(M, \mathbf{\Delta}, \mathbf{g})$. Intrinsic Riemannian extensions can be made in several cases (see for instance [14, 19, 20]). However, in general they are not known and (even if this is a non-apophantic statement) it is believed that they do not exist.

In Section 6 we answer to **Q1**, with the following theorem.

Theorem 6. *For any complement \mathbf{c} and volume ω , the macroscopic operator Δ_ω and the microscopic operator $L^{\mathbf{c}}$ have the same principal symbol and no constant term. Moreover $L^{\mathbf{c}} = \Delta_\omega$ if and only if*

$$\chi^{(\mathbf{c}, \omega)} := L^{\mathbf{c}} - \Delta_\omega = \sum_{i=1}^k \sum_{j=k+1}^n c_{ji}^{\ell} X_i + \text{grad}(\theta) = 0,$$

where $\theta = \log |\omega(X_1, \dots, X_n)|$ and c_{ij}^{ℓ} are the structural functions associated with an orthonormal frame X_1, \dots, X_k for $\mathbf{\Delta}$ and a frame X_{k+1}, \dots, X_n for \mathbf{c} .

The condition $\chi^{(\mathbf{c}, \omega)} = 0$ is a coordinate-free version of [15, Th. 5.13], which appeared online while this paper was under redaction. In [17] the same condition is obtained when \mathbf{c} is the orthogonal complement w.r.t. some Riemannian extension and ω is the corresponding Riemannian volume. The particularly simple form of Eq. 6 in our frame formalism permits us to go further in the study of its solutions.

We address **Q2** and **Q3** for Carnot groups, contact and quasi-contact structures and, more generally, corank 1 structures. Here we collect only the main results. For precise definition see the appropriate sections.

1.3.1 Contact structures

Theorem 7. *Let $(M, \mathbf{\Delta}, \mathbf{g})$ be a contact sub-Riemannian structure. For any volume ω there exists a unique complement \mathbf{c} such that $L^{\mathbf{c}} = \Delta_\omega$. In this case $\mathbf{c} = \text{span}\{X_0\}$, with*

$$X_0 = Z - J^{-1} \text{grad}(\theta), \quad \theta = \log |\omega(X_1, \dots, X_k, Z)|,$$

where Z is the Reeb vector field and $J : \mathbf{\Delta} \rightarrow \mathbf{\Delta}$ is the contact endomorphism.

Contact structures have a natural Riemannian extension, obtained by declaring the Reeb vector field a unit vector orthonormal to $\mathbf{\Delta}$. It turns out that the Riemannian volume of this extension is Popp's volume.

Corollary 8. *Let \mathcal{P} be the Popp's volume. The unique complement \mathbf{c} such that $L^{\mathbf{c}} = \Delta_{\mathcal{P}}$ is generated by the Reeb vector field. Moreover, \mathcal{P} is the unique volume (up to constant rescaling) with this property.*

Remark 6. In the results above we always use the normalization $\|J\| = 1$ (this fixes the contact form up to a sign). We stress that if we choose a different normalization the Reeb field would be different.

In Section 9.1 we also discuss the inverse problem, namely for a fixed \mathbf{c} , find a volume ω such that $L^{\mathbf{c}} = \Delta_\omega$. This is a more complicated problem (and in general has no solution). In the contact case, thanks to the non-degeneracy of J , we find explicitly a necessary and sufficient condition.

Proposition 9. *Let $\mathbf{c} = \text{span}\{X_0\}$. Define the one-form $\alpha := \frac{i_{X_0} d\eta}{\eta(X_0)}$ and the function $g = \frac{d\alpha \wedge \eta \wedge (d\eta)^{d-1}}{\eta \wedge (d\eta)^d}$. Then there exists a volume ω such that $L^{\mathbf{c}} = \Delta_\omega$ if and only if*

$$d\alpha - dg \wedge \eta - g d\eta = 0.$$

In this case, ω is unique up to constant rescaling.

1.3.2 Carnot groups

By definition, Carnot groups are particular cases of left-invariant sub-Riemannian structures (see Definition 42). Then it is natural to choose a left-invariant volume i.e. a volume proportional to the Haar's one (for example, Popp's volume \mathcal{P}) and left-invariant complements \mathbf{c} . Contrary to the contact case (where we always have existence and uniqueness of \mathbf{c} for any fixed ω) here we lose uniqueness.

Proposition 10. *For any Carnot group G , we have $\Delta_{\mathcal{P}} = L^{\mathbf{c}_0}$, where \mathbf{c}_0 is the left invariant complement*

$$\mathbf{c}_0 := \mathfrak{g}_2 \oplus \dots \oplus \mathfrak{g}_m,$$

where \mathfrak{g}_i denotes the i -th layer of the Carnot group.

Any distinct left-invariant complement is the graph of a non-trivial linear map $\ell : \mathbf{c}_0 \rightarrow \mathbf{\Delta}$, that is

$$\mathbf{c}_\ell := \{X + \ell(X) \mid X \in \mathbf{c}_0\}.$$

Proposition 11. *$\Delta_{\mathcal{P}} = L^{\mathbf{c}}$ if and only if $\mathbf{c} = \mathbf{c}_\ell$, with*

$$\text{Tr}(\ell \circ \text{ad}_{X_i}) = 0, \quad \forall i = 1, \dots, k.$$

The equation above has not unique solution ℓ in many cases, as we show in Section 7.

1.3.3 Quasi-contact structures

For all the relevant definitions, we refer to Section 10. We only stress that, analogously to the contact case, one can define a quasi-contact endomorphism $J : \blacktriangle \rightarrow \blacktriangle$ that, in the quasi-contact case, is degenerate (and its kernel is usually assumed to have dimension 1). In this case, if for some ω there is \mathbf{c} such that $L^{\mathbf{c}} = \Delta_{\omega}$, then \mathbf{c} is *never* unique. In particular, we have the following.

Theorem 12. *For any fixed ω , the space of \mathbf{c} such that $L^{\mathbf{c}} = \Delta_{\omega}$ is an affine space over $\ker J$.*

Surprisingly, we not only lose uniqueness but also existence. Consider the quasi-contact structure on $M = \mathbb{R}^4$, with coordinates (x, y, z, w) defined by $\blacktriangle = \ker \eta$ with

$$\eta = \frac{g}{\sqrt{2}} \left(dw - \frac{y}{2} dx + \frac{x}{2} dy \right), \quad \text{with} \quad g = e^z.$$

(g can be any strictly monotone, positive function). The metric is defined by the global orthonormal frame:

$$X = \frac{1}{\sqrt{g}} \left(\partial_x + \frac{1}{2} y \partial_w \right), \quad Y = \frac{1}{\sqrt{g}} \left(\partial_y - \frac{1}{2} x \partial_w \right), \quad Z = \frac{1}{\sqrt{g}} \partial_z.$$

Choose $\omega = \mathcal{P}$, the Popp's volume, that is

$$\mathcal{P} = \frac{g^{5/2}}{\sqrt{2}} dx \wedge dy \wedge dz \wedge dw.$$

Proposition 13. *For the structure above $L^{\mathbf{c}} \neq \Delta_{\mathcal{P}}$ for any choice of the complement \mathbf{c} .*

Even though this is out of the scope of the present paper, it turns out that this is the typical (generic) picture in the quasi-contact case. This result is quite surprising because, on sub-Riemannian structures with dimension smaller or equal than 4, there is a unique N-intrinsic volume up to scaling, and a unique N-intrinsic Laplacian given by $\Delta_{\mathcal{P}}$. In our example, this unique N-intrinsic Laplacian has *no* compatible complement or, in other words, the macroscopic diffusion operator has no microscopic counterpart.

On a quasi-contact manifold one can build an analogue of the Reeb vector field (actually, a one-parameter family parametrized by the distinct eigenvalues of J), that provides a standard Riemannian extension (see Section 10.1). It turns out that the Riemannian volume of these extensions is once again Popp's volume. Thus, the above non-existence result provides also an answer to **Q3**: the macroscopic operator Δ_{ω} provided by the quasi-Reeb Riemannian extension is not the microscopic operator $L^{\mathbf{c}}$ provided by the quasi-Reeb complement.

1.3.4 Convergence of random walks

Finally, in Section 11, we provide the probabilistic side of the construction of the microscopic Laplacian. Namely, as we see in equation (5), the microscopic Laplacian is built from the scaling limit of single step of a random walk. In Theorem 71, we show that this convergence of a single-step (in the sense of the convergence of the induced operator on smooth functions) can be promoted to the convergence of the random walks to the diffusion generated by the limit operator. This connects the microscopic Laplacian back to the heuristic discussion in the Section 1.1. Moreover, in Theorem 70, the convergence of a sequence of random walks to the diffusion generated by some appropriate second-order operator is established in much more generality than required for the present paper (for example, the probability measure for choosing a co-vector at each step of the walk can be supported on the entire co-tangent space), with an eye toward a wider variety of possible approximation schemes.

1.4 Comparison with recent literature

In the Riemannian setting, different authors have analyzed the convergence of geodesic random walks to the Brownian motion generated by the Laplace-Beltrami operator. In particular, in [26], Pinsky considers a process (a random flight in our terminology, that he calls the isotropic transport process) that, starting from $x \in M$, moves along a geodesic with initial vector uniformly chosen on the unit-sphere, for a randomly chosen, exponentially-distributed time, after which a new vector is chosen. The lift of this walk to the tangent bundle is a Markov process that can be understood in terms of the corresponding semi-group. Using this, the author shows that under appropriate (parabolic) rescaling, the semi-group of the random flight converges to the Brownian motion semigroup (whose generator is the Laplace-Beltrami operator). In [22], Lebeau and Michel investigate a random walk on smooth, compact, connected Riemannian manifolds, that at each step jump to a uniformly chosen (according to the Riemannian volume measure) point in the ball of radius h around the current position. A natural modification of such a random walk (based on the Metropolis algorithm) approximates Brownian motion on the manifold when h is sent to zero and time is rescaled appropriately. Moreover, the authors consider the transition operator of the random walk, and prove that its rescaled spectrum approximates the

spectrum of the Laplace-Beltrami operator. This provides a sharp rate of convergence of this random walk to its stationary distribution.

Horizontal Laplacians and diffusions on sub-Riemannian manifolds and random walk approximations to these have appeared several times in the literature recently. We make use of our notation and terminology in describing these results, in order to make the connection with the present paper clearer.

In [23], Lebeau and Michel study the spectral theory of a reversible Markov chain associated to a random walk on a sub-Riemannian manifold M , where, at each step, one walks along the integral lines (not geodesics, in general) of a fixed set of divergence-free vector fields X_1, \dots, X_k (with respect to some fixed volume ω). This random walk depends on a parameter h . In particular, they prove the convergence when $h \rightarrow 0$ to an associated hypoelliptic diffusion, which, being that the vector fields are divergence free, coincides with Δ_ω . In a similar spirit as above, they also consider the rate of convergence to equilibrium of a random walk of this type.

In [16], Gordina and Laetsch use a Riemannian extension of a sub-Riemannian metric to determine an orthogonal complement \mathbf{c}_g to \mathbf{A} , which is equivalent to determining a subspace of horizontal co-vectors. This allows them to define a horizontal Laplacian by averaging over second-derivatives in the horizontal directions. The result is, in our terminology, a microscopic Laplacian that, in fact, depends only on the choice of complement \mathbf{c}_g . They then introduce a corresponding family of horizontal random flights (in our terminology, though they call them random walks—see the discussion in Appendix 11), given by choosing a horizontal co-vector of fixed length uniformly at random, and then following the resulting geodesic at a constant speed for an exponentially distributed length of time, before repeating the procedure. They show that, under the natural parabolic scaling of the length of the co-vectors and the mean of the exponential “travel times,” these random flights converge to the diffusion generated by the horizontal Laplacian. To do this, they use resolvent formalism in order to prove the convergence of the relevant semi-groups, in a similar vein to the paper of Pinsky [26] mentioned above.

In [15], Gordina and Laetsch give a more systematic discussion of horizontal Laplacians relative to a choice of vertical complement. In particular, they take up the question of when a horizontal Laplacian relative to a choice of complement is equal to the divergence of the horizontal gradient with respect to some volume (or, as we have framed the question here, when the macroscopic Laplacian Δ_ω is equal to the microscopic Laplacian L^c). They give (see [15, Theorem 5.13]) a condition for this that is equivalent to our Theorem 57, though stated in terms of some Riemannian extension of the sub-Riemannian metric. Having given this result, they consider the application to three concrete examples of 3-dimensional Lie groups, the Heisenberg group, $SU(2)$, and the affine group.

In [17], Grong and Thalmaier define the horizontal Laplacian, corresponding to what we have called the microscopic Laplacian, by extending the sub-Riemannian metric to a Riemannian metric, considering an associated connection, and then taking the trace of a projection of the Hessian (see [17, Def. 2.1]). The resulting operator depends only on the orthogonal complement \mathbf{c}_g to \mathbf{A} with respect to the Riemannian extension \mathbf{g} . Random walks are not considered in this approach, although it also gives a construction of the horizontal Laplacian associated to a choice of horizontal subspace. They do consider the question of when their horizontal Laplacian (with respect to \mathbf{c}_g) is equal to the Laplacian determined by the Riemannian volume \mathcal{R}_g of the extension metric, and the result is essentially our Theorem 57 applied to this situation, presented in terms of the extension metric \mathbf{g} (and with a similar proof). Nonetheless, the primary interest of [17] is curvature-dimension inequalities for sub-Riemannian structures associated to certain Riemannian foliations, in the spirit of Baudoin and Garofalo [9], not in a discussion of sub-Laplacians per se. Indeed, they go on to select, among all possible complements (and then associated diffusions) special ones, that are integrable and “metric-preserving” in a way suitable for their purpose.

Finally, we note that, while this paper is primarily written from the point of view that sub-Riemannian manifolds are a generalization of Riemannian manifolds, they also arise in the theory of rough paths. In this direction, Breuillard, Friz, and Huesmann [11] have proven a Donsker-type theorem showing the convergence of Euclidean random walks to Euclidean Brownian motion in an appropriate rough path topology, which corresponds (in the spirit of this paper) to the convergence of certain random walks on Carnot groups.

In light of the above, the novelty of the present paper lies primarily in the following areas. We abandon the Riemannian formalism, and we use frame formalism, which seems better suited to the sub-Riemannian context (for instance, we get operators that depend on the choice of complement \mathbf{c} by construction). Also, having given the condition for the equivalence of macroscopic and microscopic Laplacians (“the equivalence problem”) in frame formalism (again, Theorem 57), we proceed to analyze many broad classes of examples in some detail (for example, any Carnot group and any corank-1 structure). In the contact case, we solve completely the equivalence problem. We do the same in the quasi-contact case, for the Popp volume $\omega = \mathcal{P}$. This leads to the perhaps-surprising fact that there are 4-dimensional quasi-contact structures for which the macroscopic Laplacian with respect to the Popp volume, which is in a precise sense the canonical volume, cannot be realized as a microscopic Laplacian (indeed, this is generically the case, though the complete proof is beyond the scope of the present paper). More generally, we discuss to what extent volume measures on a sub-Riemannian manifold can be thought of as canonical, which has implications for the degree to which a macroscopic Laplacian can be thought of as canonical.

On the probabilistic side, we allow random walks where the choice of co-vector is supported on the entire unit cylinder in the co-tangent space, rather than just on the horizontal subspace (relative to some choice of complement). (This really just says that sub-Riemannian geometry is rather insensitive to adding some independent vertical component to the initial co-vectors in a random walk, which is not surprising.) Also, our Theorem 70 gives a general convergence result for random walks on (sub)-Riemannian manifolds, going beyond the particular type of random walks considered elsewhere in the paper. From a technical perspective, this convergence of random walks is proved using martingale methods, rather than the semi-group approach mentioned above.

2 Preliminaries

We discuss some preliminaries in sub-Riemannian geometry. We essentially follow [3], but see also [25, 27, 21].

Definition 14. A *sub-Riemannian manifold* is a triple $(M, \mathbf{\blacktriangle}, \mathbf{g})$ where:

- M is smooth, connected manifold;
- $\mathbf{\blacktriangle} \subset TM$ is a smooth distribution of constant rank $k < n$, satisfying *Hörmander's condition*.
- \mathbf{g} is a smooth scalar product on $\mathbf{\blacktriangle}$: for all $q \in M$, \mathbf{g}_q is a positive definite quadratic form on $\mathbf{\blacktriangle}_q$, smooth as a function of q .

This definition does not include Riemannian structures, for which $k = n$ and $\mathbf{\blacktriangle} = TM$. We use the term *(sub)-Riemannian* to refer to structures $(M, \mathbf{\blacktriangle}, \mathbf{g})$ that are either Riemannian or sub-Riemannian.

Definition 15. Define $\mathbf{\blacktriangle}^1 := \mathbf{\blacktriangle}$, $\mathbf{\blacktriangle}^{i+1} := \mathbf{\blacktriangle}^i + [\mathbf{\blacktriangle}^i, \mathbf{\blacktriangle}]$, for every $i \geq 1$. A sub-Riemannian manifold is said to be *equivregular* if for each $i \geq 1$, the dimension of $\mathbf{\blacktriangle}_q^i$ does not depend on the point $q \in M$.

Smooth sections of $\mathbf{\blacktriangle}$ are called *horizontal vector fields*. Hörmander's condition guarantees that any two points in M can be joined by a Lipschitz continuous curve whose velocity is a.e. in $\mathbf{\blacktriangle}$ (Chow-Rashevskii theorem). We call such curves *horizontal*. Horizontal curves $\gamma : I \rightarrow M$ have a well-defined length, given by

$$\ell(\gamma) = \int_I \|\dot{\gamma}(t)\| dt,$$

where $\|\cdot\|$ is the norm induced by the (sub)-Riemannian scalar product. The *Carnot-Carathéodory distance* between two points $p, q \in M$ is

$$d(p, q) = \inf\{\ell(\gamma) \mid \gamma \text{ horizontal curve connecting } q \text{ with } p\}.$$

This distance turns $(M, \mathbf{\blacktriangle}, \mathbf{g})$ into a metric space that has the same topology of M .

Definition 16. The *Hamiltonian function* $H : T^*M \rightarrow \mathbb{R}$ associated with a sub-Riemannian structure is

$$H(\lambda) := \frac{1}{2} \sum_{i=1}^k \langle \lambda, X_i \rangle^2,$$

for any choice of a local orthonormal frame X_1, \dots, X_k of horizontal fields, i.e. $\mathbf{g}(X_i, X_j) = \delta_{ij}$.

On each fiber, H is non-negative quadratic form, and provides a way to measure the “length” of covectors. By setting $H_q := H|_{T_q^*M}$, one can check that $\ker H_q = \mathbf{\blacktriangle}_q^\perp$, the set of covectors that vanish on the distribution. In the Riemannian case, H is precisely the fiber-wise inverse of the metric \mathbf{g} .

Let σ the natural symplectic structure on T^*M , and π the projection $\pi : T^*M \rightarrow M$. The *Hamiltonian vector field* \vec{H} is the unique vector field on T^*M such that $dH = \sigma(\cdot, \vec{H})$. Integral lines of \vec{H} are indeed smooth curves on the cotangent bundle that satisfy Hamilton's equations $\dot{\lambda}(t) = \vec{H}(\lambda(t))$.

Definition 17. The projections $\gamma(t) = \pi(\lambda(t))$ of integral lines of \vec{H} are called *normal (sub)-Riemannian geodesics*.

Normal geodesics are indeed smooth and, as in the Riemannian case, are locally minimizing (i.e. any sufficiently small segment of $\gamma(t)$ minimizes the distance between its endpoints). For any $\lambda \in T^*M$ we consider the associated normal geodesic $\gamma_\lambda(t)$, obtained as the projection of the integral line $\lambda(t)$ of \vec{H} with initial condition $\lambda(0) = \lambda$. The *initial covector* λ plays the same role in sub-Riemannian geometry of the initial vector of Riemannian geodesics, with the important difference that an infinite number of distinct sub-Riemannian normal geodesics, with different initial covector, have the same initial vector. Notice that the Hamiltonian function, which is constant on integral lines $\lambda(t)$, measures the speed of the normal geodesic:

$$2H(\lambda) = \|\dot{\gamma}_\lambda(t)\|^2, \quad \lambda \in T^*M.$$

We are ready for the following important definition.

Definition 18. Let $D_q \subseteq [0, \infty) \times T_q^*M$ the set of the pairs (t, λ) such that the normal geodesic with initial covector λ is well defined up to time t . The (sub)-Riemannian *exponential map* (at $q \in M$) is the map $\exp_q : D_q \rightarrow M$ that associates with (t, λ) the point $\gamma_\lambda(t)$.

When clear, we suppress the initial point q in \exp_q . It is easy to show that, for any $\alpha > 0$, we have

$$\gamma_{\alpha\lambda}(t) = \gamma_\lambda(\alpha t).$$

This rescaling property, due to the fact that H is fiber-wise homogeneous of degree 2, justifies the restriction to the subset of initial covectors lying in the level set $2H = 1$.

Definition 19. The *unit cotangent bundle* is the set of initial covectors such that the associated normal geodesic has unit speed, namely

$$\square := \{\lambda \in T^*M \mid 2H(\lambda) = 1\} \subset T^*M.$$

Remark 7. We stress that, in the sub-Riemannian case, the non-negative quadratic form $H_q = H|_{T_q^*M}$ has non-trivial kernel. It follows that the fibers \square_q are cylinders and thus, non-compact, in sharp contrast with the Riemannian case (where the fibers \square_q are spheres).

For any $\lambda \in \square$, the corresponding geodesic $\gamma_\lambda(t)$ is parametrized by arc-length, namely $\ell(\gamma|_{[0, T]}) = T$. Even if \square is not compact, all arc-length parametrized geodesics are well defined for a sufficiently small time. The next Lemma is a consequence of the form of Hamilton's equations and the compactness of small balls.

Lemma 20. *There exists $\varepsilon > 0$ such that $[0, \varepsilon) \times \square_q \subseteq D_q$. In other words all arc-length parametrized normal geodesics $\gamma_\lambda(t)$ are well defined on the interval $[0, \varepsilon)$.*

In the Riemannian case, the gradient of a function is a vector field that points in the direction of the greatest rate of increase of the function. The generalization to the (sub)-Riemannian setting is straightforward.

Definition 21. Let $f \in C^\infty(M)$. The *horizontal gradient* $\text{grad}(f) \in \Gamma(\blacktriangle)$ is defined by

$$df(X) = \mathbf{g}(\text{grad}(f), X), \quad \forall X \in \Gamma(\blacktriangle).$$

Since, in the Riemannian case, it is the usual gradient, this notation will cause no confusion.

2.1 Computations with frames

If E is a smooth vector bundle over M , the symbol $\Gamma(E)$ denotes the $C^\infty(M)$ -module of smooth sections of E . Horizontal vector fields are then elements of $\Gamma(\blacktriangle)$.

In sub-Riemannian geometry, computations are most effectively done in terms of orthonormal frames (Riemannian normal coordinates are not available in general). Then, let X_1, \dots, X_k a (local) orthonormal frame for the sub-Riemannian structure. Moreover, consider some complement X_{k+1}, \dots, X_n , namely a local frame that completes X_1, \dots, X_k to a local frame for TM . Let $c_{ij}^\ell \in C^\infty(M)$ (the *structural functions*) be defined by:

$$[X_i, X_j] = \sum_{\ell=1}^n c_{ij}^\ell X_\ell, \quad i, j = 1, \dots, n.$$

Now define the functions $h_i : T^*M \rightarrow \mathbb{R}$ (that are linear on fibers) as:

$$h_i(\lambda) := \langle \lambda, X_i \rangle, \quad i = 1, \dots, n.$$

We have

$$H = \frac{1}{2} \sum_{i=1}^k h_i^2, \quad \vec{H} = \sum_{i=1}^k h_i \vec{h}_i,$$

where \vec{h}_i is the Hamiltonian vector field associated with h_i , namely $\sigma(\cdot, \vec{h}_i) = dh_i$. Indeed, for any fixed q , the restriction of $(h_1, \dots, h_n) : T_q^*M \rightarrow \mathbb{R}^n$ gives coordinates to T_q^*M , associated with the choice of X_1, \dots, X_n . In terms of these coordinates the fibers of the unit cotangent bundle are

$$\square_q := T_q^*M \cap \square = \{(h_1, \dots, h_k, h_{k+1}, \dots, h_n) \mid h_1^2 + \dots + h_k^2 = 1\} \simeq \mathbb{S}^{k-1} \times \mathbb{R}^{n-k}.$$

The last identification depends on the choice of the frame X_1, \dots, X_n . Finally, for any normal geodesic $\gamma_\lambda(t)$

$$\dot{\gamma}_\lambda(t) = \pi_* \dot{\lambda}(t) = \pi_* \vec{H}(\lambda(t)) = \sum_{i=1}^k h_i(\lambda(t)) X_i,$$

where we used the fact that $\pi_* \vec{h}_i = X_i$.

2.2 A Taylor expansion with frames

In this section we prove a Taylor expansion for smooth functions along normal geodesics. This formula will play a key role in the following. To keep a compact notation we often employ the *generalized Einstein's convention*: repeated indices are summed, albeit on different ranges: Greek indices $(\alpha, \beta, \gamma, \dots)$ from 1 to n ; Latin indices (i, j, ℓ, \dots) from 1 to k ; barred Latin indices $(\bar{i}, \bar{j}, \bar{\ell}, \dots)$ from $k+1$ to n .

Lemma 22. *Let $\phi \in C^\infty(M)$ and consider the geodesic $\gamma_\lambda(t)$, emanating from q . Then*

$$\phi(\gamma_\lambda(t)) = \phi(q) + th_i X_i(\phi)(q) + \frac{1}{2}t^2 [h_j c_{ji}^\alpha h_\alpha X_i(\phi)(q) + h_i h_j X_j(X_i(\phi))(q)] + t^3 r_\lambda(t),$$

where the initial covector $\lambda \in T_q^*M$ has coordinates (h_1, \dots, h_n) and $r_\lambda(t)$ is a remainder term. Moreover for any $q \in M$, there exists an $\varepsilon_0 > 0$ and constants $A, B_{\bar{\ell}}, C_{\bar{i}\bar{j}} \geq 0$ such that, for any $\lambda \in \square$ with $\pi(\lambda) \in B(q, \varepsilon_0)$ (the metric ball with center q and radius ε_0) we have

$$|r_\lambda(t)| \leq A + B_{\bar{\ell}}|h_{\bar{\ell}}| + C_{\bar{i}\bar{j}}|h_{\bar{i}}||h_{\bar{j}}|, \quad \forall t \leq \varepsilon_0.$$

Remark 8. In coordinates $\lambda = (h_1, \dots, h_n)$ with $h_1^2 + \dots + h_k^2 = 1$. Thus, the estimate above shows how the remainder term depends on the ‘‘unbounded’’ coordinates h_{k+1}, \dots, h_n of the initial covector.

Proof. The geodesic $\gamma_\lambda(t)$ is the projection of the integral curve $\lambda(t)$. Its initial covector $\lambda \in T_q^*M$ has coordinates $(h_1, \dots, h_n) \in \mathbb{R}^n$. We drop the subscript λ , since it's fixed. As we described above,

$$\frac{d}{dt}\phi(\gamma(t)) = \dot{\gamma}(t)(\phi) = h_i(t)X_i(\phi)(\gamma(t)), \quad (7)$$

where $h_i(t)$ is a shorthand for $h_i(\lambda(t))$. Similarly, for any function $g \in C^\infty(T^*M)$, we set $g(t) = g(\lambda(t))$. With this notation, Hamilton's equations are:

$$\dot{g}(t) = \{H, g\}(t),$$

where $\{\cdot, \cdot\}$ denotes the Poisson bracket, and the dot the derivative w.r.t. t . For $h_i : T^*M \rightarrow \mathbb{R}$ we get

$$\dot{h}_i = \{H, h_i\} = h_j \{h_j, h_i\} = h_j c_{ji}^\alpha h_\alpha,$$

where we suppressed the explicit evaluation at t . We apply the chain rule to Eq. (7) and we get

$$\frac{d^2}{dt^2}\phi(\gamma(t)) = \dot{h}_i X_i(\phi) + h_i h_j X_j(X_i(\phi)) = h_j c_{ji}^\alpha h_\alpha X_i(\phi) + h_i h_j X_j(X_i(\phi)). \quad (8)$$

Evaluating at $t = 0$ we get the second order term. Now let $t \leq \varepsilon_0$. The remainder of Taylor's expansion in Lagrange's form is:

$$r_\lambda(t) = \frac{1}{3!} \left. \frac{d^3}{dt^3} \right|_{t=t_*} \phi(\gamma(t)), \quad t_* \in [0, t].$$

To compute it we apply the chain rule to Eq. (8). By Hamilton's equations $\dot{h}_\alpha = h_i c_{i\alpha}^\beta h_\beta$, we reduce it to a polynomial in $h_1(t), \dots, h_n(t)$, structural functions $c_{\alpha\beta}^\gamma$ and their first derivatives $X_i(c_{\alpha\beta}^\gamma)$:

$$\begin{aligned} \frac{d^3}{dt^3}\phi(\gamma(t)) &= h_\ell c_{\ell j}^\beta h_j c_{ji}^\alpha h_\alpha X_i(\phi) + h_j h_\ell X_\ell(c_{ji}^\alpha) h_\alpha X_i(\phi) + h_j c_{ji}^\alpha h_\ell c_{\ell\alpha}^\beta h_\beta X_i(\phi) + h_j c_{ji}^\alpha h_\alpha h_\ell X_\ell(X_i(\phi)) + \\ &+ h_\ell c_{\ell i}^\alpha h_\alpha h_j X_j(X_i(\phi)) + h_i h_\ell c_{\ell j}^\alpha h_\alpha X_j(X_i(\phi)) + h_i h_j h_\ell X_\ell(X_j(X_i(\phi))). \end{aligned}$$

We stress that everything on the r.h.s. is computed at t . Since $\lambda \in \square_q$, and $2H$ is a constant of the motion, each $|h_i(t)| \leq 1$ for $i = 1, \dots, k$. Fix $q \in M$, and consider the sub-Riemannian closed ball $\overline{B(q, 2\varepsilon_0)}$, that is compact for sufficiently small ε_0 . If $\lambda \in \square \cap \pi^{-1}(B(q, \varepsilon_0))$, the length parametrized geodesic γ_λ does not exit the compact $\overline{B(q, 2\varepsilon_0)}$. Therefore also the structural functions and their derivatives are bounded by their maximum on $\overline{B(q, 2\varepsilon_0)}$. The only a priori uncontrolled terms are then the $h_{\bar{i}}(t)$, for $\bar{i} = k+1, \dots, n$. Thus:

$$|r_\lambda(t)| \leq P + Q_{\bar{\ell}}|h_{\bar{\ell}}(t)| + R_{\bar{i}\bar{j}}|h_{\bar{i}}(t)||h_{\bar{j}}(t)|, \quad \forall t \leq \varepsilon_0, \quad (9)$$

where $P, Q_{\bar{\ell}}, R_{\bar{i}\bar{j}} \geq 0$ are constants. Gronwall's Lemma shows that, for $t \leq \varepsilon_0$ we have

$$|h_{\bar{i}}(t)| \leq D_{\bar{i}} + E_{\bar{i}\bar{j}}|h_{\bar{j}}(0)|, \quad \bar{i} = k+1, \dots, n,$$

for constants $D_{\bar{i}}, E_{\bar{i}\bar{j}} \geq 0$. By plugging this last equation into Eq. (9), we obtain the result. \square

3 On volume forms in (sub)-Riemannian geometry

The aim of this section is to introduce a rigorous definition of N-intrinsic volume as the unique volume that, roughly speaking “depends only on the first order approximation of the (sub)-Riemannian structure” (see the discussion in Section 1).

Let M be an orientable n -dimensional smooth manifold. A (smooth) *volume form* is a positive n -form ω , that is $\omega(X_1, \dots, X_n) > 0$ for any oriented local frame X_1, \dots, X_n . Any volume form ω defines a positive measure on Borel sets of M , that we still call ω . Let $f \in C^\infty(M)$, and U a Borel set. In this way, $\int_U f\omega$ denotes the integral over U of the function f with respect to the measure induced by ω .

Definition 23. Given an m -form ω , and a vector field $X \in \Gamma(TM)$, we define the *Lie derivative* $\mathcal{L}_X\omega$ as the m -form such that its value at $q \in M$ is:

$$(\mathcal{L}_X\omega)_q = \left. \frac{d}{dt} \right|_{t=0} (P_t^*\omega)_q,$$

where P_t is the flow of X .

Definition 24. Let $X \in \Gamma(TM)$ and ω a volume form. The *divergence* $\text{div}_\omega(X)$ is the function defined by

$$\mathcal{L}_X\omega = \text{div}_\omega(X)\omega.$$

The next two lemmas are easy consequences of the definition.

Lemma 25. Let $f \in C_0^\infty(M)$ and ω a volume form. Then

$$\int_M f \text{div}_\omega(X)\omega = - \int_M df(X)\omega.$$

Lemma 26. Let $f \in C^\infty(M)$, with $f \neq 0$ and ω a volume form. Let $\omega' = f\omega$. Then

$$\text{div}_{\omega'}(X) = \text{div}_\omega(X) + X(\log |f|).$$

In particular, the definition of div_ω does not depend on the choice of orientation.

3.1 The nilpotent approximation

A key ingredient is the definition of *nilpotent approximation*: a (sub)-Riemannian structure built on the tangent space at a point T_qM which, in a sense, is the “first order approximation of the (sub)-Riemannian structure”. Let \blacktriangle be an equiregular, bracket-generating distribution. The *step* of the distribution is the first m such that $\blacktriangle_q^m = T_qM$.

Definition 27. The *nilpotentization* of \blacktriangle at a point $q \in M$ is the graded vector space

$$\widehat{M}_q := \blacktriangle_q \oplus \blacktriangle_q^2 / \blacktriangle_q \oplus \dots \oplus \blacktriangle_q^m / \blacktriangle_q^{m-1}.$$

The vector space \widehat{M}_q can be endowed with a Lie algebra structure, which respects the grading, induced by the Lie bracket, as follows. Let $X_q \in \blacktriangle_q^i / \blacktriangle_q^{i-1}$ and $Y_q \in \blacktriangle_q^j / \blacktriangle_q^{j-1}$. Let $X, Y \in \Gamma(TM)$ smooth extensions of any representative of X_q, Y_q . Then the Lie product between X_q and Y_q is

$$[X_q, Y_q] := [X, Y]_q \quad \text{mod } \blacktriangle_q^{i+j-1} \in \blacktriangle_q^{i+j} / \blacktriangle_q^{i+j-1}.$$

Under the equiregularity assumption, this product is well defined and does not depend on the choice of the representatives and of the extensions. Observe that the graded Lie algebra \widehat{M}_q is nilpotent. Then there is a unique connected, simply connected group, such that its Lie algebra is \widehat{M}_q (that we identify with the group itself). The global, left-invariant vector fields obtained by the group action on any orthonormal basis of $\blacktriangle_q \subset \widehat{M}_q$ give \widehat{M}_q the structure of a left-invariant (sub)-Riemannian manifold, which is called the *nilpotent approximation* of the sub-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ at the point q . We stress that:

- the base manifold is the vector space $\widehat{M}_q \simeq T_qM$ (the latter identification is *not* canonical);
- the distribution and the scalar product are obtained by extending $\blacktriangle_q, \mathbf{g}_q$ using the left action of the group.

It can be proved that \widehat{M}_q is isometric (as a metric space) to the Gromov tangent cone at q of the metric space (M, d) where d is the Carnot-Carathéodory distance induced by the (sub)-Riemannian structure (see [25, 10]).

3.2 Popp's volume

In this section we provide the definition of Popp's volume. Our presentation follows closely the one of [8, 25]. The definition rests on the following lemmas.

Lemma 28. *Let E be an inner product space, and let $\pi : E \rightarrow V$ be a surjective linear map. Then π induces an inner product on V such that the length of $v \in V$ is*

$$\|v\|_V = \min\{\|e\|_E \text{ s.t. } \pi(e) = v\}.$$

Lemma 29. *Let E be a vector space of dimension n with a flag of linear subspaces $\{0\} = F^0 \subset F^1 \subset F^2 \subset \dots \subset F^m = E$. Let $\mathbf{gr}(F) = F^1 \oplus F^2/F^1 \oplus \dots \oplus F^m/F^{m-1}$ be the associated graded vector space. Then there is a canonical isomorphism $\theta : \wedge^n E \rightarrow \wedge^n \mathbf{gr}(F)$.*

The proofs can be found in [8]. We report here the proof of Lemma 29 since it contains the definition of the important map θ .

Proof. We only give a sketch of the proof. For $0 \leq i \leq m$, let $k_i := \dim F^i$. Let X_1, \dots, X_n be a adapted basis for E , i.e. X_1, \dots, X_{k_i} is a basis for F^i . We define the linear map $\widehat{\theta} : E \rightarrow \mathbf{gr}(F)$ which, for $0 \leq j \leq m-1$, takes $X_{k_j+1}, \dots, X_{k_{j+1}}$ to the corresponding equivalence class in F^{j+1}/F^j . This map is indeed a non-canonical isomorphism, which depends on the choice of the adapted basis. In turn, $\widehat{\theta}$ induces a map $\theta : \wedge^n E \rightarrow \wedge^n \mathbf{gr}(F)$, which sends $X_1 \wedge \dots \wedge X_n$ to $\widehat{\theta}(X_1) \wedge \dots \wedge \widehat{\theta}(X_n)$. The proof that $\widehat{\theta}$ does not depend on the choice of the adapted basis is a straightforward check, and boils down to the fact that two different adapted basis are related by an upper triangular matrix (see [8, 25]). \square

The idea behind Popp's volume is to define an inner product on each $\blacktriangle_q^i / \blacktriangle_q^{i-1}$ which, in turn, induces an inner product on the orthogonal direct sum \widehat{M}_q . The latter has a natural volume form, such that its value on any oriented orthonormal basis is 1. Then, we employ Lemma 29 to define an element of $(\wedge^n T_q M)^*$, which is Popp's volume form computed at q .

Fix $q \in M$. Then, let $v, w \in \blacktriangle_q$, and let V, W be any horizontal extensions of v, w . Namely, $V, W \in \Gamma(\blacktriangle)$ and $V(q) = v, W(q) = w$. The linear map $\pi : \blacktriangle_q \otimes \blacktriangle_q \rightarrow \blacktriangle_q^2 / \blacktriangle_q$

$$\pi(v \otimes w) := [V, W]_q \quad \text{mod } \blacktriangle_q,$$

is well defined, and does not depend on the choice the horizontal extensions. Similarly, let $1 \leq i \leq m$. The linear maps $\pi_i : \otimes^i \blacktriangle_q \rightarrow \blacktriangle_q^i / \blacktriangle_q^{i-1}$

$$\pi_i(v_1 \otimes \dots \otimes v_i) = [V_1, [V_2, \dots, [V_{i-1}, V_i]]]_q \quad \text{mod } \blacktriangle_q^{i-1}, \quad (10)$$

are well defined and do not depend on the choice of the horizontal extensions V_1, \dots, V_i of v_1, \dots, v_i .

By the bracket-generating condition, π_i are surjective and, by Lemma 28, they induce an inner product space structure on $\blacktriangle_q^i / \blacktriangle_q^{i-1}$. Therefore, the nilpotentization of the distribution at q , namely

$$\widehat{M}_q = \blacktriangle_q \oplus \blacktriangle_q^2 / \blacktriangle_q \oplus \dots \oplus \blacktriangle_q^m / \blacktriangle_q^{m-1},$$

is an inner product space, as the orthogonal direct sum of a finite number of inner product spaces. As such, it is endowed with a canonical volume (defined up to a sign) $\omega_q \in (\wedge^n \widehat{M}_q)^*$ such that $\omega_q(v_1, \dots, v_n) = 1$ for any orthonormal basis v_1, \dots, v_n of the vector space \widehat{M}_q .

Finally, Popp's volume (computed at the point q) is obtained by transporting the volume of \widehat{M}_q to $T_q M$ through the map $\theta_q : \wedge^n T_q M \rightarrow \wedge^n \widehat{M}_q$ defined in Lemma 29. Namely

$$\mathcal{P}_q = \theta_q^*(\omega_q) = \omega_q \circ \theta_q, \quad (11)$$

where θ_q^* denotes the dual map. Eq. (11) is defined only in the domain of the chosen local frame. Since M is orientable, with a standard argument, these local n -forms can be glued to a global one, called Popp's volume \mathcal{P} . Moreover, Popp's volume is smooth, as follows, for example, by the explicit formula in [8, Theorem 1]. It is also clear that, in the Riemannian case, Popp's volume is the standard Riemannian one. Notice that Popp's volume is well defined only for equiregular structures.

3.2.1 Behavior under isometries

In the Riemannian setting, an isometry is a diffeomorphism such that its differential preserves the Riemannian scalar product. The concept is easily generalized to the sub-Riemannian case.

Definition 30. Let $(M, \blacktriangle_M, \mathbf{g}_M)$ and $(N, \blacktriangle_N, \mathbf{g}_N)$ be two (sub)-Riemannian structures. A (local) diffeomorphism $\phi : M \rightarrow N$ is a (local) *isometry* if its differential $\phi_* : TM \rightarrow TN$ preserves the (sub)-Riemannian structure, namely (i) $\phi_* \blacktriangle_M = \blacktriangle_N$, and (ii) $\phi^* \mathbf{g}_N = \mathbf{g}_M$.

The following is a trivial, but crucial property of Popp's volume that follows by its construction.

Proposition 31. (Sub)-Riemannian (local) isometries $\phi : M \rightarrow M$ preserve Popp's volume, namely $\phi^* \mathcal{P} = \mathcal{P}$.

3.3 Intrinsic volumes

An *intrinsic definition of volume* is some algorithm that associates with any (sub)-Riemannian structure, a volume form, and this association is preserved by isometries. Let us be more precise.

Definition 32. An *intrinsic definition of volume* is a map that associates, with any (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ a volume form ω_M on M such that if $\phi : M \rightarrow N$ is a (sub)-Riemannian isometry between $(M, \blacktriangle_M, \mathbf{g}_M)$ and $(N, \blacktriangle_N, \mathbf{g}_N)$, then $\phi^* \omega_N = \omega_M$.

In the following we avoid a this verbose terminology and we often talk about “(intrinsic) volume” to mean either the actual volume form ω_M or the map $M \mapsto \omega_M$.

Example 1. For any equiregular (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$, we can consider Popp's volume \mathcal{P}_M . As a consequence of Proposition 31, this is an intrinsic definition of volume. Restricted to Riemannian structures, this definition of volume is the classical, Riemannian one.

Some intrinsic volumes are more intrinsic than others. Since the nilpotent approximation is the first order approximation of the (sub)-Riemannian structure at a point q , particularly simple definition of volumes are those that, loosely speaking, depend only on the metric invariants of the nilpotent approximation.

We want to make the above statement more precise. To do this, we must go back to the very definition of nilpotent approximation. Recall that there is no canonical way to identify $T_q M$ with the nilpotentization \widehat{M}_q . Still, as an immediate consequence of Lemma 29 we can identify the top wedge product of these vector spaces.

Corollary 33. Let $(M, \blacktriangle, \mathbf{g})$ a (sub)-Riemannian manifold, $q \in M$ and \widehat{M}_q its nilpotent approximation at q . Then there exists a canonical map $\theta_q^* : (\wedge^n T_q M)^* \rightarrow (\wedge^n \widehat{M}_q)^*$.

Proof. This map is just the dual map of θ defined in the proof of Lemma 29, where $E = T_q M$ and $F^i = \blacktriangle^i$. \square

Remark 9. In other terms, at a point q , the map θ_q identifies canonically parallelotopes living in $T_q M$ and parallelotopes living in $\widehat{M}_q = T_0 \widehat{M}_q$. Notice that the (sub)-Riemannian metric does not play any role in the definition of θ_q .

Definition 34. An intrinsic definition of volume ω is *N-intrinsic* if, for any (sub)-Riemannian manifold $(M, \blacktriangle, \mathbf{g})$ and any $q \in M$ the following diagram commutes

$$\begin{array}{ccc} M & \xrightarrow{\omega} & \omega_M(q) \\ \text{nil}_q \downarrow & & \downarrow \theta_q^* \\ \widehat{M}_q & \xrightarrow{\omega} & \omega_{\widehat{M}_q}(0) \end{array}$$

where nil_q associates with any (sub)-Riemannian manifold its nilpotent approximation at q .

In other words, an intrinsic definition of volume is N-intrinsic if at any q the volume form $\omega_M(q)$ agrees with the nilpotent volume form $\omega_{\widehat{M}_q}(0)$ under the identification given by θ_q .

Example 2. Let $(M, \blacktriangle, \mathbf{g})$ be a Riemannian manifold ($\blacktriangle = TM$ and \mathbf{g} is a Riemannian metric). The Riemannian volume on M is the volume form ω_M such that, for any $q \in M$ and any oriented orthonormal parallelotope $v_1 \wedge \cdots \wedge v_n$ at q gives $\omega_M(q)(v_1 \wedge \cdots \wedge v_n) = 1$. We prove that this definition of volume is N-intrinsic. In this case, for any $q \in M$, the nilpotent approximation $\widehat{M}_q = T_q M$ with the scalar product given by \mathbf{g}_q . Then the map θ_q is just the identity map between $T_q M$ and \widehat{M}_q (that are indeed the same vector space). Thus

$$\theta_q(v_1 \wedge \cdots \wedge v_n) = v_1 \wedge \cdots \wedge v_n,$$

where the right hand side is seen as an element of $\wedge^n \widehat{M}_q = \wedge^n (T_0 \widehat{M}_q)$. Clearly $v_1 \wedge \cdots \wedge v_n$ is an orthonormal parallelotope either when seen as an element of $\wedge^n (T_q M)$ or as an element of $\wedge^n \widehat{M}_q$. By definition $\omega_M(q)(v_1 \wedge \cdots \wedge v_n) = 1$ and also $\omega_{\widehat{M}_q}(0)(\theta(v_1 \wedge \cdots \wedge v_n)) = 1$. Thus the Riemannian volume is N-intrinsic.

Example 3. Let (M, \blacktriangle, g) be a Riemannian manifold ($\blacktriangle = TM$ and g is a Riemannian metric). Let $\kappa : M \rightarrow \mathbb{R}$ be the scalar curvature function. We give a definition of volume ω as follows. For any $q \in M$ and any oriented normal parallelotope $v_1 \wedge \dots \wedge v_n$ at q :

$$\omega_M(q)(v_1 \wedge \dots \wedge v_n) = 1 + \kappa(q)^2.$$

One can check that this is an intrinsic definition of volume (due to the fact that κ is preserved by isometries). However, this is not N-intrinsic. In fact, without going into the details of the identification θ_q as above, we have that for any q , \widehat{M}_q is a flat Riemannian manifold (a vector space with inner product) thus $\widehat{\kappa} = 0$. Then

$$\omega_{\widehat{M}_q}(0)(\theta_q(v_1 \wedge \dots \wedge v_n)) = \omega_{\widehat{M}_q}(0)(v_1 \wedge \dots \wedge v_n) = 1 \neq 1 + \kappa(q)^2 = \omega_M(q)(v_1 \wedge \dots \wedge v_n).$$

Example 4. Popp's volume is N-intrinsic by construction. In fact it defined precisely by requiring the commutativity of the diagram above (see Section 3.2).

Example 5. The spherical Hausdorff volume is N-intrinsic. In fact, in [4], it is proved that the spherical Hausdorff volume is absolutely continuous w.r.t. any smooth volume form. In particular, the Radon-Nykodim derivative of the spherical Hausdorff volume w.r.t. Popp's volume is proportional to the left-invariant Haar measure of the unit ball in the nilpotent approximation. This function, in general, is not smooth, but only continuous.

We collect the results from these examples in the following propositions.

Proposition 35. *For any (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$, let \mathcal{P}_M be Popp's volume. Then the definition of volume $M \mapsto \mathcal{P}_M$ is N-intrinsic.*

Proposition 36. *For any Riemannian structure (M, \mathbf{g}) , let \mathcal{R}_M be the Riemannian volume. Let $f \in C^\infty(M)$ be any function invariant by isometries. Then the definition of volume (for Riemannian manifolds) $M \mapsto f\mathcal{R}_M$ is N-intrinsic if and only if f is constant.*

In other words, up to constant rescaling, the Riemannian one is the unique N-intrinsic definition of volume for Riemannian manifolds. This is due to the fact that the nilpotent approximation of Riemannian manifolds have no non-trivial metric invariants. In the sub-Riemannian case, this is not always true, due to the well-known presence of moduli for their nilpotent approximations. We close this section by discussing class of (sub)-Riemannian structures admitting a unique N-intrinsic definition of volume.

3.3.1 Equi-nilpotentizable structures

Equi-nilpotentizable (sub)-Riemannian structure have the same nilpotent approximation at any point.

Definition 37. A sub-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ is *equi-nilpotentizable* if for any $q, p \in M$ the nilpotent approximations \widehat{M}_q and \widehat{M}_p are isometric.

The next theorem is a generalization of Proposition 36 to the class of equi-nilpotentizable (sub)-Riemannian structures (that indeed contains the class of Riemannian structures).

Theorem 38. *Let ω and ω' two N-intrinsic definitions of volume. Then for any (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ there exists a smooth function $c_M : M \rightarrow \mathbb{R}$ such that $\omega_M = c_M \omega'_M$. The value of c_M at q depends only on the isometry class of the nilpotent approximation of the structure at q . In particular, for equi-nilpotentizable sub-Riemannian manifolds, any N-intrinsic definition of volume agrees up to a constant.*

Remark 10. By Proposition 35 Popp's one is an N-intrinsic definition of volume. Then, for equi-nilpotentizable (sub)-Riemannian structures, any N-intrinsic volume $M \mapsto \omega_M$ gives, up to a constant, Popp's one.

Proof. Assume that there are two N-intrinsic definitions of volume ω and ω' . Then, to any sub-Riemannian manifold M we associate a smooth, never vanishing function:

$$c_M := \omega_M / \omega'_M : M \rightarrow \mathbb{R}.$$

For any intrinsic definition of volume ω and isometry $\phi : M \rightarrow N$ of (sub)-Riemannian manifolds :

$$\phi^* \omega_N = \omega_M.$$

Then, if we have two intrinsic definition of volumes ω and ω' we get, for isometric structures:

$$c_M = \frac{\omega_M}{\omega'_M} = \frac{\phi^* \omega_N}{\phi^* \omega'_N} = c_N \circ \phi. \quad (12)$$

Lemma 39. *Let $(M, \blacktriangle, \mathbf{g})$ be a left-invariant structure (see Definition 42). Then c_M is constant and depends only on the isometry class of M . Namely, if $(M, \blacktriangle_M, \mathbf{g}_M)$ and $(N, \blacktriangle_N, \mathbf{g}_N)$ are two isometric left-invariant structures, then $c_M = c_N$.*

Proof. If in Eq. (12) we set $M = N$ and $\phi = L_q$, for $q \in M$, we get that c_M is a left-invariant function, and then constant. Moreover, if M and N are two isometric left-invariant structures, then the two constants c_M and c_N must be equal. \square

Let us go back to the structure $(M, \blacktriangle, \mathbf{g})$ and the two N-intrinsic definitions of volumes ω and ω' . By definition of N-intrinsic, we have:

$$c_M(q) = \frac{\omega_M(q)}{\omega'_M(q)} = \frac{\omega_{\widehat{M}_q}(0)}{\omega'_{\widehat{M}_q}(0)} = c_{\widehat{M}_q}(0), \quad \forall q \in M.$$

The sub-Riemannian structure on \widehat{M}_q is left-invariant. Then, by Lemma 39, the function $c_{\widehat{M}_q}$ depends only on the isometry class of \widehat{M}_q (that is the same for all $q \in M$). So c_M is a constant that depends only on the isometry class of \widehat{M}_q . \square

3.3.2 Homogeneous (sub)-Riemannian structures

Another class of structures that admit a unique N-intrinsic definition of volume (actually, a unique intrinsic definition of volume) is the following.

Definition 40. A (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ is *homogenous* if the group $\text{Iso}(M)$ of (sub)-Riemannian isometries of M acts transitively.

Indeed homogenous structures are equi-nilpotentizable, thus Theorem 38 applies and any two N-intrinsic definition of volumes are proportional (and proportional to Popp's one). Still, something stronger holds. In fact, for these structures, Popp's one is the unique volume form (up to scaling) preserved by (sub)-Riemannian isometries (see [8, Proposition 6]). Then we have the following corollary.

Corollary 41. *For homogeneous (sub)-Riemannian structures, any two intrinsic (definitions of) volumes are N-intrinsic, and proportional to Popp's one.*

In particular this is true for left-invariant (sub)-Riemannian structures on Lie groups.

Definition 42. Let M be a Lie group, and, for $q \in M$, let $L_q : M \rightarrow M$ denote the left multiplication. A (sub)-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ is *left-invariant* if $L_q : M \rightarrow M$ is an isometry for any $q \in M$.

Since Popp's volume on left-invariant structures is left-invariant (and then proportional to Haar's volume) we have the following.

Corollary 43. *For left-invariant (sub)-Riemannian structures, any two intrinsic (definitions of) volumes are N-intrinsic, and proportional to Haar's volume.*

4 The macroscopic Laplacian

Let $(M, \blacktriangle, \mathbf{g})$ a sub-Riemannian structure and fix a volume form ω (not necessarily intrinsic or N-intrinsic). In this setting, we have a well defined notion of gradient and divergence (see Definitions 21 and 24). The distribution is not assumed to be equiregular, unless one explicitly chooses $\omega = \mathcal{P}$.

Definition 44. The *macroscopic Laplacian* (w.r.t. ω) is the differential operator

$$\Delta_\omega = \text{div}_\omega \circ \text{grad}.$$

This is a second order differential operator, symmetric on $C_0^\infty(M)$ w.r.t. the L^2 product induced by the measure ω . A classical result by Strichartz (see [28, 29]) is the following:

Theorem 45. *If $(M, \blacktriangle, \mathbf{g})$ is complete as a metric space, then Δ_ω is essentially self-adjoint on $C_0^\infty(M)$ and the associated heat operator is given by a positive, smooth, symmetric heat kernel.*

We provide explicit formulas for Δ_ω , in terms of orthonormal frames. The proofs are routine computations.

Lemma 46. Let X_1, \dots, X_k be a (local) orthonormal frame for \blacktriangle and X_{k+1}, \dots, X_n some complement such that X_1, \dots, X_n is an oriented local frame for TM . Let $\theta \in C^\infty(M)$ defined by $\omega(X_1, \dots, X_n) = e^\theta$. Then

$$\begin{aligned} \text{grad}(f) &= \sum_{i=1}^k X_i(f)X_i, & f \in C^\infty(M), \\ \text{div}_\omega(X_i) &= \sum_{\alpha=1}^n c_{\alpha i}^\alpha + d\theta(X_i), & X \in \Gamma(TM). \end{aligned}$$

Proposition 47. The macroscopic Laplacian w.r.t. the choice of ω is

$$\Delta_\omega = \text{div}_\omega \circ \text{grad} = \sum_{i=1}^k X_i^2 + \text{div}_\omega(X_i)X_i = \sum_{i=1}^k X_i^2 + \sum_{i=1}^k \sum_{\alpha=1}^n c_{\alpha i}^\alpha X_i + \text{grad}(\theta). \quad (13)$$

where the horizontal gradient $\text{grad}(\theta)$ is seen as a derivation.

Notice that, for any choice of ω , the principal symbol of Δ_ω is the Hamiltonian function $2H : T^*M \rightarrow \mathbb{R}$.

Lemma 48 (Change of volume). Let ω be a volume form and $\omega' = e^g \omega$, with $g \in C^\infty(M)$. Then

$$\Delta_{\omega'} = \Delta_\omega + \text{grad}(g),$$

where $\text{grad}(g)$ is meant as a derivation.

Proof. It follows from the change of volume formula for the divergence. In fact, for $X \in \Gamma(TM)$,

$$\text{div}_{\omega'}(X)\omega' = \mathcal{L}_X \omega' = \mathcal{L}_X(e^g \omega) = X(g)e^g \omega + e^g \mathcal{L}_X \omega = (X(g) + \text{div}_\omega(X))\omega'.$$

Then $\text{div}_{\omega'}(X) = \text{div}_\omega(X) + X(g)$. The statement follows from the definition of $\Delta_\omega = \text{div}_\omega \circ \text{grad}$. \square

5 The microscopic Laplacian

Let \mathbf{c} be a complement (for the distribution \blacktriangle), namely a smooth sub-bundle of TM such that

$$T_q M = \blacktriangle_q \oplus \mathbf{c}_q, \quad \forall q \in M.$$

By duality we have an analogous splitting for the cotangent bundle, namely

$$T_q^* M = \mathbf{h}_q \oplus \mathbf{v}_q, \quad \forall q \in M,$$

where \mathbf{h} (resp. \mathbf{v}) denotes the annihilator bundle of \mathbf{c} (resp. of \blacktriangle). We call \mathbf{v} the space of *vertical covectors* and consequently \mathbf{h} is the space of *horizontal covectors* (see Fig. 1). The Hamiltonian function is a fiber-wise non-negative quadratic form $H_q := H|_{T_q^* M}$, with $\ker H_q = \mathbf{v}_q$. Then its restriction to \mathbf{h}_q is positive definite, thus is a scalar product on it.

Remark 11. An alternative definition, in the spirit of [25, 16, 15], is the following. Extend \mathbf{g} to a Riemannian metric such that \mathbf{c} and \blacktriangle are orthogonal. Then we have a dual co-metric on T^*M . This co-metric depends on the choice of the orthogonal extension, but its restriction on \mathbf{h} does not, and induces an Euclidean structure.

We obtain a $k - 1$ -dimensional sphere bundle over M by taking the fiber-wise intersection $\square \cap \mathbf{h}$. In fact

$$\square_q \cap \mathbf{h}_q = \{\lambda \in T_q^* M \mid 2H(\lambda) = 1, \lambda \in \mathbf{h}_q\}, \quad \forall q \in M,$$

is the Euclidean sphere \mathbb{S}_q^{k-1} on the vector space \mathbf{h}_q w.r.t. the scalar product $2H|_{\mathbf{h}_q}$.

5.1 A class of adapted measures

We define a class of probability measures on \square_q , obtained as the product of the standard probability measure on \mathbb{S}_q^{k-1} and a probability measure on the complement \mathbf{h}_q . We are interested in a sufficiently general model, where *all* geodesics on the cylinder potentially have a non-zero probability density. For this reason we include measures with non-compact support. This is indeed possible, but we need the probability to go to zero sufficiently fast at infinity along the “non compact directions” of $\square_q = \mathbb{S}_q^{k-1} \times \mathbf{v}_q$.

Definition 49. Let E be a vector space and $\alpha \in \mathbb{N}$. A Borel measure μ on E is α -*decreasing* if any linear function $f : E \rightarrow \mathbb{R}$ belongs to $L^\alpha(E, \mu)$.




Figure 1: Cotangent splitting $T_q^*M = \mathbf{v}_q \oplus \mathbf{h}_q$. Vertical covectors annihilate horizontal vectors ($\mathbf{v}_q = \mathbf{\blacktriangle}_q^\perp$), and correspond to the trivial geodesic. Horizontal covectors annihilate the transverse vectors ($\mathbf{h}_q = \mathbf{c}_q^\perp$). They correspond to normal geodesics that are the “most horizontal” w.r.t. the choice of \mathbf{c} .

If a measure is α -decreasing then it is β -decreasing for any $\beta \leq \alpha$. A compactly supported probability measure is α -decreasing for all α . Finally, we notice that one needs to check the condition only for a complete set of linear projections (the functions $(x_1, \dots, x_m) \mapsto x_m$ in terms of some basis).

Definition 50. Let μ_q be a probability measure on $\square_q = \mathbb{S}_q^{k-1} \times \mathbf{v}_q$. We say that the collection of measures $\{\mu_q\}_{q \in M}$ is *adapted to the splitting* if there exists a 2-decreasing probability measure $\mu_{\mathbf{v}_q}$ on \mathbf{v}_q such that

$$\mu_q = \mu_{\mathbb{S}_q^{k-1}} \times \mu_{\mathbf{v}_q}, \quad \forall q \in M,$$

where $\mu_{\mathbb{S}_q^{k-1}}$ denotes the standard uniform measure on $\mathbb{S}_q^{k-1} = \square_q \cap \mathbf{h}_q$. Moreover, we assume that if $f : T^*M \rightarrow \mathbb{R}$ is a continuous and fiberwise linear function, the map $q \mapsto \int_{\square_q} f^2 \mu_q$ (which is well defined) is continuous.

Remark 12. As we will see, the choice of $\alpha = 2$ is forced by the properties of the exponential map. In the Riemannian case, there is no need of this, since the fibers of the unit tangent bundle are compact. The regularity assumption is needed for the uniform convergence in the next definition.

For the next definition recall that, by Lemma 20, there exists $\varepsilon_0 > 0$ such that all arc-length parametrized geodesics $\gamma_\lambda(t) = \exp_q(t, \lambda)$, with $\lambda \in \square_q$, are well defined for $t \in [0, \varepsilon_0]$.

Definition 51. Consider a splitting $TM = \mathbf{\blacktriangle} \oplus \mathbf{c}$ and some choice of adapted measure $\{\mu_q\}_{q \in M}$. The *microscopic Laplacian* is the differential operator:

$$(L^c \phi)(q) := \lim_{t \rightarrow 0^+} \frac{2k}{t^2} \int_{\square_q} (\phi(\exp_q(t, \lambda)) - \phi(q)) \mu_q(\lambda), \quad \forall q \in M.$$

Surprisingly, this definition does not depend on the choice of the measure adapted to the splitting, and justifies the notation L^c . Moreover, the right hand side converges uniformly on compact sets for $t \rightarrow 0$. The proof of these facts are contained in the proof of Theorem 52.

Remark 13. As discussed in Section 1, this is the operator associated with the limit of a random walk, where, at each step, normal geodesics are chosen on \square_q with the given probability measure μ_q . Our construction generalizes the one given in [16, 15]. The latter can be recovered by setting

$$\mu_q = \mu_{\mathbb{S}_q^{k-1}} \times \delta_{\mathbf{v}_q}, \quad \forall q \in M,$$

where $\delta_{\mathbf{v}_q}$ is a Dirac mass centered at $0 \in \mathbf{v}$. This is certainly an α -decreasing measure $\forall \alpha \geq 0$, and satisfies the regularity assumptions of definition 50. With this choice, at each step of the associated random walk one moves only on a k -dimensional submanifold.

Theorem 52. *The microscopic sub-Laplacian L^c depends only on the choice of the complement \mathbf{c} , and the convergence of 51 is uniform on compact sets. Moreover, in terms of a local orthonormal frame X_1, \dots, X_k for $\mathbf{\blacktriangle}$ and a local frame X_{k+1}, \dots, X_n for \mathbf{c} , we have:*

$$L^c = \sum_{i=1}^k X_i^2 + \sum_{i,j=1}^k c_{ji}^j X_i, \quad (14)$$

where $c_{ij}^k \in C^\infty(M)$ are the structural functions associated with the given frame.

Proof. Let ν_1, \dots, ν_n the co-frame dual to X_1, \dots, X_n :

$$\begin{aligned} \mathbf{a} &= \text{span}\{X_1, \dots, X_k\}, & \mathbf{h} &= \text{span}\{\nu_1, \dots, \nu_k\}, \\ \mathbf{c} &= \text{span}\{X_{k+1}, \dots, X_n\}, & \mathbf{v} &= \text{span}\{\nu_{k+1}, \dots, \nu_n\}. \end{aligned}$$

Fix $q \in M$. In coordinates $(h_1, \dots, h_n) : T_q^*M \rightarrow \mathbb{R}^n$ induced by the choice of the frame, we have:

$$\square_q := \square \cap T_q^*M = \{(h_1, \dots, h_k, h_{k+1}, \dots, h_n) \mid h_1^2 + \dots + h_k^2 = 1\} \simeq \mathbb{S}^{k-1} \times \mathbb{R}^{n-k}.$$

Since q is fixed, we dub $\mathbf{v}_q = \mathbb{R}^{n-k}$ and $\mu_{\mathbf{v}_q} = \mu_{n-k}$. By plugging in the Taylor expansion of Lemma 22 in the definition of L^c we get:

$$(L^c \phi)(q) := \lim_{t \rightarrow 0^+} \frac{2k}{t^2} \int_{\square_q} \left\{ t h_i X_i(\phi) + \frac{1}{2} t^2 [h_j c_{ji}^\alpha h_\alpha X_i(\phi) + h_i h_j X_j(X_i(\phi))] + t^3 r_\lambda(t) \right\} \mu(\lambda),$$

where repeated indices are summed (according to the generalized Einstein's convention) and we suppressed some of the explicit evaluations at q . We compute the three terms of the integrand.

1. The linear in t term vanishes. In fact, the integrand depends only on the variables h_1, \dots, h_k , and

$$\int_{\square_q} h_i X_i(\phi) \mu(\lambda) = X_i(\phi) \int_{\mathbb{R}^{n-k}} \mu_{n-k} \int_{\mathbb{S}^{k-1}} h_i \mu_{\mathbb{S}^{k-1}} = 0,$$

since the integral of any linear function on the sphere vanishes.

2. For the quadratic in t term, we distinguish two contributions:

$$h_j c_{ji}^\alpha h_\alpha X_i(\phi) + h_i h_j X_j(X_i(\phi)) = \underbrace{h_j c_{ji}^{\bar{\ell}} h_{\bar{\ell}} X_i(\phi)}_1 + \underbrace{h_j c_{ji}^{\bar{\ell}} h_{\bar{\ell}} X_i(\phi) + h_i h_j X_j(X_i(\phi))}_2,$$

where, we recall that the range of summation of barred latin indices is from $k+1$ to n . The first term, not present in the Riemannian case, is the product of a linear function on \mathbb{S}^{k-1} and a linear function on \mathbb{R}^{n-k} . Since the measure μ is a product, by Fubini, we can first perform the integral on \mathbb{S}^{k-1} :

$$\int_{\square_q} h_j c_{ji}^{\bar{\ell}} h_{\bar{\ell}} X_i(\phi) = c_{ji}^{\bar{\ell}} X_i(\phi) \int_{\mathbb{S}^{k-1}} h_j \mu_{\mathbb{S}^{k-1}} \int_{\mathbb{R}^{n-k}} \mu_{n-k} = 0.$$

On the other hand, the second term does not contain any $h_{\bar{\ell}}$, for $\bar{\ell} = k+1, \dots, n$ and is the restriction on \mathbb{S}^{k-1} of a quadratic form $Q = Q_{ij} h_i h_j$, with:

$$Q_{ij} = X_i(X_j(\phi)) + c_{i\bar{\ell}}^j X_{\bar{\ell}}(\phi), \quad i, j = 1, \dots, k.$$

Recall that for any quadratic form $Q : \mathbb{R}^k \rightarrow \mathbb{R}$ we have

$$\int_{\mathbb{S}^{k-1}} Q(v) \mu_{\mathbb{S}^{k-1}}(v) = \frac{1}{k} \text{Tr}(Q).$$

Thus we get, for the quadratic in t term, the following expression:

$$k \int_{\square_q} Q_{ij} h_i h_j \mu = k \int_{\mathbb{R}^{n-k}} \mu_{n-k} \int_{\mathbb{S}^{k-1}} Q_{ij} h_i h_j \mu_{\mathbb{S}^{k-1}} = \text{Tr}(Q) = \sum_{i=1}^k X_i^2(\phi) + \sum_{i,\ell=1}^k c_{i\bar{\ell}}^i X_{\bar{\ell}}(\phi).$$

where we restored the explicit summations in the last term.

3. The final term to compute is the remainder. By Lemma 22, if t is sufficiently small, the remainder $r_\lambda(t)$ of the Taylor's expansion is uniformly bounded by a quadratic polynomial in the unbounded variables $h_{\bar{i}}$.

$$\left| \frac{k^2}{t} \int_{\square_q} r_\lambda(t) t^3 \right| \leq kt \int_{\mathbb{R}^{n-k}} (A + B_{\bar{\ell}} |h_{\bar{\ell}}| + C_{\bar{i}\bar{j}} |h_{\bar{i}}| |h_{\bar{j}}|) \mu_{n-k} \int_{\mathbb{S}^{k-1}} \mu_{\mathbb{S}^{k-1}}, \quad \forall t \leq \varepsilon_0.$$

By Lemma 22 this estimate holds uniformly on an open ball $B(q, \varepsilon_0)$. Then, for q in a compact set K :

$$\left| \frac{k^2}{t} \int_{\square_q} r_\lambda(t) t^3 \right| \leq tD \sum_{\bar{\ell}=k+1}^n \int_{\square_q} (h_{\bar{\ell}})^2 \mu_q = tg(q).$$

where D is a constant and, by our assumption on the family of measures, $g(q)$ is continuous (recall that the h_α are precisely the linear, smooth functions $\lambda \mapsto \langle \lambda, X_\alpha \rangle$ on T^*M). Then the remainder goes to zero uniformly on K . The fact that L^c depends only on the choice of the complement (and not on the adapted measure) is a consequence of Eq. (14). \square

Remark 14. We stress that equiregularity is not assumed in Theorem 52, for Lemma 22 is completely general.

5.2 Horizontal divergence

We can write the operator $L^{\mathbf{c}}$ in a more classical fashion, introducing the so-called *horizontal divergence*. As the classical divergence depends on the choice of a volume form (or a density), the horizontal one depends on the choice of the complement \mathbf{c} .

Definition 53. Let $m \in \mathbb{N}$. The bundle of *horizontal m -alternating tensors* on M is the disjoint union

$$\Lambda_{\blacktriangle}^m M := \bigsqcup_{q \in M} (\wedge^m \blacktriangle_q)^*.$$

where $(\wedge^m \blacktriangle_q)^*$ is the space of m -alternating functionals on \blacktriangle_q . Sections of $\Lambda_{\blacktriangle}^m M$ are called *horizontal m -forms*.

Any m -form induces an horizontal m -form simply by restriction. On the other hand, given the complement \mathbf{c} , the fiber-wise linear projection $\pi_{\blacktriangle} : TM \rightarrow \blacktriangle$ is well defined, and for any horizontal m -form η , we can define a m -form $\tilde{\eta} := \eta \circ \pi_{\blacktriangle}$.

Definition 54. Given an horizontal m -form η , and a vector field $q \in \Gamma(TM)$, we define the *horizontal Lie derivative* $\mathcal{L}_q^{\mathbf{c}} \eta$ as the horizontal m -form such that its value at $q \in M$ is

$$(\mathcal{L}_q^{\mathbf{c}} \eta)_q = \left. \frac{d}{dt} \right|_{t=0} (P_t^* \eta \circ \pi_{\blacktriangle})_q,$$

where P_t is the flow of q .

This is the same definition of the divergence of m -forms, with the additional ingredient of π_{\blacktriangle} , needed since η acts only on m -tuples of horizontal vectors. From the analogue formula for the divergence, one obtains readily:

$$(\mathcal{L}_q^{\mathbf{c}} \eta)(Y_1, \dots, Y_m) := q(\eta(Y_1, \dots, Y_m)) - \sum_{i=1}^m \eta(Y_m, \dots, \pi_{\blacktriangle}[q, Y_i], \dots, Y_m), \quad (15)$$

where $Y_1, \dots, Y_m \in \Gamma(\blacktriangle)$. If \blacktriangle is oriented (as a vector bundle), with $\text{rank } \blacktriangle = k$, we define a canonical horizontal k -form η such that

$$\eta(X_1, \dots, X_k) = 1,$$

for any oriented orthonormal frame $X_1, \dots, X_k \in \Gamma(\blacktriangle)$.

Definition 55. Let $q \in \Gamma(TM)$, and let η be a non-zero k form. The *horizontal divergence* $\text{div}^{\mathbf{c}}(q)$ is the function defined by

$$\mathcal{L}_q^{\mathbf{c}} \eta = \text{div}^{\mathbf{c}}(q) \eta.$$

Remark 15. The definition does not depend on the choice of orientation. A general definition for non-orientable distributions can be given in a standard way with horizontal densities (projectivized non-vanishing k forms).

Indeed $\text{div}^{\mathbf{c}}(q)$ computes the infinitesimal change of volume, under the flow of q , of the projection on \blacktriangle of the standard parallelootope of \blacktriangle . As a direct consequence of Eq. (14) and Eq. (15), we obtain the following.

Proposition 56. Let $TM = \blacktriangle \oplus \mathbf{c}$. The *microscopic Laplacian* is

$$L^{\mathbf{c}} = \text{div}^{\mathbf{c}} \circ \text{grad} = \sum_{i=1}^k X_i^2 + \text{div}^{\mathbf{c}}(X_i) X_i,$$

where X_1, \dots, X_k is an orthonormal frame, grad is the horizontal gradient and $\text{div}^{\mathbf{c}}$ is the horizontal divergence.

This formula must be compared directly with Eq (13) for the macroscopic operator, where the horizontal divergence $\text{div}^{\mathbf{c}}$ is replaced by the divergence div_{ω} w.r.t. some volume form ω on M .

6 The equivalence problem

Fix a volume form ω and a complement \mathbf{c} . We consider both $L^{\mathbf{c}}$ and Δ_{ω} as second order differential operators on the space of compactly smooth supported functions $C_0^{\infty}(M)$, with the L^2 product:

$$(f_1, f_2)_{\omega} = \int_M f_1 f_2 \omega, \quad f_1, f_2 \in C_0^{\infty}(M).$$

The operator Δ_{ω} is symmetric by construction. What about $L^{\mathbf{c}}$? The principal symbol of both operators, as a function on the cotangent bundle, is twice the Hamiltonian $2H : T^*M \rightarrow \mathbb{R}$.

Theorem 57. Let \mathbf{c} be a complement and ω a volume form. Then $L^{\mathbf{c}}$ is symmetric w.r.t. the L^2 product induced by ω if and only if $L^{\mathbf{c}} = \Delta_\omega$ or, equivalently, if and only if

$$\chi^{(\mathbf{c}, \omega)} := L^{\mathbf{c}} - \Delta_\omega = \sum_{i=1}^k \sum_{j=k+1}^n c_{ji}^j X_i + \text{grad}(\theta) = 0,$$

where $\theta = \log |\omega(X_1, \dots, X_n)|$ and $c_{ij}^\ell \in C^\infty(M)$ are the structural functions associated with an orthonormal frame X_1, \dots, X_k for \blacktriangle and a frame X_{k+1}, \dots, X_n for \mathbf{c} .

Proof. To get the explicit formula for $\chi^{(\mathbf{c}, \omega)}$ compare Eq. (13) with Eq. (14). For what concerns the symmetry, suppose that $L^{\mathbf{c}}$ is symmetric w.r.t. some choice of ω . Since Δ_ω is symmetric, $\chi^{(\mathbf{c}, \omega)}$ is also symmetric, that is (suppressing some notation):

$$(\chi(f_1), f_2)_\omega = (f_1, \chi(f_2))_\omega, \quad \forall f_1, f_2 \in C_0^\infty(M).$$

If we choose $f_2 = 1$ on some domain larger than the support of f_1 this is equivalent to $\int_M \chi(f_1) \omega = 0$ for any $f_1 \in C_0^\infty$. In particular $\chi = 0$. The converse is clear. \square

The condition $\chi^{(\mathbf{c}, \omega)} = 0$ is equivalent to the following system of PDEs:

$$X_i(\theta) + \sum_{\bar{j}=k+1}^n c_{\bar{j}i}^{\bar{j}} = 0, \quad i = 1, \dots, k.$$

In particular $L^{\mathbf{c}}$ is symmetric w.r.t. some volume ω if and only if, for fixed \mathbf{c} , the above system admits a global solution θ . If this is the case, the associated volume is given by $\omega(X_1, \dots, X_n) = e^\theta$.

Remark 16. The compatibility condition $\chi^{(\mathbf{c}, \omega)} = 0$ is the same appearing in [15, Theorem 5.13], written in a different form. We call *compatible* the pairs (\mathbf{c}, ω) solving the compatibility condition.

Finally, if $L^{\mathbf{c}}$ is symmetric w.r.t. some choice of volume ω , the latter is unique (up to constant rescaling).

Lemma 58. If $L^{\mathbf{c}} = \Delta_\omega$, then $L^{\mathbf{c}} = \Delta_{\omega'}$ if and only if $\omega = c\omega'$, where c is a non-zero constant.

Proof. Let $\omega' = e^g \omega$ with $g \in C^\infty(M)$. Then, using the change of volume formula (see Lemma 48):

$$\chi^{(\mathbf{c}, \omega')} = \Delta_{\omega'} - L^{\mathbf{c}} = \Delta_\omega + \text{grad}(g) - L^{\mathbf{c}} = \chi^{(\mathbf{c}, \omega)} + \text{grad}(g).$$

Thus both pairs are compatible iff $\text{grad}(g) = 0$, that is (by the bracket generating condition) iff g is constant. \square

7 Carnot groups

A Carnot group G of step m is a simply connected Lie group whose Lie algebra of left-invariant vector fields \mathfrak{g} admits a nilpotent stratification of step m , namely

$$\mathfrak{g} = \mathfrak{g}_1 \oplus \dots \oplus \mathfrak{g}_m, \quad \mathfrak{g}_i \neq \{0\}, \quad \forall i = 1, \dots, m,$$

with

$$[\mathfrak{g}_1, \mathfrak{g}_j] = \mathfrak{g}_{j+1}, \quad \forall 1 \leq j \leq m, \quad \text{and} \quad \mathfrak{g}_{m+1} = \{0\}.$$

A left-invariant sub-Riemannian structure on G is obtained by defining a scalar product on \mathfrak{g}_1 or, equivalently, by declaring a set $X_1, \dots, X_k \in \mathfrak{g}_1$ a global orthonormal frame. In particular, $\blacktriangle|_q = \mathfrak{g}_1|_q$, for all $q \in G$. The group exponential map,

$$\exp_G : \mathfrak{g} \rightarrow G,$$

associates with $v \in \mathfrak{g}$ the element $\gamma(1)$, where $\gamma : [0, 1] \rightarrow G$ is the unique integral line of the vector field v such that $\gamma(0) = 0$. Since G is simply connected and \mathfrak{g} is nilpotent, \exp_G is a smooth diffeomorphism. Thus we identify $G \simeq \mathbb{R}^m$, endowed with a polynomial product law. The adjoint endomorphism $\text{ad}_X : \mathfrak{g} \rightarrow \mathfrak{g}$ is:

$$\text{ad}_X(Y) := [X, Y], \quad \forall X, Y \in \mathfrak{g}.$$

Notice that, if $X \in \mathfrak{g}_j$, then

$$\text{ad}_X|_{\mathfrak{g}_\ell} : \mathfrak{g}_\ell \rightarrow \mathfrak{g}_{\ell+j}, \quad \forall \ell, j = 1, \dots, m,$$

by the graded structure of \mathfrak{g} .

Remark 17. In the literature, these structures are also referred to as *Carnot groups* of type (k, n) , where $k = \dim \blacktriangle = \text{rank } \mathfrak{g}_1$ is the *rank* of the distribution and n is the dimension of G .

By definition, Carnot groups are left-invariant sub-Riemannian structures (see Definition 42). Then it is natural to fix a left-invariant volume, that is proportional to Popp's one \mathcal{P} by Corollary 43. Moreover, we restrict to left-invariant complements \mathbf{c} . In this setting, we rewrite the compatibility condition $\chi^{(\mathbf{c}, \mathcal{P})} = 0$ in a more invariant fashion. To do this, we choose left invariant orthonormal frames X_1, \dots, X_k for $\mathbf{\Delta}$ and left-invariant frames X_{k+1}, \dots, X_n for \mathbf{c} . Thanks to the splitting $\mathbf{\Delta} \oplus \mathbf{c}$ we have the projections $\pi_{\mathbf{c}}$ and $\pi_{\mathbf{\Delta}}$ on TM .

Lemma 59. *For Carnot groups, the compatibility condition for left-invariant volumes and complements is*

$$\mathrm{Tr}(\pi_{\mathbf{c}} \circ \mathrm{ad}_{X_i}) = 0, \quad \forall i = 1, \dots, k.$$

Proof. We rewrite the compatibility condition as

$$\sum_{j=k+1}^n c_{ji}^j X_i + X_i(\theta) = \sum_{j=k+1}^n \nu^j([X_j, X_i]) = -\mathrm{Tr}(\pi_{\mathbf{c}} \circ \mathrm{ad}_{X_i} \circ \pi_{\mathbf{c}}) = -\mathrm{Tr}(\pi_{\mathbf{c}} \circ \mathrm{ad}_{X_i}) = 0,$$

where we used the cyclic property of the trace and the fact that projections are idempotent. The function θ appearing in the compatibility condition is constant, since both \mathcal{P} and X_1, \dots, X_n are left-invariant. \square

On Carnot groups, thanks to the canonical identification $\mathfrak{g}_j = \mathbf{\Delta}^j / \mathbf{\Delta}^{j-1}$, there is a natural left-invariant Riemannian extension. In fact, it is sufficient to consider the orthogonal direct sum

$$\mathfrak{g} = \mathbf{\Delta} \oplus \mathfrak{g}_2 \oplus \dots \oplus \mathfrak{g}_m,$$

where on each factor \mathfrak{g}_j , with $j \geq 2$, we have a well defined scalar product induced by the maps π_j of Eq. (10). This gives a smooth left-invariant scalar product on TM . Popp's volume \mathcal{P} is the Riemannian volume of such natural Riemannian extension. We can address questions **Q2** and **Q3** of Section 1, restricted to the class of left-invariant volumes and complements.

The most natural complement comes from the very definition of the Carnot structure:

$$\mathbf{c}_0 := \mathfrak{g}_2 \oplus \dots \oplus \mathfrak{g}_m,$$

Proposition 60. *For any Carnot group G , we have $\Delta_{\mathcal{P}} = L^{\mathbf{c}_0}$.*

Proof. By definition of \mathbf{c}_0 and the stratified structure of the Carnot algebra, we have

$$\pi_{\mathbf{c}_0} \circ \mathrm{ad}_{X_i} \circ \pi_{\mathbf{c}_0} = \mathrm{ad}_{X_i}|_{\mathbf{c}_0},$$

that is indeed nilpotent, since ad_{X_i} is, and $\mathrm{Tr}(\mathrm{ad}_{X_i}) = 0$. Then the compatibility condition is satisfied. \square

Any distinct left-invariant complement is the graph of a non-trivial linear map $\ell : \mathbf{c}_0 \rightarrow \mathbf{\Delta}$, that is

$$\mathbf{c}_{\ell} := \{X + \ell(X) \mid X \in \mathbf{c}_0\}.$$

Proposition 61. *If $\Delta_{\mathcal{P}} = L^{\mathbf{c}}$, then $\mathbf{c} = \mathbf{c}_{\ell}$, with*

$$\mathrm{Tr}(\ell \circ \mathrm{ad}_{X_i}) = 0, \quad \forall i = 1, \dots, k.$$

Proof. We use the shorthand π_{ℓ} (resp. π_0) for the two projections $\pi_{\mathbf{c}_{\ell}}$ and $\pi_{\mathbf{c}_0}$ respectively. Indeed

$$\pi_{\ell} = \pi_0 + \ell \circ \pi_0.$$

The new complement is compatible iff, for all $i = 1, \dots, k$

$$0 = \mathrm{Tr}(\pi_{\ell} \circ \mathrm{ad}_{X_i} \circ \pi_{\ell}) = \mathrm{Tr}(\pi_{\ell}^2 \circ \mathrm{ad}_{X_i}) = \mathrm{Tr}(\pi_{\ell} \circ \mathrm{ad}_{X_i}) = \mathrm{Tr}(\ell \circ \pi_0 \circ \mathrm{ad}_{X_i}) = \mathrm{Tr}(\ell \circ \mathrm{ad}_{X_i}),$$

where we used the cyclic property of the trace, the fact that projectors are idempotent, that \mathbf{c}_0 satisfies the compatibility condition and the fact that the image of ad is contained in \mathbf{c}_0 . \square

Any ℓ such that $\ell(\mathfrak{g}_2) = 0$ gives an example where the compatible complement is non-unique. This class is always non-empty for any Carnot group of step $m \geq 3$, but is trivial for step 2.

We provide an example of non-uniqueness for corank 1 Carnot groups. Choose a structure with

$$\mathfrak{g}_1 = \mathrm{span}\{X_1, \dots, X_k\}, \quad \mathfrak{g}_2 = \mathrm{span}\{X_0\},$$

$$[X_i, X_j] = A_{ij}X_0, \quad i, j = 1, \dots, k$$

for a skew symmetric matrix $A \in \mathfrak{so}(k)$ (the singular values of A define uniquely any Carnot group of type $(k, k+1)$, up to isometries, see for example [24]). Any other compatible complement is of the form:

$$\mathbf{c} = \mathrm{span}\{X_0 + \sum_{j=1}^k \ell_j X_j\}, \quad \text{with} \quad \ell \in \ker A.$$

Then the number of distinct compatible complements is equal to $\dim \ker A$. In particular, the only case in which we have uniqueness is when $\ker A = \{0\}$, that is a contact Carnot group.

8 Corank 1 structures

Consider a sub-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ with $\text{rank } \blacktriangle = \dim M - 1$. We assume that \blacktriangle has step 2, that is $\blacktriangle_q^2 = T_q M$ for all $q \in M$. This is a popular class of structures including contact, quasi-contact, and more degenerate ones. Locally $\blacktriangle = \ker \eta$ for some one-form η . The endomorphism $J_\eta : \Gamma(\blacktriangle) \rightarrow \Gamma(\blacktriangle)$ is defined by

$$g(X, JY) = d\eta(X, Y), \quad \forall X, Y \in \Gamma(\blacktriangle).$$

It is skew-adjoint w.r.t. the sub-Riemannian scalar product, namely $J^* = -J$. Moreover, $J \neq 0$. In fact, if this were the case, a simple application of Cartan's formula gives, for any local frame X_1, \dots, X_k of \blacktriangle :

$$\eta([X_i, X_j]) = -d\eta(X_i, X_j) + \underbrace{X_i(\eta(X_j))}_{\overline{X_i(\eta(X_j))}} - \underbrace{X_j(\eta(X_i))}_{\overline{X_j(\eta(X_i))}} = g(JX_j, X_i) = 0,$$

in contradiction with the step 2 condition. It follows that $\text{Tr}(JJ^*) > 0$. This observation leads the following construction. If $g \in C^\infty(M)$, then $g\eta$ gives the same distribution. On the other hand one can check that if $\eta' = g\eta$, then $d\eta'|_{\blacktriangle} = g d\eta|_{\blacktriangle}$, that is $J_{\eta'} = gJ_\eta$. We can fix η , up to a sign, with the condition

$$\|J\|^2 = \text{Tr}(JJ^*) = \sum_{i,j=1}^k g(X_i, JX_j)^2 = 1.$$

In this case, we say that the local one-form η is *normalized*. The existence of a global, normalized one-form depends on the distribution. In particular there exists a unique (up to a sign), normalized global one-form if and only if \blacktriangle is co-orientable (i.e. the quotient bundle TM/\blacktriangle is orientable or, equivalently, trivial).

Remark 18. If TM/\blacktriangle is not orientable, a global normalized one-form does not exist. This is the case, for example, for the structure on $M = \mathbb{R}^2 \times \mathbb{R}P^1$ with

$$\blacktriangle := \ker(\sin \theta dx - \cos \theta dy),$$

where we identify $\mathbb{R}P^1 \simeq \mathbb{R}/\pi\mathbb{Z}$ with the coordinate θ . One still could work with a global object $[\eta]$, section of the bundle whose fibers are $(T_q^* M \setminus \{0\})/\mathbb{Z}_2$. In particular the value $[\eta](q)$ is the equivalence class $\{\pm\eta(q)\}$. All the results and formulas appearing in the following do not depend on the choice of the representative of the class at each point. For this reason, it is not a restriction to assume that η is a well-defined global one-form.

We investigate the relation $\Delta_\omega = L^c$. We first rewrite the compatibility condition $\chi^{(c,\omega)} = 0$ in this case. Observe that any complement is (locally) generated by a vector field X_0 , that we normalize with $\eta(X_0) = 1$.

Lemma 62. *Let $\mathbf{c} = \text{span}\{X_0\}$, with $\eta(X_0) = 1$. The compatibility equation $\chi^{(\omega,\mathbf{c})} = 0$ is*

$$\sum_{i=1}^k d\eta(X_0, X_i)X_i = \text{grad}(\theta) \quad \text{with} \quad \theta = \log |\omega(X_1, \dots, X_k, X_0)|.$$

Proof. In the local adapted frame X_1, \dots, X_k, X_0 , the compatibility equation $\chi^{(\omega,\mathbf{c})} = 0$ is

$$\sum_{i=1}^k c_{0i}^0 X_i + \text{grad}(\theta) = 0.$$

This simple form is due to the fact that the structure has corank 1. Then we rewrite

$$c_{0i}^0 = \eta([X_0, X_i]) = -d\eta(X_0, X_i) + \underbrace{X_0(\eta(X_i))}_{\overline{X_0(\eta(X_i))}} - \underbrace{X_i(\eta(X_0))}_{\overline{X_i(\eta(X_0))}},$$

where we used Cartan's formula and the normalization $\eta(X_0) = 1$. □

The existence and uniqueness of compatible complements in this case depends on the dimension of $\ker J$.

Proposition 63. *Fix a volume ω , and let $p := \dim \ker J$. Then:*

- if $p = 0$, then $\exists!$ complement \mathbf{c} such that $L^c = \Delta_\omega$;
- if $p > 0$, the space of \mathbf{c} such that $L^c = \Delta_\omega$ is either empty, or an affine space over $\ker J$.

Proof. Let $p = 0$, we check that the equation $\chi^{(\omega,\mathbf{c})} = 0$ has a unique solution for fixed ω . Assume $\mathbf{c} = \text{span}\{Z + \xi\}$, where Z is transverse to \blacktriangle such that $\eta(Z) = 1$ and $\xi \in \Gamma(\blacktriangle)$. Then, by Lemma 62, the compatibility condition gives

$$\text{grad}(\theta) = d\eta(Z + \xi, X_i)X_i = d\eta(Z, X_i)X_i + g(\xi, JX_i)X_i = d\eta(Z, X_i)X_i - J\xi,$$

There is a unique ξ satisfying this equation, namely $\xi = -J^{-1}[\text{grad}(\theta) + d\eta(Z, X_i)X_i]$. Now let $p > 0$, and suppose that X_0 and X'_0 are two different (normalized) generators for complements \mathbf{c} and \mathbf{c}' , compatible with the volume ω , with $\eta(X_0) = \eta(X'_0) = 1$. The normalization implies $X_0 - X'_0 \in \blacktriangle$. According to Lemma 62

$$\sum_{i=1}^k d\eta(X_0, X_i)X_i = \text{grad}(\theta), \quad \text{and} \quad \sum_{i=1}^k d\eta(X'_0, X_i)X_i = \text{grad}(\theta)',$$

with $\theta' = \log |\omega(X_1, \dots, X_k, X'_0)|$. Since $X_0 - X'_0 \in \blacktriangle$,

$$\theta = \log |\omega(X_1, \dots, X_k, X_0)| = \log |\omega(X_1, \dots, X_k, X'_0)| = \theta'.$$

Then we have

$$d\eta(X_0 - X'_0, X_i) = 0, \quad \forall i = 1, \dots, k.$$

This implies $X_0 - X'_0 \in \ker J$. Conversely, one can check that for any (normalized) generator X_0 of a compatible complement \mathbf{c} , any other complement generated by $X_0 \oplus \ker J$ is compatible with the the same volume ω . \square

8.1 On natural Riemannian extensions and Popp's volume

We end this section with a general result about volumes of Riemannian extensions of corank 1 structures.

Proposition 64. *Let $(M, \blacktriangle, \mathbf{g})$ a corank 1 sub-Riemannian structure with normalized one-form η . Take any local vector field Z such that $\eta(Z) = 1$. Consider the (local) Riemannian extension of $(M, \blacktriangle, \mathbf{g})$ obtained by declaring Z an orthonormal vector field. Then the Riemannian volume of this Riemannian extension does not depend on the choice of Z and is equal to Popp's volume.*

Proof. This statement is an application of the explicit formula of [8, Proposition 16], stated in the contact case, whose proof holds unchanged for corank 1 sub-Riemannian structures (it only requires that $\text{Tr}(JJ^*) \neq 0$). \square

Proposition 64 gives another reason for considering Popp's volume in the corank 1 case (besides the fact, valid in general, that Popp's volume is N-intrinsic). We now specialize Proposition 63 to the contact and quasi-contact cases, giving an answer to **Q2** and **Q3** of Section 1.

9 Contact structures

Let M be a smooth manifold with $\dim M = 2d + 1$. A corank sub-Riemannian structure is *contact* if $\ker J = 0$. This is the least degenerate case. Since M is odd dimensional, there exists a unique vector Z , called *Reeb vector field*, such that $d\eta(Z, \cdot) = 0$ and $\eta(Z) = 1$.

The following corollaries of Proposition 63 give affirmative answers to **Q2** and **Q3** mentioned in Section 1.

Corollary 65. *For any volume ω there exists a unique complement \mathbf{c} such that $L^{\mathbf{c}} = \Delta_{\omega}$. This complement is generated by the vector*

$$X_0 = Z - J^{-1} \text{grad}(\theta),$$

where $\theta = \log |\omega(X_1, \dots, X_k, Z)|$.

Proof. The explicit formula is already contained in the proof of Proposition 63: here we only specified it in terms of the Reeb vector field Z . \square

Contact sub-Riemannian structures have a natural Riemannian extension, obtained by declaring Z a unit vector orthonormal to \blacktriangle . By Proposition 64, Popp's volume is the Riemannian volume of this extension.

Corollary 66. *Let \mathcal{P} be the Popp's volume. The unique complement \mathbf{c} such that $L^{\mathbf{c}} = \Delta_{\mathcal{P}}$ is generated by the Reeb vector field. Moreover, \mathcal{P} is the unique volume (up to constant rescaling) with this property.*

Proof. By [8, Rmk. 4], $\theta = 1$. Then $\text{grad}(\theta) = 0$ and $X_0 = Z$. The uniqueness follows from Lemma 58. \square

9.1 Integrability conditions

The inverse problem, namely for a fixed complement \mathbf{c} , find a volume ω such that $L^{\mathbf{c}} = \Delta_{\omega}$ is more complicated (and, in general, has no solution). In the contact case we find an explicit integrability condition.

Proposition 67. *Let $\mathbf{c} = \text{span}\{X_0\}$. Define the one-form $\alpha := \frac{i_{X_0} d\eta}{\eta(X_0)}$ and the function $g = \frac{d\alpha \wedge \eta \wedge (d\eta)^{d-1}}{\eta \wedge (d\eta)^d}$. Then there exists a volume ω such that $L^{\mathbf{c}} = \Delta_{\omega}$ if and only if*

$$d\alpha - dg \wedge \eta - gd\eta = 0.$$

In this case, ω is unique up to constant rescaling.

Remark 19. If \mathbf{c} is the Reeb direction, $\alpha = 0$, the integrability condition is satisfied and we recover Corollary 66.

Proof. We assume $\mathbf{c} = \text{span}\{Z + \xi\}$ for some $\xi \in \Gamma(\blacktriangle)$. This is equivalent to normalize X_0 in such a way that $\eta(X_0) = 1$. A volume ω is uniquely specified by the function $\theta = \log |\omega(X_1, \dots, X_k, Z)|$. The same proof of Proposition 63 leads to the compatibility condition

$$-J\xi = \text{grad}(\theta).$$

We have to solve the following problem: given an horizontal vector field X , find a function $\theta \in C^\infty(M)$ such that $X = \text{grad}(\theta)$. With X we can associate a one-form α such that

$$\alpha(Y) = g(X, Y), \quad \forall Y \in \Gamma(\blacktriangle).$$

For our case, $X = -J\xi$. Then $\alpha(Y) = g(-J\xi, Y) = -g(Y, J\xi) = -d\eta(Y, \xi) = i_\xi Y = i_{X_0} d\eta(Y)$ and in this case we set $\alpha := i_{X_0} d\eta$. In the language of forms, the above problem is equivalent to $\alpha|_{\blacktriangle} = d\theta|_{\blacktriangle}$, and has solution iff

$$\alpha + g\eta = d\theta, \quad \text{for some } g \in C^\infty(M)$$

A (local) necessary and sufficient condition for the existence of such θ is that there exists a closed representative in the class $\alpha + g\eta$. Namely g must satisfy

$$d\alpha + dg \wedge \eta + g d\eta = 0. \quad (16)$$

If such a g exists, is uniquely expressed in terms of α . Taking the appropriate wedge product with η , and then with $d\eta$ $d-1$ times, we get

$$g\eta \wedge (d\eta)^d = -\eta \wedge d\alpha \wedge (d\eta)^{d-1} \quad \Rightarrow \quad g = -\frac{d\alpha \wedge \eta \wedge (d\eta)^{d-1}}{\eta \wedge (d\eta)^d}.$$

The n -form at the denominator is never zero, it is equivalent to the non-degeneracy assumption $\ker J = \{0\}$. We have to check that such a g gives a closed representative, i.e. that f solves Eq. (16). Plugging the explicit expression of g into Eq. (16) gives the condition of our statement. Uniqueness follows from Lemma 58. \square

Remark 20. In the three dimensional case $d = 1$ and the integrability condition is simplified. Fix a basis X_1, X_2 . We can assume $d\eta(X_1, X_2) = 1$ and we get for the function g

$$g = -\frac{d\alpha \wedge \eta \wedge (d\eta)^{d-1}}{\eta \wedge (d\eta)^d} = -\frac{d\alpha \wedge \eta(X_1, X_2, Z)}{\eta \wedge d\eta(X_1, X_2, Z)} = -d\alpha(X_1, X_2)$$

One can check that, when restricted to \blacktriangle , Eq. (16) is always satisfied. Then the only non-trivial condition is given by taking the contraction with the Reeb field: $i_Z d\alpha + d(d\alpha(X_1, X_2)) = 0$.

10 Quasi-contact structures

Let M a smooth manifold with $\dim M = 2d+2$. A corank 1 sub-Riemannian structure $(M, \blacktriangle, \mathbf{g})$ is *quasi-contact* if $\blacktriangle_0 := \ker J$ has positive dimension. We assume *minimal degeneration*, that is $\dim \blacktriangle_0 = 1$.

10.1 The quasi-Reeb vector field

We discuss the construction of a ‘‘canonical’’ vector transverse to \blacktriangle , analogous to the Reeb vector field, in the quasi-contact case. We learned this construction from [13] (where a different normalization is used).

In the construction that follows we always assume that $J_q : \blacktriangle_q \rightarrow \blacktriangle_q$ has distinct eigenvalues at all points on M . This is true for a generic quasi-contact structure and outside a codimension 3 closed stratified subset. Thus, J has distinct imaginary eigenvalues $\pm i\lambda_j$ with

$$0 = \lambda_0 < \lambda_1 < \dots < \lambda_d \in \mathbb{R}.$$

where λ_i are smooth functions on M . The normalization condition gives $2 \sum_{i=1}^d \lambda_i^2 = 1$.

For any j , denote by $\blacktriangle_j \subset \blacktriangle$ the 2-dimensional (real) eigenspace associated with the eigenvalues $\pm i\lambda_j$ (we use the same notation for the eigenspaces and the associated sub-bundle). By definition of J , the \blacktriangle_j are mutually orthogonal, for all $i = 0, 1, \dots, d$. We can choose generators X_j, Y_j of each distribution \blacktriangle_j by taking smooth families of (generalized) eigenvectors of J . Namely

$$J(X_j + iY_j) = i\lambda_j(X_j + iY_j), \quad \forall j = 1, \dots, d.$$

10.2 An example of non-existence

Proposition 63 states that, in the contact case (i.e. when J is non-degenerate), for any volume ω there exists a unique complement \mathbf{c} . In the quasi-contact case the situation changes dramatically. Compatible complements are never unique as soon as $\ker J \neq \{0\}$. Surprisingly, compatible complements might not exist. We discuss an example where for a given volume (Popp's one) there are *no* compatible complements.

Example 6. Consider the quasi-contact structure on $M = \mathbb{R}^4$, with coordinates (x, y, z, w) , defined by

$$\eta = \frac{g}{\sqrt{2}} \left(dw - \frac{y}{2} dx + \frac{x}{2} dy \right),$$

where g is any monotone, strictly positive function (e.g. $g = e^z$). The metric is defined by the following global orthonormal frame for \blacktriangle :

$$X = \frac{1}{\sqrt{g}} \left(\partial_x + \frac{1}{2} y \partial_w \right), \quad Y = \frac{1}{\sqrt{g}} \left(\partial_y - \frac{1}{2} x \partial_w \right), \quad Z = \frac{1}{\sqrt{g}} \partial_z.$$

This is essentially $\mathbb{H}_3 \oplus \mathbb{R}$ with a scaled metric. One can check that

$$d\eta = \frac{\dot{g}}{\sqrt{2}} \left(dz \wedge dw + \frac{y}{2} dx \wedge dz - \frac{x}{2} dy \wedge dz \right) + \frac{g}{\sqrt{2}} dx \wedge dy,$$

where the $\dot{g} = \partial_z g$. The representative matrix A of $d\eta$ in coordinates (x, y, z, w) is

$$A \simeq \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & g & \frac{y}{2} \dot{g} & 0 \\ -g & 0 & -\frac{x}{2} \dot{g} & 0 \\ -\frac{y}{2} \dot{g} & \frac{x}{2} \dot{g} & 0 & \dot{g} \\ 0 & 0 & -\dot{g} & 0 \end{pmatrix}, \quad \text{with} \quad \det A = \frac{g^2 \dot{g}^2}{4} > 0.$$

In particular, $\ker d\eta = \{0\}$. Moreover, one can check that

$$JX = -\frac{1}{\sqrt{2}} Y, \quad JY = \frac{1}{\sqrt{2}} X, \quad JZ = 0.$$

Thus we have the correct normalization

$$\|J\|^2 = \text{Tr}(JJ^*) = 1.$$

Choose $\omega = \mathcal{P}$, the Popp's volume. We look for a complement $\mathbf{c} = \text{span}\{X_0\}$ such that $L^{\mathbf{c}} = \Delta_{\omega}$. We can assume, rescaling X_0 , that $\eta(X_0) = 1$. Using the explicit formula from [8], we have

$$\theta_{\mathcal{P}} = \log |\mathcal{P}(X_1, \dots, X_k, X_0)| = \frac{1}{\sqrt{\sum_{i,j=1}^k d\eta(X_i, X_j)^2}} = \frac{1}{\|J\|} = 1.$$

We look for solutions of the equation $\chi^{(\mathcal{P}, \mathbf{c})} = 0$. Since $\theta_{\mathcal{P}} = 1$, we get from Lemma 62

$$d\eta(X_0, X_i) = 0, \quad \forall i = 1, \dots, k.$$

This implies $X_0 \in \ker d\eta$, but $d\eta$ has trivial kernel.

Remark 25. It can be shown that $L^{\mathbf{c}} \neq \Delta_{\mathcal{P}}$ is a generic property for quasi-contact structures. The precise statement and the proof involves techniques from transversality theory and is out of the scope of this paper.

This result in the case $d = 1$ (that is $\dim M = 4$, the lowest dimension for quasi-contact structures) is particularly surprising, for the following reason. The nilpotent approximation of any corank 1 sub-Riemannian structure is a corank 1 Carnot group. This is a sub-Riemannian structure on \mathbb{R}^n that, in coordinates $(x, z) \in \mathbb{R}^{n-1} \times \mathbb{R}$ is generated by the global orthonormal frame:

$$X_i = \frac{\partial}{\partial x_i} - \frac{1}{2} \sum_{j=1}^{n-1} A_{ij} x_j \frac{\partial}{\partial z}, \quad i = 1, \dots, n-1, \quad (17)$$

where $A \in \mathfrak{so}(n-1)$. Isometry classes of corank 1 Carnot groups are determined by the string $0 \leq \alpha_1 < \dots < \alpha_p$ of the p distinct non-negative singular values of A , up to multiplication for a global constant, and their geometric multiplicities (see [24] for corank 1 structures, and [7, Rmk. 1] for higher corank). In particular, for $n = 4$, $A \in \mathfrak{so}(3)$ there is a unique such a string: $(0, 1)$ where 0 has multiplicity 1 and 1 has multiplicity 2. In particular, any corank 1 Carnot group in dimension $n = 4$ is isometric to the one defined by the frame (17), with:

$$A = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

It follows that any corank 1 sub-Riemannian in dimension 4 is quasi-contact and equi-nilpotentizable. In this case, we proved that there exists a unique N-intrinsic volume (up to scaling) and is given by Popp's one, and a unique N-intrinsic Laplacian $\Delta_{\mathcal{P}}$ (see Theorem 38). In our example, this unique N-intrinsic Laplacian has *no* compatible complement or, in other words, the macroscopic diffusion operator has no microscopic counterpart.

11 Convergence of random walks

We let M be a smooth, geodesically complete (sub)-Riemannian manifold. Our goal here is to give fairly general conditions for a sequence of random walks on M to converge. In particular, we will work with a larger class of random walks than those treated elsewhere in the paper, which will require some differences in notation. In this section we work in the general (sub)-Riemannian framework, so that geodesics are specified by the choice of an initial covector $\lambda \in T^*M$ and the (sub)-Riemannian distribution has rank $k \leq n = \dim M$.

11.1 A general class of random walks

By a random walk, in this section, we mean a process X_t that travels along a geodesic (with constant speed) for a fixed length of time (say $h > 0$), at which time a new geodesic is chosen at random (and independently of the particle's past), which is then traveled for time h , and this procedure repeats indefinitely. The result is that X_0, X_h, X_{2h}, \dots is a Markov chain on M , and for $t \in (ih, (i+1)h)$, X_t interpolates between X_{ih} and $X_{(i+1)h}$ along a geodesic between them. We distinguish this from the related scheme of following a randomly chosen geodesic for a random, exponentially distributed length of time, at which point a new geodesic is chosen and a new (independent) exponential clock is started, and so on, which we will refer to as a random flight. (The logic being that a walk involves regular steps, and one does not turn midstep, whereas a flight seems to evoke a mode of travel which can change direction at any time.) A random flight has the nice property that, because the exponential is memory-less, it is a time-homogeneous process (assuming the choice of geodesic is time-homogeneous, of course).

More concretely, for $h > 0$, we have a family of probability measures μ_q^h on the cotangent spaces T_q^*M . Then $X_{h(i+1)}^h = \exp_{X_{hi}^h}(h, \lambda/h)$ where λ is a covector chosen according to $\mu_{X_{hi}^h}^h$ independently of the previous steps, and the path travels from X_{hi}^h to $X_{h(i+1)}^h$ along the geodesic determined by λ , at constant speed given by $\|\lambda\|/h$ (and for a distance given by $\|\lambda\|$). That is, for $t \in [hi, h(i+1)]$, we have $X_t^h = \exp_{X_{hi}^h}((t-hi), \lambda/h)$.

Note that, unlike earlier in the paper, here we index our random walks by the size of the time-step, instead of by the size of the spacial-step. This is because we no longer require that μ_q^h be supported on covectors of fixed length. Indeed, earlier the paper, we have considered the case when μ_q is a probability measure on covectors of length 1, and μ_q^h is given by parabolic scaling of μ_q ; that is, if f_h is the fiber-preserving map that takes a covector λ to $\sqrt{2hk}\lambda$, then μ_q^h is the pushforward of μ_q by f_h . (Also compare the present set-up with Definition 51.) However, here we work in the generality of random walks as just described.

We are interested in the convergence of a sequence of random walks X_t^h (on M) as $h \rightarrow 0$ to a (continuous time) diffusion. A key role is played by the following associated operator (which we have already seen). For any $\phi \in C_0^\infty(M)$, let

$$L_h\phi(q) = \frac{1}{h} (\mathbb{E} [\phi(X_h^h) | X_0^h = q] - \phi(q)) = \frac{1}{h} \left(\int_{T_q^*M} \phi(\exp_q(h, \lambda/h)) \mu_q^h(\lambda) - \phi(q) \right).$$

The idea is that, under suitable assumptions, the random walks will converge to the diffusion generated by $\lim_{h \rightarrow 0} L_h$ as $h \rightarrow 0$. That the behavior in the limit is governed only by the second-order operator, and not any other features of the random walks, should be seen as a version of Donsker invariance.

11.2 Background on diffusions and pathspace

Next, we recall some basic facts about the diffusions which will arise as limits of our random walks. In what follows, we assume that L is a smooth second-order operator with non-negative-definite principal symbol and without zeroth-order term. This means that, in local coordinates x_1, \dots, x_n , L can be written as

$$\sum_{i,j=1}^n a_{ij} \partial_{x_i} \partial_{x_j} + \sum_{i=1}^n b_i \partial_{x_i},$$

where the a_{ij} and b_i are smooth functions, and the matrix $[a_{ij}]$ is symmetric and non-negative definite. Alternatively, L can be locally written as

$$Y_0 + \sum_{i=1}^n Y_i^2$$

for smooth vector fields Y_0, Y_1, \dots, Y_n . Recall that, in this case, there exists a unique diffusion generated by L , although this diffusion may explode in finite time (that is, with positive probability, a path might exit every compact subset of M in finite time). For simplicity, we assume that the diffusion generated by L does not explode (indeed, our analysis is fundamentally local). Then if we let P_q be the measure on Ω_M corresponding to this diffusion (starting from q), P_q is uniquely characterized by the fact that

$$\phi(\omega_t) - \int_0^t L\phi(\omega_s) ds \quad \text{is a } P_q\text{-martingale with } P_q(\omega_0 = q) = 1.$$

A flexible and powerful approach to the convergence of random walks, based on martingale theory, was essentially already provided by Stroock and Varadhan in [30] (specifically, in Section 11.2). However, they develop their approach in the setting of Euclidean space. To generalize to the Riemannian or sub-Riemannian setting is relatively straight-forward. The main things one needs to do are to replace the systematic use of Euclidean coordinates by the use of coordinate-free or local coordinate-based expressions and to replace the linear interpolation from X_{ih} and $X_{(i+1)h}$ via the appropriate geodesic segment. In regard to the local versus global nature of the problem, in principle we could even allow for the limiting diffusion to explode in finite time, as just described. However, this requires a more substantial reworking of the underlying framework (M should be compactified by adding a point at infinity— the Alexandrov one-point compactification, the space of continuous paths on M should be suitably modified, etc.), which would be a substantial digression from the main thrust of this work. Indeed, we are most interested in understanding various choices of sub-Laplacian on a sub-Riemannian manifold, and this is really a local problem. Thus, there is no harm in restricting our attention to the situation when the underlying diffusion on M does not explode.

Following Section 1.3 of [30], let Ω_M be the space of continuous paths from $[0, \infty)$ to M . For $\omega \in \Omega_M$, let ω_t be the position of ω at time t . Then we can define a metric on Ω by

$$d_{\Omega_M}(\omega, \tilde{\omega}) = \sum_{i=1}^{\infty} \frac{1}{2^i} \frac{\sup_{0 \leq t \leq i} d_M(\omega_t, \tilde{\omega}_t)}{1 + \sup_{0 \leq t \leq i} d_M(\omega_t, \tilde{\omega}_t)}. \quad (18)$$

This metric makes Ω_M into a Polish space, and the notion of convergence that it induces is that of uniform convergence on bounded time intervals. We equip Ω_M with its Borel σ -algebra \mathcal{M} and give it the natural filtration \mathcal{M}_t generated by $\{\omega_s : 0 \leq s \leq t\}$. For probability measures on Ω_M , we will always work with respect to the weak topology (in probabilists' terminology; this is the weak* topology in the language of functional analysis, corresponding to convergence against bounded, continuous test functions). We will generally be interested in a (Markov) family of probability measures, indexed by points of M . (Indeed, we have already seen that our random walks correspond to such families.) In that case, we will write them as P_q for $q \in M$.

The case of $M = \mathbb{R}^n$ (with the standard Euclidean metric) is the one considered by Stroock and Varadhan. Rather than generalize their approach directly (which would perhaps be more mathematically satisfying, but would also take longer), we exploit a standard trick to make the general case a special case of the Euclidean case. In particular, first note that the sub-Riemannian metric on M can be extended to a Riemannian metric (not necessarily in any canonical way, but that doesn't matter here). Then the Nash embedding theorem implies that M can be isometrically embedded into some Euclidean space of high enough dimension, say \mathbb{R}^N (and this embedding is smooth, although not necessarily proper if M is not compact).

11.3 Control of the interpolating paths

Assuming that the operators L_h converge to an operator L of the type just discussed, in the sense that, for any $\phi \in C_0^\infty(M)$, we have $L_h\phi \rightarrow L\phi$ uniformly on compacts as $h \rightarrow 0$, is almost the only condition that we need. (Indeed, for $M = \mathbb{R}^N$ this is enough.) However, the convergence of the L_h depends only on the positions of the random walks at the sequences of times $0, h, 2h, \dots$, and does not necessarily control the path traveled between these times. To give a simple example of what can go wrong, consider $M = \mathbb{S}^1$ with the standard metric. If $X_0^h = q$, let the walk travel a full circle either clockwise or counter-clockwise, each with probability $1/2$, at constant speed in time h . In other words, if θ is the standard coordinate, we let $\mu_q^h = \pm 2\pi d\theta$ where the sign has probability $1/2$ of being positive or negative. Then L^h is the zero operator for all h (since $q = X_0^h = X_h^h = X_{2h}^h = \dots$), but it's clear from Eq. (18) that the walks are not converging. This is in contrast to Euclidean space, where the entire path X_t^h for $t \in [ih, (i+1)h]$ is determined by X_{ih}^h and $X_{(i+1)h}^h$. (Indeed, in [30], each step of the random walk is determined by choosing a point in \mathbb{R}^N according to, in slightly modified notation, a probability measure Π_q^h on \mathbb{R}^N . Then the path of the particle during this step is given by linear interpolation. The appropriate geometric generalization of this is our measure μ_q^h on the co-tangent space, which we can think of as giving the next step in the walk along with a path for the particle to travel to get there.)

Thus, we will also need to assume that for any $\rho > 0$, any compact $Q \subset M$, and any $\alpha > 0$, there exists $h_0 > 0$ such that

$$\frac{1}{h} P_q^h \left[\sup_{0 \leq s \leq h} d_M(X_0^h, X_s^h) \leq \rho \right] > 1 - \alpha \quad (19)$$

whenever $q \in Q$ and $h < h_0$. This assumption is not especially restrictive. First of all, on Euclidean space, Lemmas 11.2.1 and 11.2.2 of [30] show that if $L_h \phi \rightarrow L \phi$ uniformly on compacts as $h \rightarrow 0$ for any $\phi \in C_0^\infty(\mathbb{R}^N)$, then, under some global uniformity assumptions on the μ_q^h , for any $\rho > 0$ we have

$$\lim_{h \rightarrow 0} \sup_{\substack{q \in \mathbb{R}^N \\ i \in \{0, 1, 2, \dots\}}} \frac{1}{h} P_q^h \left[d_{\mathbb{R}^N}(X_{ih}^h, X_{(i+1)h}^h) > \rho \right] = 0,$$

and thus for any $0 < T < \infty$,

$$\lim_{h \rightarrow 0} \sup_{q \in \mathbb{R}^N} P_q^h \left[\sup_{\substack{i: 0 \leq hi \leq T \\ 0 \leq s \leq h}} d_{\mathbb{R}^N}(X_{ih}^h, X_{ih+s}^h) \leq \rho \right] = 1.$$

Thus, we are essentially assuming that the behavior of X_t^h for $t \in (ih, (i+1)h)$ is comparable to the behavior of the step $X_{(i+1)h}^h | X_{ih}^h$, which as already mentioned, is automatic in the Euclidean case but not necessarily in other situations.

Further, we note that the condition of Eq. (19) follows from assuming that $\frac{1}{h} \mu_q^h [\|\lambda\| \leq \rho]$ approaches 1 as $h \rightarrow 0$ uniformly for q in any compact set (here λ is a covector in T_q^*M). In particular, for random walks of the type we consider earlier in the paper, this condition is trivially satisfied by taking $h < \rho^2/2k$, because by construction a step of X_t^h consists in traveling a geodesic for distance $\sqrt{2hk}$; that is, μ_q^h is supported on covectors of length $\sqrt{2hk}$ for all q (again, see Definition 51).

11.4 Convergence results

We can now prove the following general convergence theorem. Note that the argument is essentially soft, since all of the serious estimates are implicitly already taken care of in [30].

Theorem 70. *Let M be a (sub)-Riemannian manifold with a smooth second-order operator L with non-negative-definite principal symbol and without zeroth-order term. Further, suppose that the diffusion generated by L , which we call X^0 , does not explode, and let P_q be the corresponding probability measure on Ω_M starting from q . Similarly, let P_q^h be the probability measures on Ω_M corresponding to a sequence of random walks X_t^h as above with $X_0^h = q$, and let L_h be the associated operators. Suppose that, for any $\phi \in C_0^\infty(M)$, we have that*

$$L_h \phi \rightarrow L \phi \quad \text{uniformly on compacts as } h \rightarrow 0, \quad (20)$$

and also suppose that the condition of Eq. (19) holds for the X_t^h . Then if $q_h \rightarrow q$ as $h \rightarrow 0$, we have that $P_{q_h}^h \rightarrow P_q$ as $h \rightarrow 0$.

Proof. First, suppose that M is compact. Then M can be realized as a compact (isometrically embedded) submanifold of \mathbb{R}^N via Nash embedding. This makes Ω_M a closed subset of $\Omega_{\mathbb{R}^N}$. Further, we claim that there is an increasing, continuous function $u : [0, \infty) \rightarrow [0, \infty)$ with $u(0) = 0$ such that, for $q, \tilde{q} \in M$,

$$d_{\mathbb{R}^N}(q, \tilde{q}) \leq d_M(q, \tilde{q}) \leq u(d_{\mathbb{R}^N}(q, \tilde{q})). \quad (21)$$

In particular, the topology on Ω_M induced as a subset of $\Omega_{\mathbb{R}^N}$ agrees with the original topology on Ω_M . To see that this is true, first observe that the inequality $d_{\mathbb{R}^N}(q, \tilde{q}) \leq d_M(q, \tilde{q})$ is immediate from the isometric embedding and the optimality of Euclidean geodesics in \mathbb{R}^N . Next, we let \bar{d}_M be the Riemannian distance on M (induced by the Riemannian extension metric of the sub-Riemannian metric). Then [21, Theorem 1.2] implies that every point of M has a neighborhood on which we can find a function u as above for which $d_M(q, \tilde{q}) \leq u(\bar{d}_M(q, \tilde{q}))$. One only needs to check that the constants in [21, Theorem 1.2] can be chosen uniformly, and this follows from the fact that, in [21, Lemma 1.1], the commutators X_{I_i} can be used on the entire neighborhood (assuming it's small enough) and then the function $\varphi(t_1, \dots, t_n)$ depends continuously on the base point. Since M is compact, this u can be chosen for all of M . Further, since the Riemannian distance \bar{d}_M and the Euclidean distance $d_{\mathbb{R}^N}$ are Lipschitz-comparable under the embedding, potentially multiplying u by a positive constant allows us to replace \bar{d}_M by $d_{\mathbb{R}^N}$. This establishes Eq. (21) (and in fact, u can be taken to be Hölder-continuous, though we don't need that here).

Next, there is no problem extending L to a smooth operator on all of \mathbb{R}^N with bounded coefficients (where we mean the coefficients used when writing L with respect to the Euclidean coordinates). We also extend the

family of random walks X_t^h to a family of Euclidean random walks \tilde{X}^h so that we have a random walk starting from any point of \mathbb{R}^N in such a way that the convergence of Eq. (20) and the condition of Eq. (19) still hold (with X_t^h replaced by \tilde{X}_t^h). Here \tilde{X}_t^h is not exactly an extension, since we interpolate between \tilde{X}_{ih}^h and $\tilde{X}_{(i+1)h}^h$ by a Euclidean geodesic. However, we still have that, if $X_0^h = \tilde{X}_0^h \in M$, then $X_{ih}^h = \tilde{X}_{ih}^h$ for all i , so that the underlying Markov chains on M agree. Further, it's easy to see that these extensions can be performed in such a way that the assumptions of Theorem 11.2.3 of [30] are satisfied (indeed, this just means some extra global assumptions on the random walks and operators, but since we start from a compact set, we can make L and \tilde{X}_t^h be anything we'd like outside of some large ball, just by using a bump function). We note that if f is a smooth function on \mathbb{R}^N obtained as an extension of a smooth function on M , then Lf and $L_h f$, evaluated at a point on M , depend only on the restriction $f|_M$, and also, the diffusion associated to the extended operator L on \mathbb{R}^N , started from a point on M , stays on M for all time (and, of course, agrees with the original diffusion on M). Thus the extension to \mathbb{R}^N is compatible with the original intrinsic formulation on M , and we are justified in (with a slight abuse of notation) continuing to use L for the extended operator, etc.

We let \tilde{P}^h be the measures on $\Omega_{\mathbb{R}^N}$ corresponding to the \tilde{X}^h . If we now apply Theorem 11.2.3 of [30] to the situation at hand, see that $\tilde{P}_{q_h}^h \rightarrow P_q$ as $h \rightarrow 0$ (as probability measures on $\Omega_{\mathbb{R}^N}$). It remains only to see that the same holds for $P_{q_h}^h$, that is, that whether or not we go from X_{ih} to $X_{(i+1)h}$ via a Euclidean geodesic or via the original M -geodesic doesn't matter in the limit.

We now pass to any discrete subsequence h_j such that $h_j \rightarrow 0$ as $j \rightarrow \infty$. However, for simplicity, we continue just to write h . By Theorem 1.3.1 of [30] and the fact that $\{\tilde{P}_{q_h}^h\}$ is precompact (indeed, it converges as $h \rightarrow 0$), we see that, for every $\rho > 0$ and $0 < T < \infty$,

$$\liminf_{\delta \searrow 0} \inf_h \tilde{P}_{q_h}^h \left[\sup_{\substack{0 \leq s \leq t \leq T \\ t-s \leq \delta}} d_{\mathbb{R}^N}(\tilde{X}_t^h, \tilde{X}_s^h) \leq \rho \right] = 1.$$

(Here the h on both the measure \tilde{P}^h and the corresponding random path \tilde{X}^h is somewhat redundant, but clear nonetheless. In a moment we will make better use of this notation.)

By assumption, the condition of Eq. (19) holds. Since M is compact and X_{ih}^h is a Markov chain, simply summing the probability of the distance exceeding ρ at each step of the walk gives that for any $0 < T < \infty$, any $\rho > 0$, and any $\alpha > 0$, there exists $h_0 > 0$ such that

$$\sup_{q \in M} P_q^h \left[\sup_{\substack{i: 0 \leq hi \leq T \\ s < h}} d_M(X_{ih}^h, X_{ih+s}^h) > \rho \right] \leq \frac{T+1}{h} \sup_{q \in M} P_q^h \left[\sup_{0 \leq s \leq h} d_M(X_0^h, X_s^h) > \rho \right] \leq (T+1)\alpha$$

whenever $h < h_0$. By Eq. (21), this also holds if we replace d_M with $d_{\mathbb{R}^N}$, perhaps by taking h smaller. Thus, for any $0 < T < \infty$, any $\rho > 0$, and any $\alpha > 0$, we have that

$$P_{q_h}^h \left[\sup_{\substack{i: 0 \leq hi \leq T \\ s < h}} d_{\mathbb{R}^N}(X_{ih}^h, X_{ih+s}^h) \leq \rho \right] > 1 - \alpha$$

for all small enough h .

In order to compare $\tilde{P}_{q_h}^h$ and $P_{q_h}^h$, it is convenient to realize them as push-forwards of a single measure on some probability space. In some sense, this is how we have described these random walks as being generated, in terms of a sequence of independent random variables which we use to draw the new cotangent vector at every step. However, it is more direct just to note that the paths \tilde{X}_t can be recovered from the paths X_t . More concretely, for any path ω on M that is piecewise geodesic with respect to the times $0, h, 2h, \dots$, let $F^h(\omega)$ be the path in \mathbb{R}^N that interpolates between ω_{ih} and $\omega_{(i+1)h}$ by the appropriate Euclidean geodesic for each i . Then F^h is defined on the support of $P_{q_h}^h$, and $\tilde{P}_{q_h}^h$ is the pushforward of $P_{q_h}^h$ under F^h for each h . So it is natural to view both X^h and \tilde{X}^h as random variables under $P_{q_h}^h$, with $\tilde{X}^h = F^h(X^h)$.

It now follows from the above (recall in particular that $\tilde{P}_{q_h}^h$ and $P_{q_h}^h$ have the same marginals on the sequence of times $0, h, 2h, \dots$) that, given $0 < T < \infty$, $\rho > 0$, and $\alpha > 0$, for all sufficiently small h we have

$$P_{q_h}^h \left[\sup_{0 \leq t \leq T} d_{\mathbb{R}^N}(\tilde{X}_t^h, X_t^h) \leq \rho \right] > 1 - \alpha.$$

In light of Eq. (18), we conclude that, for any $i \in \{1, 2, \dots\}$, and $\rho > 0$ and any $\alpha > 0$,

$$P_{q_h}^h \left[d_{\mathbb{R}^N}(\tilde{X}^h, X^h) \leq \rho + \frac{1}{2^i} \right] > 1 - \alpha$$

for all sufficiently large h .

Next, let Φ be a bounded, uniformly continuous function on $\Omega_{\mathbb{R}^N}$, and choose any $\delta > 0$. Then there exists $\eta > 0$ such that $|\Phi(\omega) - \Phi(\tilde{\omega})| < \delta$ whenever $d_{\mathbb{R}^N}(\omega, \tilde{\omega}) < \eta$. Now choose i large enough so that $\sum_{j=i}^{\infty} 1/2^j < \eta/2$ and let $\rho < \eta/2$ and $\alpha < \delta$ in the above. Then there exists $h_0 > 0$ such that

$$P_{q_h}^h [d_{\mathbb{R}^N}(\tilde{X}^h, X^h) > \eta] < \delta,$$

and thus

$$\mathbb{E}^{P_{q_h}^h} [|\Phi(\tilde{X}^h) - \Phi(X^h)|] < \delta + 2\delta\|\Phi\|_{\infty}$$

for $h < h_0$ (here $\|\Phi\|_{\infty}$ is the L^{∞} -norm of Φ). Further, using $\tilde{P}_{q_h}^h \rightarrow P_q$, after perhaps decreasing h_0 , we have

$$\left| \mathbb{E}^{P_{q_h}^h} [\Phi(\tilde{X}^h)] - \mathbb{E}^{P_q} [\Phi(X^0)] \right| < \delta$$

for $h < h_0$. Then linearity of expectation and the triangle inequality imply that

$$\left| \mathbb{E}^{P_{q_h}^h} [\Phi(X^h)] - \mathbb{E}^{P_q} [\Phi(X^0)] \right| < 2\delta(1 + \|\Phi\|_{\infty})$$

for $h < h_0$, and thus, since $\delta > 0$ is arbitrary

$$\lim_{h \rightarrow 0} \mathbb{E}^{P_{q_h}^h} [\Phi(X^h)] = \mathbb{E}^{P_q} [\Phi(X^0)].$$

This convergence holds for any bounded, uniformly continuous function Φ on $\Omega_{\mathbb{R}^N}$, and this is sufficient (by the portemanteau theorem), to show that $P_{q_h}^h \rightarrow P_q$. Here there is some ambiguity as to whether or not we think of $P_{q_h}^h$ and P_q as probability measures on Ω_M or $\Omega_{\mathbb{R}^N}$, but the point is that it doesn't matter. As measures on $\Omega_{\mathbb{R}^N}$, they are nonetheless supported on Ω_M . Further, weak convergence over $\Omega_{\mathbb{R}^N}$ implies weak convergence over Ω_M , and thus the convergence pulls back to Ω_M . Since this convergence holds for any discrete sequence $h_j \rightarrow 0$, it holds as $h \rightarrow 0$ via all positive reals, and this completes the proof in the case when M is compact.

If M is not compact, we proceed by exhaustion, using Lemma 11.1.1 of [30]. In this context, note that this lemma is stated for the case when Ω is $\Omega_{\mathbb{R}^N}$, but it is proved by a general and elementary method, and thus the statement and proof apply, as written, to the case when $\Omega = \Omega_M$.

Let $A_k \subset M$ be an exhaustion of M by compact subsets (with smooth boundary, possible by Sard's theorem). We assume that q and all of the q_h are contained in A_1 . If τ_k is the first hitting time of the complement of the interior of A_k , then τ_k is a non-decreasing sequence of lower semi-continuous stopping times that increases to ∞ for each $\omega \in \Omega_M$. Now let M_k be a sequence of compact Riemannian manifolds into which A_k can be isometrically included. (In other words, we truncate, in a geometrically reasonable way, M outside of A_k to get M_k .) For each k , extend L to a smooth operator L_k on M_k , and extend the random walk to M_k in such a way that the probability measures $P_{q_h}^{h,k}$ corresponding to the random walks on M_k converge to the the probability measure corresponding to the diffusion generated by L_k , which we denote Q_q^k . The previous argument for compact M makes it clear that this is possible.

By construction, the probability measure $P_{q_h}^{h,k}$ agrees with $P_{q_h}^h$ on \mathcal{M}_{τ_k} . Also, by the previous argument, for any $k \geq 1$ and any discrete sequence $h_j \rightarrow 0$, the family $\{P_{q_{h_j}}^{h_j,k} : j \geq 1\}$ converges to Q_q^k , and Q_q^k equals P_q on \mathcal{M}_{τ_k} . These are exactly the assumptions of Lemma 11.1.1 of [30], and so we conclude that $P_{q_{h_j}}^{h_j} \rightarrow P_q$ as $j \rightarrow \infty$. As before, since this holds for any discrete sequence $h_j \rightarrow 0$ we have that $P_{q_h}^h \rightarrow P_q$ as $h \rightarrow 0$. \square

Remark 26. Note that nowhere in the proof did we use that the path from $X_{h_i}^h$ to $X_{h_{(i+1)}}^h$ was a geodesic. Indeed, all that matters is that the path is continuous and satisfies the condition of Eq. (19) (and goes from $X_{h_i}^h$ to $X_{h_{(i+1)}}^h$, of course). Thus, Theorem 70 holds in slightly more generality, where we extend the class of random walks under consideration to include random walks where the interpolations from $X_{h_i}^h$ to $X_{h_{(i+1)}}^h$ are not necessarily geodesics. Perhaps the most natural such situation would be where each step of the random walk is given by flowing along the integral curve of some horizontal vector field for a small time (as opposed to traveling along a geodesic for a small time).

Finally, we return to the special case of the type of random walks discussed earlier in the paper. Then we can show that the convergence of a single step of the random walk, used to (first) define the microscopic Laplacian with respect to a splitting, also implies the convergence of the random walk to the natural limiting diffusion.

Theorem 71. *For a (sub)-Riemannian manifold M , consider a splitting $TM = \blacktriangle \oplus \mathbf{c}$ (if M is Riemannian, the splitting is necessarily trivial) and some choice of adapted measure $\{\mu_q\}_{q \in M}$, and let $L^{\mathbf{c}}$ be the associated microscopic Laplacian, as in Definition 51. Then there is a unique diffusion $X_t^{0,q}$ generated by $L^{\mathbf{c}}$ starting from any q . Let $X_t^{h,q}$ be the random walk, starting from q , determined by the parabolic scaling of the family $\{\mu_q\}$, as*

discussed above. Assume that $X_t^{0,q}$ does not explode, for any q . If $q_h \rightarrow q_0$ as $h \rightarrow 0$, then the random walks X_t^{h,q_h} converge to the diffusion X_t^{0,q_0} as $h \rightarrow 0$, in the sense that the corresponding probability measures on Ω_M converge weakly.

Proof. By Theorem 52, the operator L^c satisfies the assumptions of Theorem 70. Also, the existence and uniqueness of $X_t^{0,q}$ follows. Now let L_h be the operator associated to the random walks $X_t^{h,\cdot}$, as described above. Then for any $\phi \in C_0^\infty(M)$, we see that $L_h\phi \rightarrow L^c\phi$ uniformly (on all of M) as $h \rightarrow 0$, by Theorem 52 and the fact that $L_h\phi$, for all small h , has support in a fixed compact set by construction. Finally, we have already noted that the condition of Eq. (19) holds trivially for the random walks $X_t^{h,q}$. Thus we can apply Theorem 70, which gives the desired convergence of randoms walks. \square

Acknowledgments

The authors are grateful to Frédéric Jean for his advice on non-equiregular structures, and in particular, for explaining how to derive Eq. (21). We also thank Andrei Agrachev, Grégoire Charlot, Erlend Grong, and Jose Veloso for useful and stimulating discussions. We finally thank the Institut Henri Poincaré (Paris) for the hospitality during the Trimester “Geometry, Analysis and Dynamics on sub-Riemannian manifolds,” where most of this research has been carried out.

This research has been supported by the European Research Council, ERC StG 2009 “GeCoMethods,” contract number 239748, by the iCODE institute, research project of the Idex Paris-Saclay and by the EU FP7 project QUANT, grant agreement no. 297861.

References

- [1] A. Agrachev. Any sub-Riemannian metric has points of smoothness. *Dokl. Akad. Nauk*, 424(3):295–298, 2009.
- [2] A. Agrachev. Some open problems. In *Geometric control theory and sub-Riemannian geometry*, volume 5 of *Springer INdAM Ser.*, pages 1–13. Springer, Cham, 2014.
- [3] A. Agrachev, D. Barilari, and U. Boscain. Introduction to Riemannian and sub-Riemannian geometry (Lecture Notes), http://people.sissa.it/agrachev/agrachev_files/notes.html. 2012.
- [4] A. Agrachev, D. Barilari, and U. Boscain. On the Hausdorff volume in sub-Riemannian geometry. *Calc. Var. Partial Differential Equations*, 43(3-4):355–388, 2012.
- [5] A. Agrachev, D. Barilari, and L. Rizzi. The curvature: a variational approach. *ArXiv e-prints*, June 2013.
- [6] A. Agrachev, U. Boscain, J.-P. Gauthier, and F. Rossi. The intrinsic hypoelliptic Laplacian and its heat kernel on unimodular Lie groups. *J. Funct. Anal.*, 256(8):2621–2655, 2009.
- [7] A. Agrachev, A. Gentile, and A. Lerario. Geodesics and admissible-path spaces in Carnot Groups. *to appear on Geometry & Topology*.
- [8] D. Barilari and L. Rizzi. A formula for Popp’s volume in sub-Riemannian geometry. *Anal. Geom. Metr. Spaces*, 1:42–57, 2013.
- [9] F. Baudoin and N. Garofalo. Curvature-dimension inequalities and Ricci lower bounds for sub-Riemannian manifolds with transverse symmetries. *to appear in Journal of the European Mathematical Society*.
- [10] A. Bellaïche. The tangent space in sub-Riemannian geometry. In *Sub-Riemannian geometry*, volume 144 of *Progr. Math.*, pages 1–78. Birkhäuser, Basel, 1996.
- [11] E. Breuillard, P. Friz, and M. Huesmann. From random walks to rough paths. *Proc. Amer. Math. Soc.*, 137(10):3487–3496, 2009.
- [12] R. W. Brockett. Nonlinear control theory and differential geometry. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Warsaw, 1983)*, pages 1357–1368, Warsaw, 1984. PWN.
- [13] G. Charlot. Quasi-contact S-R metrics: normal form in \mathbf{R}^{2n} , wave front and caustic in \mathbf{R}^4 . *Acta Appl. Math.*, 74(3):217–263, 2002.
- [14] M. M. Diniz and J. M. M. Veloso. k -step sub-Riemannian manifold whose sub-Riemannian metric admits a canonical extension to a Riemannian metric. *J. Dyn. Control Syst.*, 16(4):517–538, 2010.
- [15] M. Gordina and T. Laetsch. Sub-Laplacians on sub-Riemannian manifolds. *ArXiv e-prints*, Nov. 2014.

- [16] M. Gordina and T. Laetsch. Weak Convergence to Brownian Motion on Sub-Riemannian Manifolds. *ArXiv e-prints*, Mar. 2014.
- [17] E. Grong and A. Thalmaier. Curvature-dimension inequalities on sub-riemannian manifolds obtained from riemannian foliations, part 1. 2014.
- [18] E. Grong and A. Thalmaier. Curvature-dimension inequalities on sub-riemannian manifolds obtained from riemannian foliations, part 2. 2014.
- [19] R. K. Hladky. Connections and curvature in sub-Riemannian geometry. *Houston J. Math.*, 38(4):1107–1134, 2012.
- [20] R. K. Hladky. Intrinsic complements of equiregular sub-Riemannian manifolds. *Geom. Dedicata*, 173:89–103, 2014.
- [21] F. Jean. *Control of Nonholonomic Systems: from Sub-Riemannian Geometry to Motion Planning*. Springer International Publishing, SpringerBriefs in Mathematics, 2014.
- [22] G. Lebeau and L. Michel. Semi-classical analysis of a random walk on a manifold. *Ann. Probab.*, 38(1):277–315, 2010.
- [23] G. Lebeau and L. Michel. Spectral analysis of hypoelliptic random walks. *Journal of the Institute of Mathematics of Jussieu*, FirstView:1–41, 2 2015.
- [24] A. Lerario and L. Rizzi. How many geodesics join two points on a contact sub-Riemannian manifold? *ArXiv e-prints*, May 2014.
- [25] R. Montgomery. *A tour of subriemannian geometries, their geodesics and applications*, volume 91 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2002.
- [26] M. A. Pinsky. Isotropic transport process on a Riemannian manifold. *Trans. Amer. Math. Soc.*, 218:353–360, 1976.
- [27] L. Rifford. *Sub-Riemannian geometry and optimal transport*. Springer, 2014.
- [28] R. S. Strichartz. Sub-Riemannian geometry. *J. Differential Geom.*, 24(2):221–263, 1986.
- [29] R. S. Strichartz. Corrections to: “Sub-Riemannian geometry” [*J. Differential Geom.* **24** (1986), no. 2, 221–263; MR0862049 (88b:53055)]. *J. Differential Geom.*, 30(2):595–596, 1989.
- [30] D. W. Stroock and S. R. S. Varadhan. *Multidimensional diffusion processes*. Classics in Mathematics. Springer-Verlag, Berlin, 2006. Reprint of the 1997 edition.