> **This is a recreation of the original chat by Codex using the data export supplied by the user.**

Global Chart and Analytic Manifolds

Conversation ID: 6a2bc31d-df00-83e8-9e51-32333c81985a

73 recorded user and assistant messages • Dates shown in Asia/Shanghai (UTC+08:00), with source timestamps.

Source: ChatGPT data export

Source: supplied ChatGPT data export dated 2026-10-25. Conversation URL: https://chatgpt.com/c/6a2bc31d-df00-83e8-9e51-32333c81985a
This PDF reproduces all 73 saved user questions and assistant text responses: 73 on the saved active conversation path.
Messages retain their original wording. Mathematics and Markdown are typeset; citation chips are reconstructed from exported reference metadata.
Message dates use the original creation timestamps, displayed in Asia/Shanghai (UTC+08:00); UTC source timestamps are also shown.
The exact source message text is embedded in this PDF as a UTF-8 transcript and JSON record. Original uploaded files and available cited file assets are also embedded; use a PDF reader with an attachments panel.
This is a record of the conversation, including any errors and later corrections in the original answers.
Message date range: 2026-06-12 16:28:40 (Asia/Shanghai, UTC+08:00) to 2026-06-16 11:28:14 (Asia/Shanghai, UTC+08:00)

Main conversation

1. User

2026-06-12 16:28:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:28:40.800000Z

if a finite dimensional manifold has a global chart to R^m, does that automatically make it an analytic riemannian manifold?

2. Assistant

2026-06-12 16:28:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:28:40.197476Z

Yes, **provided you are allowed to choose the analytic structure and the Riemannian metric using that chart**.

Suppose M has a global chart

$$\psi : M \to \mathbb{R}^m$$

or more generally $\psi : M \to U \subseteq \mathbb{R}^m$. Then:

1. **Analytic manifold structure:**
   The one-chart atlas $\{(M, \psi)\}$ makes M into a real analytic manifold. There are no nontrivial transition maps to check, so analyticity is automatic.
2. **Analytic Riemannian metric:**
   You can pull back the Euclidean metric on $\mathbb{R}^m$:

$$g_x(v, w) = \langle d\psi_x(v), d\psi_x(w) \rangle_{\mathbb{R}^m}.$$

In the coordinate chart ψ, the metric tensor is just the identity matrix:

$$g_{ij} = \delta_{ij},$$

so it is certainly analytic. Thus (M, g) is an analytic Riemannian manifold.

The main caveat is: if M already has some **given** smooth Riemannian metric h, then h need not be analytic in this chart. A global chart lets you **choose** an analytic Riemannian metric canonically by pulling back the Euclidean one, but it does not make every preexisting smooth metric analytic.

So the answer is:

$$\boxed{\text{A global chart lets you equip } M \text{ with an analytic manifold structure and an analytic Riemannian metric.}}$$

But

$$\boxed{\text{An arbitrary existing smooth Riemannian metric need not become analytic automatically.}}$$

3. User

2026-06-12 16:33:23 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:33:23.576000Z

what is a riemannian manifold?

4. Assistant

2026-06-12 16:33:22 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:33:22.313527Z

A **Riemannian manifold** is a smooth manifold M together with a smoothly varying inner product on each tangent space.

More explicitly, for every point $p \in M$, the tangent space T_pM is a vector space. A Riemannian metric assigns an inner product

$$g_p(\cdot,\cdot) : T_pM \times T_pM \to \mathbb{R}$$

to each tangent space, in a way that varies smoothly with p.

So a Riemannian manifold is a pair

$$(M, g),$$

where M is a smooth manifold and g is a Riemannian metric.

The point of g is that it lets you measure geometric quantities on M, such as:

$$\text{lengths of tangent vectors}, \quad \text{angles between tangent vectors}, \quad \text{lengths of curves}, \quad \text{distances between points}, \quad \text{volume}.$$

For example, ordinary Euclidean space $\mathbb{R}^m$ is a Riemannian manifold with the standard dot product

$$g_p(v, w) = v \cdot w.$$

A surface in $\mathbb{R}^3$, such as a sphere, also becomes a Riemannian manifold by using the dot product from $\mathbb{R}^3$ restricted to the tangent plane at each point.

---

5. User

2026-06-12 16:34:11 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:34:11.116000Z

what does it mean for the inner product to vary smoothly?

---

6. Assistant

2026-06-12 16:34:09 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:34:09.630429Z

It means that when you write the inner product in local coordinates, its coefficients are smooth functions of the point.

Suppose M is an m-dimensional smooth manifold, and on some coordinate chart you write tangent vectors using the coordinate basis

$$\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^m}.$$

The Riemannian metric g is determined by the functions

$$g_{ij}(x) = g_x\left(\frac{\partial}{\partial x^i}, \frac{\partial}{\partial x^j}\right).$$

Then saying that the inner product varies smoothly means exactly that each coefficient function

$$x \mapsto g_{ij}(x)$$

is a smooth function in the coordinate variables.

So if

$$v = \sum_i v^i \frac{\partial}{\partial x^i}, \qquad w = \sum_j w^j \frac{\partial}{\partial x^j},$$

then

$$g_x(v, w) = \sum_{i,j} g_{ij}(x) v^i w^j.$$

The dependence on x should be smooth.

For example, on $\mathbb{R}^2$, the metric

$$g_x(v, w) = e^{x_1}(v_1 w_1) + (1 + x_2^2)(v_2 w_2)$$

is smooth, because the coefficient functions e^{x_1} and $1 + x_2^2$ are smooth. In matrix form,

$$g(x) = \begin{pmatrix} e^{x_1} & 0 \\ 0 & 1 + x_2^2 \end{pmatrix}.$$

At each point x, this matrix defines an inner product on $T_x\mathbb{R}^2$, and the entries of the matrix change smoothly as x moves.

---

7. User

2026-06-12 16:38:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:38:21.822000Z

do we usually use the term 'analytic riemannian manifold' or 'analytic manifold' when the smooth structure is actually analytic?

---

8. Assistant

2026-06-12 16:38:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:38:20.413114Z

Usually:

“analytic manifold” means the **manifold structure** is analytic: the transition maps between charts are real analytic.

“analytic Riemannian manifold” usually means more: the manifold is real analytic **and** the Riemannian metric is analytic in analytic coordinates.

So the distinction is:

$$\text{analytic manifold} = M \text{ has an analytic atlas.}$$

$$\text{analytic Riemannian manifold} = (M, g) \text{ where } M \text{ is analytic and } g \text{ is analytic.}$$

If the smooth structure is analytic but the metric is only smooth, I would say:

> M is a real analytic manifold equipped with a smooth Riemannian metric.

I would **not** usually call that an analytic Riemannian manifold unless the metric coefficients g_{ij} are analytic functions in analytic charts.

In your earlier situation, if you have a global chart $\psi : M \to \mathbb{R}^m$ and pull back the Euclidean metric, then both the chart structure and the metric are analytic. So $(M, \psi^* g_{\text{Eucl}})$ is naturally a **real analytic Riemannian manifold**.

---

9. User

2026-06-12 16:43:57 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:43:57.396000Z

what does it mean for the metric to be analytic?

---

10. Assistant

2026-06-12 16:43:55 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:43:55.928571Z

It means the metric coefficients are **real analytic functions** in analytic local coordinates.

Suppose M is a real analytic manifold and g is a Riemannian metric. In a local analytic chart with coordinates

$$(x^1, \dots, x^m),$$

write

$$g_{ij}(x) = g_x\left(\frac{\partial}{\partial x^i}, \frac{\partial}{\partial x^j}\right).$$

Then g is called **analytic** if every coefficient function

$$x \mapsto g_{ij}(x)$$

is real analytic.

Equivalently, in coordinates,

$$g_x(v,w) = \sum_{i,j=1}^{m} g_{ij}(x) v^i w^j,$$

and the functions g_{ij} are analytic functions of x.

Here “analytic” means locally given by a convergent power series. So near each point x_0,

$$g_{ij}(x) = \sum_{\alpha} c_\alpha (x - x_0)^\alpha$$

for x sufficiently close to x_0.

For example, on $\mathbb{R}^2$,

$$g = e^x\,dx^2 + (1+y^2)\,dy^2$$

is analytic, because the coefficient functions e^x and $1+y^2$ are analytic. But

$$g = \left(1 + e^{-1/x^2}\right) dx^2 + dy^2$$

with the usual smooth extension at $x = 0$, is generally smooth but **not analytic** at $x = 0$, because e^{-1/x^2} is smooth but not real analytic there.

---

11. User

2026-06-12 16:45:44 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:45:44.229000Z

so is this the same as the inner product being analytic?

---

12. Assistant

2026-06-12 16:45:43 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:45:43.158090Z

Yes — **that is exactly the idea**, with one small precision.

A Riemannian metric is a choice of inner product

$$g_x(\cdot,\cdot)$$

on each tangent space T_xM. Saying the metric is **analytic** means that this family of inner products depends analytically on the point x.

In local analytic coordinates, this means that the matrix entries

$$g_{ij}(x) = g_x\left(\frac{\partial}{\partial x^i}, \frac{\partial}{\partial x^j}\right)$$

are analytic functions of x.

So informally:

$$\boxed{\text{analytic metric} = \text{the tangent-space inner product varies analytically from point to point.}}$$

The phrase “the inner product is analytic” is understandable, but slightly imprecise because each individual g_x is just a bilinear form on one vector space. The analyticity refers to how the whole family $\{g_x\}_{x \in M}$ changes as x moves.

---

13. User

2026-06-12 16:48:35 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:48:35.629000Z

given an inner product on the manifold, how do we compute the distance between two points?

14. Assistant

2026-06-12 16:48:34 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:48:34.307648Z

Given a Riemannian metric g, the distance between two points $p, q \in M$ is defined by minimizing the lengths of all smooth, or piecewise smooth, paths from p to q.

If

$$\gamma : [a, b] \to M$$

is a path with $\gamma(a) = p$ and $\gamma(b) = q$, then its length is

$$L(\gamma) = \int_a^b \sqrt{g_{\gamma(t)}(\gamma'(t), \gamma'(t))}\, dt.$$

Here $\gamma'(t) \in T_{\gamma(t)}M$, and the metric $g_{\gamma(t)}$ measures the squared length of that tangent vector.

Then the Riemannian distance is

$$d(p, q) = \inf_{\gamma} L(\gamma),$$

where the infimum is over all piecewise smooth curves γ from p to q.

So the metric gives you **infinitesimal lengths**, and the distance between two points is the shortest possible accumulated length along a curve connecting them.

For example, in $\mathbb{R}^m$ with the usual Euclidean metric,

$$g_x(v, w) = v \cdot w,$$

we get

$$L(\gamma) = \int_a^b \|\gamma'(t)\|\, dt,$$

and minimizing over curves gives the usual Euclidean distance

$$d(p, q) = \|p - q\|.$$

On a sphere, the same definition gives distance along the surface, so the shortest paths are arcs of great circles rather than straight lines through the ambient space.

---

15. User

2026-06-12 16:50:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:50:21.121000Z

do we really need smoothness for this? Isnt C^1 enough to calculate distance in this fashion?

---

16. Assistant

2026-06-12 16:50:19 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:50:19.760129Z

Yes. ** C^1 is more than enough** if all you want is to define lengths of curves and the induced distance.

For a path $\gamma : [a, b] \to M$, the length formula is

$$L(\gamma) = \int_a^b \sqrt{g_{\gamma(t)}(\gamma'(t), \gamma'(t))}\, dt.$$

To make this integral well behaved for C^1 or piecewise C^1 curves, it is enough that the metric coefficients $g_{ij}(x)$ vary **continuously** in local coordinates. Then the integrand is continuous, hence integrable.

So the hierarchy is roughly:

$$\text{continuous metric} \quad \Rightarrow \quad \text{lengths and distances make sense.}$$

$$C^1 \text{ metric} \quad \Rightarrow \quad \text{lengths and distances make sense, plus some first-order differential structure.}$$

$$C^\infty \text{ metric} \quad \Rightarrow \quad \text{standard smooth Riemannian geometry.}$$

Smoothness is not needed just to define the Riemannian distance. It is usually assumed because one wants more geometry: Levi-Civita connection, geodesic equations, curvature, smooth dependence of geodesics, normal coordinates, comparison theorems, etc.

For example, if M is a smooth manifold and g is only a **continuous positive-definite** inner product on each tangent space, people may call this a **continuous Riemannian metric** or a C^0 Riemannian metric. The distance is still defined by

$$d_g(p,q) = \inf_\gamma L_g(\gamma).$$

So yes: for the distance formula alone, C^1 is plenty; even C^0 is usually enough.

---

17. User

2026-06-12 16:52:28 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:52:28.381000Z

I agree that C^0, i.e., continuity of the family of inner products is enough, but we would need to assume that the manifold is at least C^1 in order to be able to take a derivative of the curves, or to even be able to define C^1 curves, right?

---

18. Assistant

2026-06-12 16:52:27 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:52:27.087934Z

Yes, exactly. In the usual coordinate-free Riemannian setup, the manifold should be at least a ** C^1 manifold**.

Reason: to define

$$\gamma'(t) \in T_{\gamma(t)}M,$$

you need a notion of differentiable curve and tangent vector. That requires the coordinate transition maps to be differentiable, at least C^1. Otherwise the derivative of a curve may depend on the chosen chart.

So the minimal usual setup for defining Riemannian length is something like:

$$M \text{ is a } C^1 \text{ manifold},$$

and

$$g = \{g_x\}_{x \in M}$$

is a **continuous** positive-definite inner product on the tangent bundle TM.

Then for a piecewise C^1 curve γ,

$$L_g(\gamma) = \int_a^b \sqrt{g_{\gamma(t)}(\gamma'(t), \gamma'(t))}\, dt$$

is well-defined, and the induced distance is

$$d_g(p,q) = \inf_\gamma L_g(\gamma).$$

So yes:

$$\boxed{C^1 \text{ structure on } M + \text{continuous Riemannian metric } g \text{ is enough to define distance.}}$$

Smoothness of M and g is needed for the fuller apparatus of Riemannian geometry—connections, curvature, geodesic equations, etc.—but not merely for defining lengths and distances.

19. User

2026-06-12 16:54:05 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:54:05.819000Z

is a riemannian metric a metric induced by a family of continuously varying inner products, or does it have to be smoothly varying?

20. Assistant

2026-06-12 16:54:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T08:54:04.348661Z

In standard differential geometry, a **Riemannian metric** usually means a **smoothly varying** family of inner products on the tangent spaces.

So if M is a smooth manifold, a Riemannian metric is a smooth section

$$g \in \Gamma^\infty(S^2T^*M)$$

such that each

$$g_p : T_pM \times T_pM \to \mathbb{R}$$

is an inner product.

Equivalently, in smooth local coordinates,

$$g = \sum_{i,j} g_{ij}(x)\, dx^i dx^j$$

with the matrix $(g_{ij}(x))$ symmetric positive definite and the coefficient functions g_{ij} smooth.

That said, people also study weaker versions:

$$C^0 \text{ Riemannian metric}$$

means the inner products vary **continuously**. This is enough to define lengths of C^1 curves and the induced distance.

$$C^k \text{ Riemannian metric}$$

means the coefficients g_{ij} are C^k.

$$\text{analytic Riemannian metric}$$

means the coefficients g_{ij} are real analytic in analytic coordinates.

So the terminology is:

$$\boxed{\text{“Riemannian metric” usually means smooth unless otherwise specified.}}$$

If the metric is only continuous, it is better to say **continuous Riemannian metric** or ** C^0 Riemannian metric**.

21. User

2026-06-12 17:11:01 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:11:01.397000Z

Is a quotient manifold of an analytic riemannian manifold again an analytic riemannian manifold?

22. Assistant

2026-06-12 17:11:00 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:11:00.224988Z

Not automatically. There are two separate issues:

1. **Is the quotient an analytic manifold?**
2. **Does the Riemannian metric descend to the quotient?**

Suppose (M, g) is an analytic Riemannian manifold and a group Γ acts on M. If Γ acts **freely, properly discontinuously, and analytically**, then the quotient

$$M/\Gamma$$

is naturally an analytic manifold.

But for the Riemannian metric g to descend to M/Γ, the action also needs to preserve g. That is, each $\gamma \in \Gamma$ should act by an isometry:

$$\gamma^* g = g.$$

If this holds, then there is a unique Riemannian metric $\bar{g}$ on M/Γ such that the quotient map

$$\pi : M \to M/\Gamma$$

is locally an isometry. In that case, $(M/\Gamma, \bar{g})$ is again an analytic Riemannian manifold.

So the clean statement is:

$$\boxed{\text{A quotient of an analytic Riemannian manifold by a free, proper, analytic, isometric action is again an analytic Riemannian manifold.}}$$

Without the isometry assumption, the metric generally does **not** descend.

For example, if $x \sim x + 1$ on $\mathbb{R}$, then the Euclidean metric descends to the circle $\mathbb{R}/\mathbb{Z}$. But if the identification distorted the metric, then the inner product assigned to an equivalence class could depend on which representative you choose, so it would not be well-defined.

In nilmanifold language: if G/Γ is formed from an analytic Lie group G and a discrete subgroup Γ, then G/Γ is naturally an analytic manifold. A Riemannian metric on G descends to G/Γ if it is invariant under the Γ-action used to form the quotient.

---

23. User

2026-06-12 17:13:41 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:13:41.849000Z

indeed, I wanted to apply this to the nilmanifold setting. So we pull back the euclidean metric from R^m to G using a malcev basis map psi_mathcal{X} and this gives an analytic riemannian metric on G. However, I don't think that this metric is right invariant? I guess we have to just do some pullback at the identity and then spread it around using the group operation to create right invariance, right? Can you clarify this picture for me.

---

24. Assistant

2026-06-12 17:13:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:13:40.689162Z

The key distinction is between two different metrics on G: the coordinate-pullback metric from Mal'cev coordinates, and a group-invariant metric built from an inner product on $\mathfrak{g}$. They usually do not coincide.

---

25. Assistant

2026-06-12 17:13:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:13:40.689162Z

One immediate consequence: for the quotient G/Γ, the metric on G has to be invariant under the right Γ-action. A fully right-invariant metric on G guarantees that.

---

26. Assistant

2026-06-12 17:13:40 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:13:40.689162Z

Yes, your picture is right.

There are **two different constructions** one can make from a Mal'cev coordinate map

$$\psi_{\mathcal{X}} : G \to \mathbb{R}^m.$$

They agree at the identity, but they generally give different metrics away from the identity.

1. Coordinate-pullback metric

You can pull back the Euclidean metric from $\mathbb{R}^m$ to G:

$$g^{\mathrm{coord}}_x(v,w) = \langle d\psi_{\mathcal{X}}|_x(v), d\psi_{\mathcal{X}}|_x(w)\rangle_{\mathbb{R}^m}.$$

This is an **analytic Riemannian metric** on G, because $\psi_{\mathcal{X}}$ is an analytic global chart.

But this metric is generally **not right-invariant**. Right-invariance would require

$$g^{\mathrm{coord}}_{xh}\big(dR_h v, dR_h w\big) = g^{\mathrm{coord}}_x(v,w)$$

for all $x, h \in G$. In coordinates, this would say that the coordinate expression of right translation is an isometry of Euclidean space. But in Mal'cev coordinates, right translations are typically polynomial maps, not Euclidean isometries.

For example, in a Heisenberg-type group, right translation introduces terms like

$$z \mapsto z + c + xb,$$

so its derivative is not usually an orthogonal matrix. Thus the coordinate-pullback metric is analytic, but not right-invariant.

---

2. Right-invariant metric induced by the Mal'cev basis

To get a right-invariant metric, you should only use the Mal'cev chart **at the identity**.

The differential

$$d\psi_{\mathcal{X}}|_e : T_eG = \mathfrak{g} \to \mathbb{R}^m$$

lets you define an inner product on the Lie algebra:

$$\langle Y, Z\rangle_e = \langle d\psi_{\mathcal{X}}|_e(Y), d\psi_{\mathcal{X}}|_e(Z)\rangle_{\mathbb{R}^m}.$$

Equivalently, if $\mathcal{X} = (X_1, \dots, X_m)$, this usually means declaring the Mal'cev basis vectors X_i to be orthonormal.

Then spread this inner product around G by right translations:

$$g^R_x(v,w) = \langle dR_{x^{-1}}v, dR_{x^{-1}}w\rangle_e.$$

This metric is right-invariant by construction. Indeed, for $h \in G$,

$$g^R_{xh}(dR_h v, dR_h w) = \big\langle dR_{(xh)^{-1}}dR_h v, dR_{(xh)^{-1}}dR_h w\big\rangle_e = \langle dR_{x^{-1}}v, dR_{x^{-1}}w\rangle_e.$$

So

$$g^R_{xh}(dR_h v, dR_h w) = g^R_x(v,w).$$

Because the group operations are analytic, this right-invariant metric is also analytic.

Passing to the nilmanifold G/Γ

For the quotient map

$$\pi : G \to G/\Gamma,$$

the equivalence relation is

$$x \sim x\gamma, \qquad \gamma \in \Gamma.$$

So to descend a metric from G to G/Γ, the metric on G must be invariant under right translations by elements of Γ. A fully right-invariant metric on G certainly satisfies this.

Therefore g^R descends to an analytic Riemannian metric $\bar{g}$ on G/Γ, characterized by saying that π is locally an isometry.

So the clean picture is:

$$\boxed{\text{Mal'cev coordinates give an inner product on } \mathfrak{g}.}$$

$$\boxed{\text{Right-translating that inner product gives a right-invariant analytic metric on } G.}$$

$$\boxed{\text{Because it is right-invariant, it descends to an analytic metric on } G/\Gamma.}$$

One more useful point: the Riemannian volume measure of this right-invariant metric on G is a right Haar measure. Since nilpotent Lie groups are unimodular, this is also a left Haar measure. After descending to compact G/Γ and normalizing, it gives the usual Haar probability measure on the nilmanifold.

---

27. User

2026-06-12 17:28:18 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:28:18.947000Z

if we worked with cocompact lattices in connected simply connected finite dimensional solvable lie groups, would everything regarding the riemannian metrics work out in the same fashion? I think this just boils down to whether or not the group operation is still analytic if we only assume solvable instead of nilpotent.

---

28. Assistant

2026-06-12 17:28:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:28:17.903012Z

Yes, the metric construction itself is not special to nilpotent groups. The main nilpotent-specific convenience is the existence of global polynomial/exponential Mal'cev coordinates, not the analytic/right-invariant Riemannian metric part.

---

29. Assistant

2026-06-12 17:28:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T09:28:17.903012Z

Yes. For this Riemannian-metric discussion, **nilpotent is not essential**. The same construction works for any finite-dimensional Lie group, in particular for a connected simply connected finite-dimensional solvable Lie group G.

The key facts are:

1. Every finite-dimensional real Lie group has a compatible **real analytic** structure.
2. In that analytic structure, the group operations

$$G \times G \to G, \qquad (g,h) \mapsto gh,$$

and

$$G \to G, \qquad g \mapsto g^{-1},$$

are analytic. 3. Therefore right translations

$$R_h : G \to G, \qquad x \mapsto xh$$

are analytic diffeomorphisms.

So if $\Gamma \leq G$ is a cocompact lattice and you form

$$G/\Gamma,$$

then the right action of Γ on G is analytic, free, and properly discontinuous. Hence G/Γ is naturally an analytic manifold.

Now choose any inner product on the Lie algebra

$$\mathfrak{g} = T_eG.$$

Define a right-invariant Riemannian metric on G by

$$g_x(v,w) = \langle dR_{x^{-1}}v, dR_{x^{-1}}w \rangle_{\mathfrak{g}}.$$

Because right translation is analytic, this is an **analytic Riemannian metric** on G. Because it is right-invariant, for every $\gamma \in \Gamma$,

$$R_\gamma^* g = g.$$

Therefore g descends to an analytic Riemannian metric $\bar{g}$ on

$$G/\Gamma.$$

So the clean statement is:

$$\boxed{\text{If } G \text{ is a finite-dimensional Lie group and } \Gamma \leq G \text{ is discrete, then any right-invariant analytic metric on } G \text{ descends to } G/\Gamma.}$$

For a cocompact lattice in a connected simply connected solvable Lie group, yes, everything works the same way.

The nilpotent case has extra conveniences: the exponential map is a global analytic diffeomorphism, and Mal'cev coordinates often give very explicit polynomial coordinate formulas. For general solvable groups, the exponential map need not be a global diffeomorphism, and the coordinate formulas need not be polynomial. But the analytic Lie group structure and the right-invariant metric construction still work.

One more useful point: the volume form of a right-invariant Riemannian metric on G is a right Haar measure. A Lie group admitting a lattice is unimodular, so for your cocompact lattice situation this is also left Haar. Thus, after descending and normalizing on compact G/Γ, the Riemannian volume measure is the usual Haar probability measure on the solvmanifold.

---

30. User

2026-06-12 18:49:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T10:49:20.209000Z

Before Theorem 3.1 the author describes the various quantities related to R. What is R? What is R_{ijkl}?

> **Attachments recorded in source:**
>
> - TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf

---

31. Assistant

2026-06-12 18:50:06 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T10:50:06.429362Z

In Gray's notation, R is the **Riemannian curvature operator/tensor**.

Earlier in Section 2, he defines the curvature operator by

$$R_{XY} = \nabla_{[X,Y]} - [\nabla_X, \nabla_Y].$$

So for vector fields X, Y, R_{XY} is an operator acting on another vector field Z. In more expanded notation,

$$R_{XY} Z$$

is the curvature applied to X, Y, Z. Be careful: this sign convention is the negative of one common modern convention.

The components R_{ijkl} are defined using an orthonormal normal-coordinate frame $X_1, \ldots, X_n$ at the point m. Gray writes

$$\langle R_{X_i X_j} X_k, X_\ell \rangle = R_{ijk\ell}.$$

So $R_{ijk\ell}$ is just the (i, j, k, ℓ)-component of the curvature tensor in that frame. Equivalently,

$$R_{ijk\ell} = g\left(R(X_i, X_j) X_k, X_\ell\right),$$

with Gray's sign convention for R. On page 5 he introduces this notation, and just before Theorem 3.1 on page 9 he uses it to define curvature invariants such as $\|R\|^2$, $\|\rho(R)\|^2$, and $\tau(R)$. [TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf]

The other quantities are then:

$$R_{ij} = \sum_{k=1}^{n} R_{ijkk},$$

the Ricci curvature components in his convention,

$$\tau(R) = \sum_{i=1}^{n} R_{ii},$$

the scalar curvature,

$$\|R\|^2 = \sum_{i,j,k,\ell=1}^{n} R_{ijk\ell}^2,$$

the squared norm of the full curvature tensor, and

$$\|\rho(R)\|^2 = \sum_{i,j=1}^{n} R_{ij}^2,$$

the squared norm of the Ricci tensor. The term ΔR in Theorem 3.1 means the Laplacian of the scalar curvature, i.e.

$$\Delta R = \sum_{i} \nabla_{ii}^2 \tau(R).$$

So: R is the full curvature tensor/operator, and R_{ijkl} are its coordinate components in a normal orthonormal frame at the point under consideration.

---

32. User

2026-06-12 19:01:42 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:01:42.342000Z

In my nilmanifold setting, if I have a fixed ambient nilmanifold G/Gamma, and then some rational subnilmanifolds whose mal'cev bases mathcal{X}_i' have each element being a M-rational combination of elements of the original Mal'cev basis mathcal{X}, then can we bound tau(R) and all these other functions of R by a single function of M? Maybe it can be M, the real dimension m of the outer nilmanifold, and the nilpotence step s, but that's all, I don't want any dependence on the subnilmanifolds other than M.

---

33. Assistant

2026-06-12 19:01:43 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:01:43.464513Z

Yes, **with one important qualification**: the bound also depends on the ambient Lie bracket in the original Mal'cev basis, unless that bracket-complexity is already being included in your M-notation.

Let

$$\mathcal{X} = (X_1, \dots, X_m)$$

be the fixed Mal'cev basis of $\mathfrak{g}$, and write

$$[X_a, X_b] = \sum_{c=1}^{m} c_{ab}^c X_c.$$

For a rational subnilmanifold G'/Γ', suppose its Mal'cev basis is

$$\mathcal{X}' = (Y_1, \dots, Y_r), \qquad r \le m,$$

with

$$Y_i = \sum_{a=1}^{m} q_{ia} X_a,$$

where the q_{ia} are M-rational, say with numerators and denominators bounded by M.

Now define the right-invariant metric on G' by declaring $Y_1, \dots, Y_r$ orthonormal at the identity and translating by right multiplication. This descends to G'/Γ'.

In the Y-basis, write

$$[Y_i, Y_j] = \sum_{k=1}^{r} d_{ij}^k Y_k.$$

Because the Y_i's are M-rational combinations of the fixed ambient X_a's, the coefficients d_{ij}^k are bounded by some function of

$$M, \quad m, \quad s, \quad \max_{a,b,c} |c_{ab}^c|.$$

This is just linear algebra plus the fact that the bracket coefficients in the ambient basis are fixed. Since the dimension is bounded by m, Cramer's rule gives a uniform bound once the rational height of the basis vectors is bounded.

Then the Levi-Civita connection coefficients are algebraic expressions in the d_{ij}^k. For example, up to the right/left invariant sign convention, the Koszul formula gives

$$2\langle \nabla_{Y_i} Y_j, Y_k \rangle = \langle [Y_i, Y_j], Y_k \rangle - \langle [Y_j, Y_k], Y_i \rangle + \langle [Y_k, Y_i], Y_j \rangle.$$

So the connection coefficients are bounded by the same kind of function.

The curvature components

$$R_{ijk\ell} = \langle R_{Y_i Y_j} Y_k, Y_\ell \rangle$$

are then polynomial expressions of degree at most 2 in the structure constants d_{ij}^k. Hence

$$|R_{ijk\ell}| \le C(M, m, s, \mathcal{X}).$$

Once the individual curvature components are bounded, the scalar curvature, Ricci norm, full curvature norm, and the other Gray-type invariants are bounded too. In Gray's notation, the quantities appearing before Theorem 3.1 include

$$\tau(R) = \sum_i R_{ii}, \qquad \|R\|^2 = \sum_{i,j,k,\ell} R_{ijk\ell}^2, \qquad \|\rho(R)\|^2 = \sum_{i,j} R_{ij}^2,$$

and these are finite sums over at most m indices, so they are bounded by the same type of uniform function. [TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf]

Moreover, because the metric on G' is right-invariant, the scalar curvature is constant on G' and hence also on G'/Γ'. Therefore

$$\Delta \tau(R) = 0.$$

So for the quantities in Gray's small-ball expansion, you get a uniform bound of the form

$$|\tau(R)| + \|R\|^2 + \|\rho(R)\|^2 + |\Delta \tau(R)| \le C(M, m, s, \mathcal{X}).$$

Since the ambient nilmanifold and ambient Mal'cev basis are fixed, it is reasonable in your setting to abbreviate this as

$$O_{M,m,s}(1),$$

provided your convention allows constants to depend on the fixed ambient nilmanifold.

The caveat is that one cannot have a bound depending only on M, m, s **uniformly over all ambient nilmanifolds**. For example, take the three-dimensional Heisenberg Lie algebra with orthonormal basis X, Y, Z and

$$[X, Y] = NZ.$$

This is still 2-step nilpotent and 3-dimensional, but the scalar curvature of the corresponding invariant metric grows like $-N^2$. Taking the whole group as the "subnilmanifold" gives $M = 1$, but the curvature is unbounded as $N \to \infty$. So some dependence on the ambient bracket complexity is unavoidable unless it has already been normalized into M.

---

34. User

2026-06-12 19:10:37 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:10:37.697000Z

and this function C, can we take it to be a polynomial function?

---

35. Assistant

2026-06-12 19:10:38 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:10:38.871264Z

Yes — **assuming "M-rational" means rational coordinates of height $\le M$**, you can take the curvature bound to be **polynomial in M**, with exponent depending only on the ambient dimension m — and with constants depending on the fixed ambient Mal'cev basis.

More precisely, let

$$[X_a, X_b] = \sum_c c_{ab}^c X_c$$

be the structure constants of the fixed ambient Lie algebra in the fixed Mal'cev basis $\mathcal{X}$. Let

$$Y_i = \sum_{a=1}^m q_{ai} X_a, \qquad i = 1, \dots, r,$$

be the Mal'cev basis of a rational subalgebra, with each $q_{ai} \in \mathbb{Q}$ of height $\leq M$. Declare $Y_1, \dots, Y_r$ orthonormal and spread this inner product by right translation on G'.

Then in the Y-basis,

$$[Y_i, Y_j] = \sum_k d_{ij}^k Y_k.$$

The coefficients d_{ij}^k are bounded polynomially in M. The reason is:

$$[Y_i, Y_j] = \sum_{a,b,c} q_{ai} q_{bj} c_{ab}^c X_c,$$

so its ambient X-coordinates are polynomially bounded in M. To rewrite this vector back in the Y-basis, one inverts an $r \times r$ rational minor of the coordinate matrix (q_{ai}). Since the entries have rational height $\leq M$ and $r \leq m$, any nonzero determinant is bounded below by $M^{-O_m(1)}$. Thus the inverse minor has entries bounded by $M^{O_m(1)}$. Therefore

$$|d_{ij}^k| \leq C_{\mathcal{X},m} M^{A_m}.$$

The Levi-Civita connection coefficients for a right-invariant metric are linear expressions in the d_{ij}^k, and the curvature components are quadratic expressions in the d_{ij}^k. Hence

$$|R_{ijk\ell}| \leq C_{\mathcal{X},m} M^{A_m}.$$

After increasing A_m, the scalar curvature, Ricci norm, and full curvature norm satisfy a polynomial bound of the form

$$|\tau(R)| + \|\rho(R)\|^2 + \|R\|^2 \leq C_{\mathcal{X},m} M^{A_m}.$$

In Gray's notation, these are exactly the curvature quantities appearing before Theorem 3.1, where $R_{ijk\ell}$ are the curvature tensor components, $\tau(R)$ is scalar curvature, $\|\rho(R)\|^2$ is the squared Ricci norm, and $\|R\|^2$ is the squared full curvature norm. [TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf]

Also, since the metric is right-invariant, these scalar curvature invariants are constant on G' and descend to constants on G'/Γ'. Therefore

$$\Delta\tau(R) = 0.$$

So for your purposes, yes, you can write something like

$$|\tau(R)| + \|\rho(R)\|^2 + \|R\|^2 + |\Delta\tau(R)| \leq C_{\mathcal{X},m} M^{A_m}.$$

If the ambient nilmanifold and original Mal'cev basis are fixed once and for all, you can abbreviate this as

$$O_m(M^{A_m})$$

or

$$M^{O_m(1)}.$$

If you want to keep the nilpotence step visible, you can write

$$M^{O_{m,s}(1)},$$

although for this curvature bound the dependence on s is not really essential once m is fixed. The dependence on the ambient Lie bracket is essential unless your ambient basis has already been normalized to have uniformly bounded structure constants.

36. User

2026-06-12 19:29:29 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:29:29.278000Z

how do we calculate the curvature and all of these constants associated to R in the original nilmanifold? Can we bound them by the rationality of the Mal'cev basis?

---

37. Assistant

2026-06-12 19:30:36 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:30:36.162929Z

Yes. For the **original nilmanifold** G/Γ, once you choose the right-invariant metric obtained by declaring the original Mal'cev basis

$$\mathcal{X} = (X_1, \dots, X_m)$$

orthonormal at the identity, all the curvature quantities are computed purely from the **Lie bracket structure constants** of $\mathfrak{g}$ in that basis.

Write

$$[X_i, X_j] = \sum_{k=1}^{m} c_{ij}^k X_k.$$

Let E_i be the corresponding right-invariant vector fields on G. Then $E_1, \dots, E_m$ is a global orthonormal frame for the right-invariant metric. Because they are right-invariant,

$$[E_i, E_j] = -\sum_k c_{ij}^k E_k.$$

The minus sign is just the usual right-invariant versus left-invariant convention, and it does not matter for size bounds.

The Levi-Civita connection coefficients are given by the Koszul formula. If

$$\nabla_{E_i} E_j = \sum_k \Gamma_{ij}^k E_k,$$

then

$$2\Gamma_{ij}^k = \langle [E_i, E_j], E_k \rangle - \langle [E_j, E_k], E_i \rangle + \langle [E_k, E_i], E_j \rangle.$$

So every Γ_{ij}^k is a linear expression in the structure constants c_{ij}^k.

Then the curvature components are computed from

$$R(E_i, E_j)E_k$$

by the usual formula. With one common convention,

$$R(E_i, E_j)E_k = \nabla_{E_i}\nabla_{E_j}E_k - \nabla_{E_j}\nabla_{E_i}E_k - \nabla_{[E_i,E_j]}E_k.$$

Since the Γ_{ij}^k are constant in the invariant frame, this gives

$$R_{ijk}^{\ell} = \sum_p \left(\Gamma_{jk}^p \Gamma_{ip}^{\ell} - \Gamma_{ik}^p \Gamma_{jp}^{\ell} - a_{ij}^p \Gamma_{pk}^{\ell} \right),$$

where

$$[E_i, E_j] = \sum_p a_{ij}^p E_p.$$

Thus each curvature component $R_{ijk\ell}$ is a **quadratic polynomial** in the structure constants c_{ij}^k.

Gray's notation is

$$R_{ijk\ell} = \langle R_{E_i E_j} E_k, E_\ell \rangle,$$

and he defines the Ricci components by

$$R_{ij} = \sum_{k=1}^{m} R_{ijik}.$$

Then

$$\tau(R) = \sum_i R_{ii}, \qquad \|R\|^2 = \sum_{i,j,k,\ell} R_{ijk\ell}^2, \qquad \|\rho(R)\|^2 = \sum_{i,j} R_{ij}^2.$$

These are the quantities appearing before Theorem 3.1 in Gray's paper.
[TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf]

So if

$$B := \max_{i,j,k} |c_{ij}^k|,$$

then very crudely,

$$|\Gamma_{ij}^k| \le C_m B,$$

$$|R_{ijk\ell}| \le C_m B^2,$$

and hence

$$|\tau(R)| \le C_m B^2,$$

$$\|R\|^2 + \|\rho(R)\|^2 \le C_m B^4.$$

Also, because the metric is right-invariant, the scalar curvature is constant on G, hence also on G/Γ. Therefore

$$\Delta\tau(R) = 0.$$

So yes: the curvature invariants are bounded by the bracket complexity of the original Mal'cev basis.

If your Mal'cev basis is M_0-rational in the quantitative sense that the structure constants c_{ij}^k have rational height at most M_0, then in particular

$$|c_{ij}^k| \le M_0,$$

and therefore

$$|\tau(R)| + \|R\|^2 + \|\rho(R)\|^2 + |\Delta\tau(R)| \le C_m M_0^{A_m}$$

for some exponent A_m. In fact, from the crude estimates above one can take something like $A_m = 4$ for the squared-norm quantities, after absorbing dimension-dependent factors into C_m.

The lattice Γ itself does **not** affect these local curvature quantities once the metric on G is fixed. It affects global geometry, such as injectivity radius, but not the pointwise curvature tensor of the descended metric on G/Γ.

---

38. User

2026-06-12 19:33:54 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:33:54.724000Z

do green and tao show that the d_mathcal{X} metric is locally equivalent to the right invariant riemannian metric induced by pulling back the l^2 metric at the identity? I think they analyzed the pull back (not necessarily right invariant) of the l^infty metric, but not this other metric, right?

---

39. Assistant

2026-06-12 19:34:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:34:17.838724Z

Yes — Green and Tao **do explicitly mention exactly this right-invariant Riemannian metric**, but their actual working metric is the artificial Mal'cev-coordinate metric $d_{\mathcal{X}}$, built from the ℓ^∞-norm.

They define

$$d_{\mathcal{X}}(x, y)$$

to be the largest metric satisfying

$$d_{\mathcal{X}}(x,y) \le |\psi_{\mathcal{X}}(xy^{-1})|_\infty.$$

Equivalently, they give it as a path-type metric using chains and increments measured by $|\psi_{\mathcal{X}}(\cdot)|_\infty$. This metric is right-invariant by construction. Annals of Mathematics

But in the footnote immediately after Definition 2.2, they say that at small scales d agrees with, or at least should be understood as equivalent to, the distance function of the **unique right-invariant Riemannian metric on** G whose value at the identity is obtained by pulling back the Euclidean metric on $\mathbb{R}^m$ via ψ. They then refer to Lemma A.4. Annals of Mathematics

Your instinct is still right in an important sense: the estimates they actually prove and use are not developed in Riemannian language. Lemma A.4 compares $d_{\mathcal{X}}$ with coordinate differences such as

$$|\psi_{\mathcal{X}}(x) - \psi_{\mathcal{X}}(y)|$$

on bounded regions, with $Q^{O(1)}$-type constants when the bases are Q-rational. Annals of Mathematics

So the picture is:

$$d_{\mathcal{X}} \quad \text{is the Green–Tao working metric, built from } \ell^\infty \text{ Mal'cev coordinates.}$$

$$d_{\mathrm{Riem}} \quad \text{is the right-invariant Riemannian distance induced by declaring } d\psi_e(X_i) \text{ orthonormal.}$$

Locally, these are bi-Lipschitz equivalent. Indeed, near the identity,

$$d_{\mathcal{X}}(x,y) \asymp |\psi_{\mathcal{X}}(xy^{-1})|_\infty,$$

while the right-invariant Riemannian distance satisfies

$$d_{\mathrm{Riem}}(x,y) \asymp |\psi_{\mathcal{X}}(xy^{-1})|_2.$$

Since ℓ^∞ and ℓ^2 norms are equivalent in fixed dimension m,

$$|\cdot|_\infty \le |\cdot|_2 \le \sqrt{m}\,|\cdot|_\infty,$$

one gets

$$d_{\mathcal{X}}(x,y) \asymp_m d_{\mathrm{Riem}}(x,y)$$

locally.

One small caveat: if one interprets their word "agrees" literally as exact equality with the usual ℓ^2-Riemannian distance, that cannot be right in general, because even in the abelian case $G = \mathbb{R}^m$, their construction with ℓ^∞ gives the ℓ^∞ distance, not the Euclidean ℓ^2 distance. So the safe interpretation is **local equivalence/comparability**, which is all their arguments need.

Thus for your purposes, yes:

$$\boxed{d_{\mathcal{X}} \text{ is locally equivalent to the right-invariant Riemannian distance induced from } d\psi_e^*(\ell^2).}$$

But the detailed Green--Tao estimates are formulated for their ℓ^∞-based coordinate metric $d_{\mathcal{X}}$, not for the Riemannian metric directly.

---

40. User

2026-06-12 19:40:52 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:40:52.619000Z

if we pull back the l^infty metric from R^m to G at the identity, and then spread it around with right invariance, do we see get an analytic riemannian metric, or does it have to be the l^2 metric? My guess is that it has to be l^2 because l^infty is not induced by an inner product.

---

41. Assistant

2026-06-12 19:40:53 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-12T11:40:53.573922Z

Yes, your guess is right.

To get a **Riemannian metric**, the infinitesimal norm must come from an inner product. The ℓ^2-norm does:

$$\|v\|_2^2 = \langle v, v \rangle.$$

But the ℓ^∞-norm

$$\|v\|_\infty = \max_i |v_i|$$

does **not** come from an inner product. Its unit ball is a cube, not an ellipsoid, and it fails the parallelogram law. So pulling back ℓ^∞ at the identity and spreading by right translations does **not** give a Riemannian metric.

What it gives is a right-invariant **norm structure** on TG:

$$\|v\|_{x,\infty} = \left\| d\psi_\mathcal{X}|_e \big(dR_{x^{-1}} v \big) \right\|_\infty .$$

This lets you define lengths of curves by

$$L_\infty(\gamma) = \int_a^b \|\gamma'(t)\|_{\gamma(t),\infty}\, dt,$$

and hence a right-invariant path metric. But it is not Riemannian.

It is closer to a **right-invariant Finsler metric**, except that the ℓ^∞-norm is not smooth away from the origin along directions where two coordinates tie for the maximum. So it is not a smooth or analytic Finsler metric either. It is a perfectly good continuous right-invariant length structure, though.

By contrast, if you use ℓ^2, then at the identity you get an inner product

$$\langle Y, Z \rangle_e = \langle d\psi_\mathcal{X}|_e(Y), d\psi_\mathcal{X}|_e(Z) \rangle_{\ell^2} ,$$

and spreading by right translations gives a genuine right-invariant analytic Riemannian metric:

$$g_x(v, w) = \langle dR_{x^{-1}} v, dR_{x^{-1}} w \rangle_e .$$

So:

$$\boxed{\ell^2 \Rightarrow \text{right-invariant analytic Riemannian metric.}}$$

$$\boxed{\ell^\infty \Rightarrow \text{right-invariant length/Finsler-type metric, but not Riemannian.}}$$

For Green–Tao-style quantitative purposes, this distinction is usually harmless because in fixed dimension

$$\|v\|_\infty \le \|v\|_2 \le \sqrt{m}\, \|v\|_\infty.$$

Thus the ℓ^∞-based metric and the ℓ^2-based Riemannian metric are locally bi-Lipschitz equivalent, with constants depending only on m and the chosen coordinate control. But only the ℓ^2 construction gives a genuine analytic Riemannian metric.

---

42. User

2026-06-13 11:35:33 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:35:33.790000Z

is there a better/more modern reference than the alfred gray paper which will let me show that in a analytic riemannian manifold with all curvature related constants bounded by M, we have for small enough balls of radius r \in (0,f(M)) that the volume of the ball is between f_1(M)r^d and f_2(M)r^d, where d is the dimension of the manifold, and f,f_1, and f_2 are all polynomial functions.

---

43. Assistant

2026-06-13 11:36:42 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:36:42.122253Z

Yes — but I would **not** use Gray for the estimate you want. Gray is useful if you want the actual asymptotic expansion

$$\operatorname{vol} B(p, r) = \omega_d r^d \left(1 - \frac{\tau(p)}{6(d+2)} r^2 + \cdots \right),$$

but for the rough two-sided estimate

$$f_1(M) r^d \le \operatorname{vol} B(p, r) \le f_2(M) r^d,$$

standard **comparison geometry** is cleaner. Gray's Theorem 3.1 is much more precise than necessary. [TheVolumeOfASmallGeodesicBallOfARiemannianManifoldByAGray.pdf]

The reference I would use is something like **Lee, *Introduction to Riemannian Manifolds***, or **Gallot–Hulin–Lafontaine, *Riemannian Geometry***. Lee's book is a modern graduate text focused on curvature, geodesics, and comparison tools, while Gallot–Hulin–Lafontaine has the standard Riemannian metric/curvature/Ricci comparison material. Springer Springer For a dedicated comparison-geometry reference, Cheeger–Ebin's *Comparison Theorems in Riemannian Geometry* is the classic source; the AMS description says its theme is the interaction between curvature and global geometry. American Mathematical Society

The theorem you want should be stated with one extra hypothesis: **you need an injectivity-radius scale**, not just curvature bounds. Curvature bounds alone do not prevent collapse. For example, a flat torus with one circle factor of length ε has curvature 0, but balls of radius $r \gg \varepsilon$ have volume comparable to εr^{d-1}, not r^d. So no lower bound of the form cr^d can follow from curvature alone.

A clean usable statement is:

Let (N^d, g) be a smooth Riemannian manifold. Suppose

$$|\sec_g| \le K$$

and

$$\operatorname{inj}(N, g) \ge \iota.$$

Then there are constants $c_d, C_d > 0$ such that for every $p \in N$ and every

$$0 < r \le c_d \min\{\iota, K^{-1/2}\},$$

one has

$$c_d r^d \le \operatorname{vol}_g B(p, r) \le C_d r^d.$$

More precisely, for $r < \operatorname{inj}(p)$, comparison of Jacobi fields gives estimates of the form

$$\omega_{d-1} \int_0^r \left(\frac{\sin(\sqrt{K}t)}{\sqrt{K}} \right)^{d-1} dt \le \operatorname{vol} B(p, r) \le \omega_{d-1} \int_0^r \left(\frac{\sinh(\sqrt{K}t)}{\sqrt{K}} \right)^{d-1} dt,$$

with the usual interpretation when $K = 0$. For $r \le c/\sqrt{K}$, both integrals are comparable to r^d. This is the standard Bishop–Günther/Bishop–Gromov comparison mechanism; Wei's survey states Bishop–Gromov volume comparison explicitly, and Eschenburg's comparison notes describe volume and length estimates arising from curvature comparison. Department of Mathematics Math Dept - University of Toronto

For your nilmanifold application, the right formulation is therefore:

$$|R| \le M^A \quad \text{and} \quad \operatorname{inj}(G'/\Gamma') \ge M^{-B}$$

imply

$$c_d r^d \le \operatorname{vol} B(x, r) \le C_d r^d$$

for

$$0 < r \le c_d M^{-C},$$

where C depends on A, B, d. The constants c_d, C_d can be chosen independent of the particular subnilmanifold. In your desired language, $1/f(M)$ is polynomial in M, while f_1, f_2 may even be taken to be positive dimension-dependent constants, hence degree-zero polynomials.

So the main remaining thing to prove in your setting is not the volume estimate itself; it is the **polynomial lower bound on the injectivity radius** of the rational subnilmanifolds. The curvature bound follows from the rational Mal'cev-basis control as we discussed, and the small-ball volume estimate then follows from standard comparison geometry.

---

44. User

2026-06-13 11:45:03 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:45:03.576000Z

can you explain to me what is the injectivity radius?

45. Assistant

2026-06-13 11:45:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:45:04.904344Z

The **injectivity radius** is the largest radius up to which geodesic coordinates around a point behave without ambiguity.

Let (M, g) be a Riemannian manifold and $p \in M$. The exponential map at p is

$$\exp_p : T_pM \to M.$$

Given a tangent vector $v \in T_pM$, $\exp_p(v)$ means: start at p, travel along the geodesic with initial velocity v, and stop at time 1.

The **injectivity radius at** p is

$$\mathrm{inj}_p(M) = \sup\left\{r > 0 : \exp_p \text{ is a diffeomorphism from } B_{T_pM}(0, r) \text{ onto its image}\right\}.$$

Informally, it is the largest radius such that every point within distance $< r$ from p is reached from p by a **unique shortest geodesic**, and normal coordinates centered at p are valid.

The **global injectivity radius** is

$$\mathrm{inj}(M) = \inf_{p \in M} \mathrm{inj}_p(M).$$

So if

$$r < \mathrm{inj}(M),$$

then every radius-r ball in M is "honestly geodesic": it does not wrap around the manifold, and it is parametrized by a genuine ball in the tangent space via $\exp_p$.

For example:

Euclidean space

In $\mathbb{R}^d$,

$$\mathrm{inj}(\mathbb{R}^d) = \infty.$$

Straight lines never stop being unique shortest paths.

Circle of circumference L

For the circle S^1 of circumference L,

$$\mathrm{inj}(S^1) = \frac{L}{2}.$$

If you go less than halfway around the circle, there is a unique shortest arc. At exactly halfway, there are two shortest arcs.

Flat torus

For a flat torus

$$\mathbb{R}^d/\Lambda,$$

the injectivity radius is

$$\mathrm{inj}(\mathbb{R}^d/\Lambda) = \frac{1}{2} \min_{\lambda \in \Lambda \setminus \{0\}} \|\lambda\|.$$

So it is half the length of the shortest nontrivial closed geodesic. If the torus has a very short circle direction, the injectivity radius is very small even though the curvature is zero.

This is why injectivity radius matters for volume estimates. Curvature bounds alone control the local shape of geodesic balls, but they do not prevent the manifold from wrapping around itself at very small scales.

In your nilmanifold setting, the analogy is similar. On

$$G/\Gamma,$$

small balls in G/Γ look like balls in G as long as their radius is below the scale at which two different Γ-translates become close. Roughly, the injectivity radius measures the shortest nontrivial way of returning near the same point because of the lattice Γ.

For small-ball volume estimates, one usually wants

$$0 < r < c\,\mathrm{inj}(M)$$

and also r small compared to the curvature scale. Then the geodesic ball $B(p,r)$ is well-controlled and has volume comparable to r^d.

---

46. User

2026-06-13 11:45:51 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:45:51.822000Z

can we prove that the injectivity radius of the subnilmanifolds is bounded by a polynomial function of the parameter M from the earlier discussion?

---

47. Assistant

2026-06-13 11:45:52 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T03:45:52.960419Z

Yes — in the nilmanifold setting you should be able to prove an **inverse-polynomial lower bound** for the injectivity radius. More precisely, the useful statement is not

$$\mathrm{inj}(G'/\Gamma') \le \mathrm{poly}(M),$$

but rather

$$\boxed{\mathrm{inj}(G'/\Gamma') \ge c\,M^{-A}}$$

or equivalently

$$\boxed{\mathrm{inj}(G'/\Gamma')^{-1} \le CM^{A}.}$$

Here A, C, c may depend on the fixed ambient nilmanifold G/Γ, the original Mal'cev basis $\mathcal{X}$, the ambient dimension m, and the step s, but **not** on the particular subnilmanifold except through M.

Let me describe the proof.

Let

$$H/\Lambda$$

be one of your rational subnilmanifolds, where $H \le G$ is a connected rational subgroup and $\Lambda = H \cap \Gamma$. Suppose

$$\mathcal{Y} = (Y_1, \dots, Y_d)$$

is a Mal'cev basis for H/Λ, and each Y_i is an M-rational combination of the original ambient basis $\mathcal{X} = (X_1, \dots, X_m)$.

Declare $Y_1, \dots, Y_d$ orthonormal at the identity and spread this inner product by right translations. This gives a right-invariant Riemannian metric on H, descending to H/Λ.

Write the Lie bracket in the $\mathcal{Y}$-basis as

$$[Y_i, Y_j] = \sum_{k=1}^{d} d_{ij}^k Y_k.$$

As discussed earlier, the M-rationality assumption implies

$$|d_{ij}^k| \le C_{\mathcal{X},m,s} M^A.$$

Let

$$B := 1 + \max_{i,j,k} |d_{ij}^k|.$$

Then $B \le CM^A$.

The key geometric lemma is:

$$\boxed{\mathrm{inj}(H/\Lambda) \geq c_{d,s} B^{-A_{d,s}}.}$$

Once we prove this, the desired bound follows immediately:

$$\mathrm{inj}(H/\Lambda) \geq cM^{-A'}.$$

The proof has two ingredients.

First, the simply connected cover H has a polynomially controlled local injectivity radius. In Mal'cev coordinates associated to $\mathcal{Y}$, the group law is given by BCH polynomials. Since the nilpotence step is at most s, these polynomials have bounded degree, and their coefficients are polynomial functions of the structure constants d_{ij}^k. Therefore the right-invariant metric coefficients, Christoffel symbols, and the local geodesic equation have coefficients bounded by some polynomial in B.

By a quantitative inverse-function-theorem argument applied to the Riemannian exponential map,

$$\exp_e^{\mathrm{Riem}} : T_e H \to H,$$

one gets

$$\mathrm{inj}_H(e) \geq cB^{-A}.$$

Because the metric is right-invariant, the same bound holds at every point of H.

Second, the lattice Λ has no nontrivial element too close to the identity. Since $\mathcal{Y}$ is a Mal'cev basis for H/Λ, every $\lambda \in \Lambda$ has integer Mal'cev coordinates:

$$\psi_{\mathcal{Y}}(\lambda) \in \mathbb{Z}^d.$$

Thus if $\lambda \neq e$, then

$$\|\psi_{\mathcal{Y}}(\lambda)\|_\infty \geq 1.$$

On the other hand, for sufficiently small radii, say

$$r \leq cB^{-A},$$

the Riemannian ball $B_H(e, r)$ is contained in a small Mal'cev coordinate box, for instance

$$\|\psi_{\mathcal{Y}}(h)\|_\infty < \frac{1}{2}.$$

This follows from the same polynomial control of the right-invariant frame in Mal'cev coordinates. Therefore no nontrivial $\lambda \in \Lambda$ can lie inside $B_H(e, cB^{-A})$. Hence

$$\inf_{\lambda \in \Lambda \setminus \{e\}} d_H(e, \lambda) \geq cB^{-A}.$$

Finally, since

$$H \to H/\Lambda$$

is a Riemannian covering and right translations by elements of Λ are deck transformations, the injectivity radius of the quotient satisfies

$$\mathrm{inj}(H/\Lambda) \geq \min\left\{ \mathrm{inj}(H), \frac{1}{2} \inf_{\lambda \neq e} d_H(e, \lambda) \right\}.$$

Both terms on the right are bounded below by cB^{-A}, so

$$\mathrm{inj}(H/\Lambda) \geq cB^{-A}.$$

Since $B \leq CM^A$, this gives

$$\boxed{\mathrm{inj}(H/\Lambda) \geq c_{\mathcal{X},m,s} M^{-A_{\mathcal{X},m,s}}.}$$

So yes: in your setting, the injectivity radius is controlled **inverse-polynomially** in the rationality parameter M.

One caveat: this uses that $\mathcal{Y}$ is genuinely a Mal'cev basis adapted to Λ, so that Λ has integer Mal'cev coordinates in the $\mathcal{Y}$-coordinates. If you only knew that H had an M-rational spanning basis but not that it was adapted to the lattice, then this argument would be missing the crucial lattice-separation step.

48. User

2026-06-13 14:38:19 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T06:38:19.196000Z

Why does the cover H even need an injectivity radius? Is the exponential map not globally injective?

---

49. Assistant

2026-06-13 14:38:20 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T06:38:20.856991Z

There is an ambiguity in the word **exponential** here.

For a simply connected nilpotent Lie group H, the **Lie-group exponential**

$$\exp_{\mathrm{Lie}} : \mathfrak{h} \to H$$

is indeed a global analytic diffeomorphism. So in that sense there is no obstruction.

But the **injectivity radius** uses the **Riemannian exponential**

$$\mathrm{Exp}_p^g : T_pH \to H,$$

which sends a tangent vector v to the endpoint at time 1 of the Riemannian geodesic starting at p with initial velocity v. This is a different map.

They agree in flat Euclidean space, and more generally for bi-invariant metrics, but your right-invariant metric on a nonabelian nilpotent group is usually **not** bi-invariant. In fact, a nonabelian nilpotent Lie group cannot carry a positive-definite bi-invariant Riemannian metric. So geodesics need not be one-parameter subgroups, and

$$\mathrm{Exp}_e^g(v) \neq \exp_{\mathrm{Lie}}(v)$$

in general.

So the answer is:

$$\boxed{\exp_{\mathrm{Lie}} : \mathfrak{h} \to H \text{ is globally injective, but this does not imply that } \mathrm{Exp}_e^g : T_eH \to H \text{ is globally injective.}}$$

That is why I mentioned the injectivity radius of the cover H. The quotient injectivity radius is controlled by two possible obstructions:

1. **local Riemannian cut/conjugacy phenomena already present in** H, measured by inj_H;
2. **wrapping around by the lattice**, measured by the shortest nontrivial displacement

$$\inf_{\lambda \in \Lambda \setminus \{e\}} d_H(e, \lambda).$$

For a Riemannian covering $H \to H/\Lambda$, one expects a lower bound of the form

$$\mathrm{inj}(H/\Lambda) \geq \min\left\{ \mathrm{inj}(H), \frac{1}{2} \inf_{\lambda \neq e} d_H(e, \lambda) \right\}.$$

The lattice-separation term is the “nilmanifold wrapping” obstruction. The $\mathrm{inj}(H)$ term is the possible failure of Riemannian normal coordinates on the cover.

That said, for your nilmanifold volume argument, you may not actually need to discuss $\mathrm{inj}(H)$ explicitly. A more nilpotent-coordinate proof can use the global Lie exponential / Mal’cev coordinates directly:

$$\psi_{\mathcal{Y}} : H \to \mathbb{R}^d.$$

Because $\mathcal{Y}$ is a Mal’cev basis for H/Λ, nontrivial lattice elements have nonzero integer Mal’cev coordinates. Hence no nontrivial $\lambda \in \Lambda$ lies in the coordinate box

$$\|\psi_{\mathcal{Y}}(h)\|_\infty < 1/2.$$

Then, using polynomial control of the right-invariant Riemannian metric in those coordinates, one proves that for

$$r \leq M^{-A},$$

the Riemannian ball $B_H(e, r)$ lies inside that coordinate box. Therefore the projection

$$H \to H/\Lambda$$

is injective on $B_H(e,r)$, and the quotient ball has the same volume as the lifted ball.

So the corrected picture is:

$$\boxed{\text{The Lie exponential is globally injective and helps with coordinates/lattice separation.}}$$

$$\boxed{\text{The Riemannian exponential controls geodesic normal coordinates and injectivity radius.}}$$

$$\boxed{\text{For your application, one can often avoid the cover injectivity-radius issue by working directly in Mal'cev coordinates.}}$$

My earlier reference to $\mathrm{inj}(H)$ was appropriate for a standard comparison-geometry proof, but in the nilpotent setting a direct Mal'cev-coordinate proof is probably cleaner.

---

50. User

2026-06-13 14:55:17 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T06:55:17.134000Z

I understand now that there is a difference between the lie bracket exponential and the exponential for the riemannian metric, but I don't understand your proposition for how to avoid the exponential of the riemannian metric and calculate the volume of small balls just with the lie bracket exponential. Can you explain this part to me again in more detail? Also, don't box your answers, because they go off screen and I can't read the end.

---

51. Assistant

2026-06-13 14:55:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T06:55:21.288903Z

Yes, let me separate the two ideas more cleanly.

I should not have said "calculate the volume" in the sense of getting the exact shape of the Riemannian ball. The point is more modest: **we can estimate the volume of small Riemannian balls using Mal'cev/Lie coordinates, without using the Riemannian exponential map.**

The reason is that for volume estimates of order r^d, we do not need geodesic normal coordinates. We only need a coordinate chart in which the Riemannian metric tensor is uniformly comparable to the Euclidean metric.

Let H be a simply connected nilpotent Lie group with lattice Λ, and let

$$\mathcal{Y} = (Y_1, \dots, Y_d)$$

be a Mal'cev basis for H/Λ. Use Mal'cev coordinates

$$\psi_{\mathcal{Y}} : H \to \mathbb{R}^d,$$

for instance

$$h = \exp(t_1 Y_1) \cdots \exp(t_d Y_d), \qquad \psi_{\mathcal{Y}}(h) = (t_1, \dots, t_d).$$

This is a global analytic coordinate chart. It comes from the Lie group exponential/product coordinates, not from the Riemannian exponential.

Now put the right-invariant Riemannian metric on H by declaring $Y_1, \dots, Y_d$ orthonormal at the identity. Let E_i be the corresponding right-invariant vector field whose value at the identity is Y_i. Then $E_1, \dots, E_d$ is a global orthonormal frame.

In Mal'cev coordinates $t = (t_1, \dots, t_d)$, each E_i can be written as

$$E_i(t) = \sum_{a=1}^{d} A_i^a(t) \frac{\partial}{\partial t_a}.$$

The matrix $A(t) = (A_i^a(t))$ is analytic, and in the nilpotent case it is polynomial in t. Its coefficients are controlled by the structure constants of the Lie algebra in the basis $\mathcal{Y}$. Also

$$A(0) = I.$$

Therefore, for $\|t\|_\infty \le \eta$, with η sufficiently small depending polynomially on the structure constants, $A(t)$ remains close to I. In particular, the Riemannian metric tensor in the t-coordinates is uniformly comparable to the Euclidean metric:

$$C^{-1}\|\xi\|_2^2 \le g_t(\xi,\xi) \le C\|\xi\|_2^2$$

for all coordinate tangent vectors ξ, as long as $\|t\|_\infty \le \eta$.

This is the main replacement for Riemannian normal coordinates.

Once we have this, we can compare Riemannian distance with coordinate distance. If $h \in H$ has Mal'cev coordinate $t = \psi_{\mathcal{Y}}(h)$ and $\|t\|_\infty \le \eta$, then for sufficiently small t,

$$c\|t\|_2 \le d_H(e,h) \le C\|t\|_2.$$

The upper bound comes from taking the straight-line path in coordinates from 0 to t. The lower bound comes from the fact that every path from 0 to t has Euclidean coordinate length at least $\|t\|_2$, and Riemannian length dominates Euclidean length up to the factor c.

Thus the Riemannian ball is sandwiched between coordinate balls:

$$\{h : \|\psi_{\mathcal{Y}}(h)\|_2 \le cr\} \subseteq B_H(e,r) \subseteq \{h : \|\psi_{\mathcal{Y}}(h)\|_2 \le Cr\},$$

provided r is small enough that these coordinate balls remain inside $\|t\|_\infty \le \eta$.

Now look at volume. In Mal'cev coordinates, the Riemannian volume form has the form

$$d\,\mathrm{vol}_g = J(t)\,dt_1 \cdots dt_d,$$

where

$$J(t) = \sqrt{\det(g_{ab}(t))}.$$

Since $g_{ab}(t)$ is uniformly comparable to the Euclidean metric in our small coordinate box, $J(t)$ is bounded above and below:

$$c' \le J(t) \le C'$$

for $\|t\|_\infty \le \eta$.

Therefore,

$$\mathrm{vol}_H(B_H(e,r)) \asymp r^d$$

for all sufficiently small r. More explicitly,

$$c_1 r^d \le \mathrm{vol}_H(B_H(e,r)) \le c_2 r^d.$$

This proof never uses the Riemannian exponential map. It uses only:

1. Mal'cev coordinates coming from the Lie group structure;
2. right-invariance of the Riemannian metric;
3. local comparability of the metric tensor to the Euclidean tensor;
4. the volume density $J(t) = \sqrt{\det g(t)}$.

Now we still need to descend to the quotient H/Λ. Since $\mathcal{Y}$ is a Mal'cev basis adapted to Λ, lattice elements have integer Mal'cev coordinates:

$$\psi_{\mathcal{Y}}(\Lambda) = \mathbb{Z}^d$$

up to the usual convention for Mal'cev coordinates.

So nontrivial lattice elements cannot lie in a sufficiently small coordinate box around the identity. More quantitatively, because the group law and inverse map are polynomial in Mal'cev coordinates with polynomially controlled coefficients, there is some

$$\eta \ge M^{-A}$$

such that if

$$\|\psi_{\mathcal{Y}}(h_1)\|_\infty,\ \|\psi_{\mathcal{Y}}(h_2)\|_\infty \le \eta,$$

then

$$\|\psi_{\mathcal{Y}}(h_2^{-1}h_1)\|_\infty < \frac{1}{2}.$$

If $h_1\Lambda = h_2\Lambda$, then $h_2^{-1}h_1 \in \Lambda$. But a nontrivial lattice element has some integer coordinate nonzero, so its ℓ^∞-coordinate norm is at least 1. Hence $h_2^{-1}h_1 = e$, and therefore $h_1 = h_2$. Thus the quotient map

$$H \to H/\Lambda$$

is injective on this small coordinate box.

Combining this with the ball containment

$$B_H(e,r) \subseteq \{\|\psi_{\mathcal{Y}}(h)\|_\infty \le \eta\},$$

valid for $r \le cM^{-A}$, the projection is injective on $B_H(e,r)$. Therefore the quotient ball of radius r has the same volume as the lifted ball in H, at least for such small r.

So the final estimate is:

$$c_1 r^d \le \mathrm{vol}_{H/\Lambda}(B(x,r)) \le c_2 r^d$$

for every $x \in H/\Lambda$ and every

$$0 < r \le cM^{-A}.$$

The constants c_1, c_2, c, A depend on the ambient dimension, step, and the fixed ambient nilmanifold data, but not on the particular rational subnilmanifold except through M.

The role of the Lie/Mal'cev coordinates is not to identify the Riemannian ball exactly. Rather, they give a controlled coordinate system in which the metric tensor and the volume density are bounded. That is enough to sandwich the Riemannian ball between ordinary coordinate balls and get the r^d volume estimate.

---

52. User

2026-06-13 15:11:04 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:11:04.483000Z

what is this g_ab(t) of which you are taking a determinant?

---

53. Assistant

2026-06-13 15:11:06 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:11:06.544446Z

Here $g_{ab}(t)$ means the **matrix of the Riemannian metric in the Mal'cev coordinate basis**.

Let

$$\psi_{\mathcal{Y}} : H \to \mathbb{R}^d$$

be Mal'cev coordinates, and write

$$h(t) = \psi_{\mathcal{Y}}^{-1}(t), \qquad t = (t_1, \dots, t_d).$$

The coordinate vector fields are

$$\frac{\partial}{\partial t_1}, \dots, \frac{\partial}{\partial t_d}.$$

At the point $h(t)$, the Riemannian metric is an inner product on $T_{h(t)}H$. So we define

$$g_{ab}(t) = g_{h(t)}\left(\frac{\partial}{\partial t_a}, \frac{\partial}{\partial t_b}\right).$$

Thus

$$G(t) = (g_{ab}(t))_{a,b=1}^d$$

is the matrix representing the Riemannian metric in the Mal'cev coordinate chart.

If a tangent vector is written in coordinates as

$$v = \sum_a \xi^a \frac{\partial}{\partial t_a},$$

then its squared Riemannian norm is

$$g_{h(t)}(v,v) = \sum_{a,b} g_{ab}(t)\xi^a\xi^b.$$

That is the object whose determinant appears in the volume form:

$$d\,\mathrm{vol}_g = \sqrt{\det(g_{ab}(t))}\, dt_1 \cdots dt_d.$$

So the determinant is simply the usual Jacobian factor converting coordinate Lebesgue measure $dt_1 \cdots dt_d$ into Riemannian volume.

Now connect this to the right-invariant frame. Let $E_1, \ldots, E_d$ be the right-invariant vector fields obtained from the basis $Y_1, \ldots, Y_d$. By construction,

$$E_1, \ldots, E_d$$

is an orthonormal frame for the right-invariant Riemannian metric. But the coordinate fields

$$\partial_{t_1}, \ldots, \partial_{t_d}$$

are generally not orthonormal.

Suppose

$$E_i(t) = \sum_a A_i^a(t) \frac{\partial}{\partial t_a}.$$

Equivalently, $A(t)$ is the change-of-basis matrix from the orthonormal right-invariant frame to the coordinate frame. Since the E_i's are orthonormal, the coordinate metric matrix $G(t) = (g_{ab}(t))$ is determined by

$$A(t)^T G(t) A(t) = I.$$

So

$$G(t) = A(t)^{-T} A(t)^{-1}.$$

In particular,

$$\sqrt{\det G(t)} = \frac{1}{|\det A(t)|}.$$

This is the density of Riemannian volume in Mal'cev coordinates.

At the identity, $t = 0$, if the Mal'cev coordinate basis is normalized so that

$$\left.\frac{\partial}{\partial t_i}\right|_0 = Y_i,$$

then

$$G(0) = I.$$

Away from the identity, $G(t)$ is usually not the identity, because the coordinate vector fields and the right-invariant vector fields differ. But for t sufficiently small, $G(t)$ remains close to I. That is the local comparison we use.

In short:

$$g_{ab}(t)$$

is not a new curvature quantity. It is just the local coordinate expression of the Riemannian metric:

$$g_{ab}(t) = \langle \partial_{t_a}, \partial_{t_b} \rangle_{h(t)}.$$

The reason we take $\det(g_{ab}(t))$ is that Riemannian volume in coordinates is

$$\sqrt{\det(g_{ab}(t))}\, dt_1 \cdots dt_d.$$

For the volume estimate, we want to know that on a small Mal'cev coordinate box,

$$c \le \sqrt{\det(g_{ab}(t))} \le C.$$

Then Riemannian volume is comparable to ordinary Lebesgue measure in those coordinates.

---

54. User

2026-06-13 15:22:16 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:22:16.499000Z

but I think that it is not enough to have an upper and lower bound with c and C, because these constants depend on the manifold. when I pass to the subnilmanifold, I need the constants to be global and then maybe also have a factor of M^a, where M is the rationality parameter from before.

---

55. Assistant

2026-06-13 15:22:18 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:22:18.625166Z

You are exactly right: a statement with constants c, C is useless for your application unless we know how c, C depend on the subnilmanifold. The point is that, in the Mal'cev-coordinate proof, we can arrange the constants to be **uniform**, after shrinking the allowed radius by an inverse-polynomial function of M. If you use a slightly different metric, then the upper/lower constants may also acquire polynomial factors in M.

Let me spell this out.

Suppose H/Λ is one of your rational subnilmanifolds, with Mal'cev basis

$$\mathcal{Y} = (Y_1, \dots, Y_d), \qquad d \le m,$$

and suppose each Y_i is M-rational in the ambient Mal'cev basis. Put the right-invariant metric on H by declaring the Y_i's orthonormal at the identity.

Write the bracket in the $\mathcal{Y}$-basis as

$$[Y_i, Y_j] = \sum_k c_{ij}^k Y_k.$$

From the M-rationality assumption, we have

$$\max_{i,j,k} |c_{ij}^k| \le C_{\mathcal{X},m,s} M^{A_{\mathcal{X},m,s}}.$$

Call this bound B. So

$$B \le C M^A.$$

Now use Mal'cev coordinates

$$h(t) = \exp(t_1 Y_1) \cdots \exp(t_d Y_d).$$

Let $E_1, \dots, E_d$ be the right-invariant orthonormal frame. In coordinates,

$$E_i(t) = \sum_a A_i^a(t) \frac{\partial}{\partial t_a}.$$

The matrix $A(t)$ satisfies

$$A(0) = I.$$

Moreover, because the group law is given by BCH polynomials of degree at most s, the entries of $A(t)$ are polynomials in t, whose coefficients are polynomially bounded in the structure constants. Therefore there is some exponent $L = L(m, s)$ such that, for $\|t\|_\infty \le 1$,

$$\|A(t) - I\| \le C_{m,s} B^L \|t\|_\infty.$$

So if we restrict to

$$\|t\|_\infty \le \eta, \qquad \eta := c_{m,s} B^{-L},$$

then

$$\|A(t) - I\| \le \frac{1}{2}.$$

At that scale, all singular values of $A(t)$ are bounded between, say, $1/2$ and 2, uniformly in the subnilmanifold. Hence the metric matrix

$$G(t) = (g_{ab}(t))$$

satisfies

$$C_d^{-1}|\xi|^2 \le \sum_{a,b} g_{ab}(t)\xi_a\xi_b \le C_d|\xi|^2$$

for $\|t\|_\infty \le \eta$, with C_d depending only on d, hence only on m.

Also,

$$d\,\mathrm{vol}_g = \sqrt{\det G(t)}\,dt.$$

Since the eigenvalues of $G(t)$ are uniformly bounded above and below, we get

$$C_d^{-1} \le \sqrt{\det G(t)} \le C_d$$

on the same box. This is now a **uniform** estimate. The dependence on the particular subnilmanifold has been absorbed into the size of the coordinate box:

$$\eta \ge cM^{-A'}.$$

So the constants in the volume comparison do not secretly depend on H/Λ; only the allowed radius does.

Next, because the metric tensor is uniformly comparable to the Euclidean metric in this coordinate box, the Riemannian ball is sandwiched between coordinate balls:

$$\{\|t\|_2 \le c_d r\} \subseteq B_H(e,r) \subseteq \{\|t\|_2 \le C_d r\},$$

provided

$$r \le c_d\eta.$$

Therefore

$$c_d r^d \le \mathrm{vol}_H(B_H(e,r)) \le C_d r^d$$

for all

$$0 < r \le cM^{-A'}.$$

These constants c_d, C_d are genuinely uniform.

Now we need to make sure the ball does not wrap around the quotient. Since $\mathcal{Y}$ is a Mal'cev basis adapted to Λ, nontrivial elements of Λ have integer Mal'cev coordinates, hence cannot lie in

$$\|\psi_{\mathcal{Y}}(h)\|_\infty < 1/2.$$

The BCH formula gives a polynomially controlled estimate for $h_2^{-1}h_1$ in terms of the coordinates of h_1, h_2. Thus, after possibly shrinking η again by another inverse-polynomial factor in B, the quotient map

$$H \to H/\Lambda$$

is injective on the coordinate box $\|t\|_\infty \le \eta$. Consequently, for

$$r \le cM^{-A''},$$

the quotient ball is represented injectively upstairs, and

$$\mathrm{vol}_{H/\Lambda}(B_{H/\Lambda}(x,r)) = \mathrm{vol}_H(B_H(\tilde{x},r)).$$

Using right-invariance upstairs, the same volume estimate holds at every lift $\tilde{x}$. Therefore

$$c_m r^d \le \mathrm{vol}_{H/\Lambda}(B(x,r)) \le C_m r^d$$

for all $x \in H/\Lambda$ and all

$$0 < r \le cM^{-A}.$$

So in the intrinsic metric where the subnilmanifold's Mal'cev basis $\mathcal{Y}$ is declared orthonormal, you can take

$$f(M) = cM^{-A},$$

and

$$f_1(M) = c_m, \qquad f_2(M) = C_m.$$

The upper and lower constants can be degree-zero polynomials in M. The polynomial dependence enters mainly in the allowable radius.

There is one variant where f_1, f_2 may need actual M-powers. If instead of declaring $\mathcal{Y}$ orthonormal, you restrict the ambient right-invariant metric from G to H, then the Gram matrix

$$\big(\langle Y_i, Y_j\rangle_{\mathcal{X}}\big)_{i,j}$$

may have eigenvalues depending on M. But because the Y_i's have rational coordinates of height $\le M$, Cramer-type estimates give only polynomial degeneration:

$$M^{-A} \le \lambda_{\min} \le \lambda_{\max} \le M^{A}.$$

In that case the volume density bounds can acquire polynomial factors:

$$M^{-A} r^d \le \mathrm{vol}(B(x,r)) \le M^{A} r^d,$$

for

$$0 < r \le cM^{-A}.$$

But for the metric normalized by the sub-Mal'cev basis $\mathcal{Y}$, the cleanest result is stronger:

$$\mathrm{vol}(B(x,r)) \asymp_m r^d$$

uniformly, as long as $r \le M^{-O(1)}$.

---

56. User

2026-06-13 15:33:55 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:33:55.644000Z

instead of proving all of this from scratch in my paper, is there anything I can cite about volumes of small balls in nilmanifolds and subnilmanifolds and their relation to the structure constants and rationality of the mal cev basis?

---

57. Assistant

2026-06-13 15:35:11 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:35:11.119845Z

I do not know a single reference that states exactly your desired lemma with the words "volume of small balls in rational subnilmanifolds is polynomially controlled by the rationality of the Mal'cev basis." The closest standard source is **Green–Tao, "The quantitative behaviour of polynomial orbits on nilmanifolds," Appendix A**, plus a basic reference on Haar measure in Mal'cev coordinates.

The best citation strategy is probably:

1. cite **Green–Tao Appendix A** for all the quantitative Mal'cev-coordinate and rationality input;
2. cite a nilpotent Lie group reference for the fact that Mal'cev coordinates preserve Haar measure;
3. include a short lemma in your paper saying that these facts imply the small-ball volume estimate.

Green–Tao explicitly define their Mal'cev-coordinate metric $d_{\mathcal{X}}$, using the ℓ^∞-norm in coordinates, and the descended quotient metric on G/Γ. They also define Q-rational Mal'cev bases by bounding the heights of the structure constants, and Q-rational subgroups by requiring a basis made of rational combinations of the ambient Mal'cev basis. Annals of Mathematics

The key technical pieces in Green–Tao are:

- **Lemma A.3**: multiplication and inversion in Mal'cev coordinates are polynomial maps with coefficients of height $Q^{O(1)}$ when the basis is Q-rational. Annals of Mathematics

- **Lemma A.4**: their metric $d_{\mathcal{X}}$ is quantitatively comparable to differences in Mal'cev coordinates, with $Q^{O(1)}$-losses. Annals of Mathematics
- **Proposition A.10**: if $G_0 \leq G$ is a Q-rational subgroup, then G_0/Γ_0 has a Mal'cev basis whose vectors are $Q^{O(1)}$-rational combinations of the ambient basis. Annals of Mathematics
- **Lemma A.17**: the Green–Tao metrics on G/Γ and G_0/Γ_0 are comparable with $Q^{O(1)}$-losses. Annals of Mathematics

For the measure side, a clean modern reference is Le Donne's chapter on nilpotent Lie groups. It states that in simply connected nilpotent groups, exponential and Mal'cev coordinates are global, the group product is polynomial, left and right translations have Jacobian one, and therefore Lebesgue measure in these coordinates is a bi-invariant Haar measure. Springer

Putting these together, you can state something like this.

Let G/Γ be a fixed m-dimensional nilmanifold with Q_0-rational Mal'cev basis $\mathcal{X}$. Let $H \leq G$ be a connected M-rational subgroup, $\Lambda = H \cap \Gamma$, and let $d = \dim H$. Then, after choosing the Green–Tao Mal'cev basis $\mathcal{Y}$ of H/Λ, there are constants $A, C > 0$, depending only on the ambient nilmanifold data, m, and the step, such that for all $x \in H/\Lambda$ and all

$$0 < r \leq C^{-1}M^{-A},$$

one has

$$C^{-1}M^{-A}r^d \leq \mu_{H/\Lambda}\big(B(x,r)\big) \leq CM^{A}r^d.$$

Here $B(x,r)$ can be taken with respect to either the Green–Tao metric $d_{\mathcal{Y}}$, or with respect to the right-invariant Riemannian metric obtained by declaring $\mathcal{Y}$ orthonormal, after changing A, C. If you use the $\mathcal{Y}$-orthonormal right-invariant Riemannian metric, the constants in front of r^d can often be made dimension-only; the main M-dependence is then in the allowed radius $r \leq M^{-O(1)}$. But allowing $M^{\pm A}$ in the volume bounds is safer and easier to cite/prove.

So the honest answer is: **Green–Tao gives almost exactly the quantitative Mal'cev technology you need, but not packaged as a small-ball volume lemma.** You will probably still need to include a short lemma, but you can make the proof very short by saying it follows from Green–Tao Lemmas A.3, A.4, A.17, Proposition A.10, and the fact that Haar measure is Lebesgue measure in Mal'cev coordinates.

A good way to present it is:

> By Green–Tao's quantitative Mal'cev-coordinate estimates, all coordinate changes, multiplication maps, inverse maps, and metric comparisons on M-rational subnilmanifolds have $M^{O(1)}$-controlled constants. Since Haar measure is Lebesgue measure in Mal'cev coordinates, coordinate boxes of side r have measure $\asymp r^d$. Therefore sufficiently small metric balls, with radius at most $M^{-O(1)}$, are sandwiched between coordinate boxes of side $M^{-O(1)}r$ and $M^{O(1)}r$, giving the stated polynomial volume bounds.

That should avoid reproving the whole coordinate technology from scratch while still making the small final inference explicit.

---

58. User

2026-06-13 15:46:25 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:46:25.480000Z

can you elaborate on the relationship between lebesgue measure in mal cev coordinates and the bi invariant haar measure? This is not clear to me?

---

59. Assistant

2026-06-13 15:46:27 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T07:46:27.947990Z

Yes. The key point is that in a simply connected nilpotent Lie group, the usual Mal'cev coordinates are not just arbitrary coordinates: translations act by polynomial maps whose Jacobian determinant is 1. Therefore ordinary Lebesgue measure in those coordinates is Haar measure.

Let H be a connected simply connected nilpotent Lie group with Lie algebra $\mathfrak{h}$, and let

$$\mathcal{Y} = (Y_1, \dots, Y_d)$$

be a Mal'cev basis. Use Mal'cev coordinates

$$\Phi : \mathbb{R}^d \to H, \qquad \Phi(t_1, \dots, t_d) = \exp(t_1 Y_1) \cdots \exp(t_d Y_d).$$

This is a global analytic diffeomorphism. So we can transfer Lebesgue measure from $\mathbb{R}^d$ to H:

$$\mu_H(A) = \mathrm{Leb}\big(\Phi^{-1}(A)\big).$$

The claim is that μ_H is Haar measure on H.

Why? Fix $a \in H$. Left translation is

$$L_a(h) = ah.$$

In Mal'cev coordinates this becomes a map

$$T_a : \mathbb{R}^d \to \mathbb{R}^d, \qquad T_a(t) = \psi_y(a\Phi(t)).$$

Because the group is nilpotent, the Baker–Campbell–Hausdorff formula terminates, so T_a is a polynomial map. More importantly, because the basis is ordered in the Mal'cev way, this polynomial map is triangular with diagonal entries equal to 1. Schematically,

$$T_a(t)_i = t_i + \text{terms depending on } a \text{ and other coordinates}.$$

The precise dependence on earlier/later coordinates depends on the convention for ordering the Mal'cev basis, but the important fact is:

$$\det DT_a(t) = 1.$$

Therefore T_a preserves Lebesgue measure on $\mathbb{R}^d$. Hence

$$\mu_H(aA) = \mu_H(A).$$

So μ_H is left-invariant.

Similarly, right translation

$$R_a(h) = ha$$

also becomes a polynomial triangular map in Mal'cev coordinates with Jacobian determinant 1. Hence

$$\mu_H(Aa) = \mu_H(A).$$

So this same measure is both left- and right-invariant. That is what it means here to say it is the bi-invariant Haar measure.

A quick concrete example is the Heisenberg group. In coordinates (x, y, z), one convention gives

$$(x, y, z)(x', y', z') = (x + x',\ y + y',\ z + z' + xy').$$

For fixed (x, y, z), left translation sends

$$(x', y', z') \mapsto (x + x',\ y + y',\ z + z' + xy').$$

The Jacobian matrix with respect to (x', y', z') is triangular with diagonal entries $1, 1, 1$, so the determinant is 1. Thus $dx\,dy\,dz$ is left-invariant. The same happens for right translations. So ordinary coordinate Lebesgue measure is Haar measure.

Now connect this with the Riemannian metric. Suppose we define a right-invariant Riemannian metric on H by declaring

$$Y_1, \dots, Y_d$$

orthonormal at the identity. Let $E_1, \dots, E_d$ be the corresponding right-invariant vector fields. Then $E_1, \dots, E_d$ is a global orthonormal frame. The Riemannian volume form is the volume form dual to this orthonormal frame.

Because the metric is right-invariant, this Riemannian volume form is right-invariant. Thus the Riemannian volume measure is a right Haar measure.

Since H is nilpotent, it is unimodular, so right Haar and left Haar measures agree. Therefore the Riemannian volume measure is the same as the bi-invariant Haar measure, up to an overall constant.

With the normalization that $Y_1, \dots, Y_d$ is orthonormal and the Mal'cev coordinates are built from exactly the same basis, that constant is 1. So in these coordinates,

$$d\,\mathrm{vol}_g = dt_1 \cdots dt_d.$$

This is actually stronger than the estimate I described earlier using

$$d\,\mathrm{vol}_g = \sqrt{\det(g_{ab}(t))}\,dt.$$

In the right-invariant nilpotent setting with Mal'cev coordinates from the same orthonormal basis, one gets

$$\sqrt{\det(g_{ab}(t))} = 1$$

identically. The coordinate metric matrix $g_{ab}(t)$ is not usually the identity, but its determinant is 1. So lengths and distances can be distorted in coordinates, but volume is not.

For the quotient H/Λ, if $\mathcal{Y}$ is a Mal'cev basis adapted to the lattice Λ, then the fundamental domain is represented by the unit cube

$$[0,1)^d$$

in Mal'cev coordinates, up to boundary issues. The Haar probability measure on H/Λ is the pushforward of Lebesgue measure on this cube. Thus for a measurable set $A \subseteq H/\Lambda$,

$$\mu_{H/\Lambda}(A) = \mathrm{Leb}\big(\psi_{\mathcal{Y}}(\widetilde{A} \cap [0,1)^d)\big),$$

where $\widetilde{A}$ is the lift to H.

So the relationship is:

Lebesgue measure in Mal'cev coordinates on H is Haar measure on H. After quotienting by Λ, Lebesgue measure on a Mal'cev fundamental cube gives Haar probability measure on H/Λ. If the right-invariant Riemannian metric declares the Mal'cev basis orthonormal, its Riemannian volume measure is exactly this same Haar measure.

---

60. User

2026-06-13 16:25:55 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T08:25:55.878000Z

does this mean that the pull back of lebesgue measure on R^m via psi_mathcal{X} is the/a haar measure on G?

---

61. Assistant

2026-06-13 16:25:58 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T08:25:58.012352Z

Yes. In the nilpotent simply connected setting, if

$$\psi_{\mathcal{X}} : G \to \mathbb{R}^m$$

is a Mal'cev coordinate map, then the pullback of Lebesgue measure on $\mathbb{R}^m$ is a Haar measure on G.

More explicitly, define μ_G by

$$\int_G F(g)\, d\mu_G(g) = \int_{\mathbb{R}^m} F(\psi_{\mathcal{X}}^{-1}(t))\, dt.$$

Equivalently, for measurable $A \subseteq G$,

$$\mu_G(A) = \mathrm{Leb}(\psi_{\mathcal{X}}(A)).$$

Then μ_G is Haar measure on G.

The reason is that left and right translations become polynomial maps in Mal'cev coordinates with Jacobian determinant 1. For example, for fixed $a \in G$, left translation gives

$$T_a(t) = \psi_{\mathcal{X}}(a \psi_{\mathcal{X}}^{-1}(t)).$$

In Mal'cev coordinates, T_a is triangular/unipotent, so

$$|\det DT_a(t)| = 1.$$

Therefore Lebesgue measure in coordinates is preserved by left translation. Similarly right translations preserve it. Hence μ_G is both left and right Haar.

So yes:

$$\psi_{\mathcal{X}}^*(dt_1 \cdots dt_m)$$

is a Haar measure on G.

There is only a normalization issue. Haar measure is unique up to scalar. The choice of Mal'cev basis fixes that scalar. If you change the basis, the corresponding coordinate Lebesgue measure may change by a constant factor, but it is still Haar.

For the quotient G/Γ, if $\mathcal{X}$ is adapted to Γ, then the fundamental domain is represented by

$$[0,1)^m$$

in Mal'cev coordinates. The Haar probability measure on G/Γ is the pushforward of Lebesgue measure on this cube. In other words, for nice $F : G/\Gamma \to \mathbb{R}$,

$$\int_{G/\Gamma} F(x)\, d\mu(x) = \int_{[0,1)^m} F(\psi_{\mathcal{X}}^{-1}(t)\Gamma)\, dt.$$

This also matches the Riemannian volume measure coming from the right-invariant metric that declares the Mal'cev basis $\mathcal{X}$ orthonormal at the identity, after the same normalization. The metric volume is right Haar, and since nilpotent groups are unimodular, right Haar equals left Haar.

---

62. User

2026-06-13 16:38:24 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T08:38:24.361000Z

if we normalize the haar measure on G so that a fundamental domain for G/Gamma has measure 1, then can we show that the multiplicative constant between the haar measure on G and the noramlized haar measure on R^m (so that [0,1)^m has measure 1) is bounded by the rationality of the mal cev basis?

---

63. Assistant

2026-06-13 16:38:26 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T08:38:26.835475Z

Yes — in the usual nilmanifold/Mal'cev-basis normalization, the multiplicative constant is actually **exactly** 1.

Let

$$\psi_{\mathcal{X}} : G \to \mathbb{R}^m$$

be Mal'cev coordinates for G/Γ, adapted to Γ, so that

$$F_{\mathcal{X}} := \psi_{\mathcal{X}}^{-1}([0,1)^m)$$

is a fundamental domain for G/Γ.

Define the Haar measure $\mu_{\mathcal{X}}$ on G by pulling back Lebesgue measure:

$$\mu_{\mathcal{X}}(A) = \mathrm{Leb}\big(\psi_{\mathcal{X}}(A)\big).$$

As we discussed, $\mu_{\mathcal{X}}$ is a Haar measure on G. Since

$$\mu_{\mathcal{X}}(F_{\mathcal{X}}) = \mathrm{Leb}([0,1)^m) = 1,$$

this Haar measure is already normalized so that a fundamental domain for G/Γ has measure 1. Therefore it descends to the Haar probability measure on G/Γ.

So in the adapted Mal'cev basis situation,

$$d\mu_G = \psi_{\mathcal{X}}^*(dt_1 \cdots dt_m)$$

with no extra multiplicative factor.

Equivalently, for a function $f : G/\Gamma \to \mathbb{C}$,

$$\int_{G/\Gamma} f(g\Gamma)\, d\mu_{G/\Gamma}(g\Gamma) = \int_{[0,1)^m} f\big(\psi_{\mathcal{X}}^{-1}(t)\Gamma\big)\, dt.$$

The constant is 1, not merely bounded polynomially in the rationality parameter.

The caveat is important: this uses that $\mathcal{X}$ is not merely a rational basis for the Lie algebra, but a Mal'cev basis **adapted to** Γ in the sense that $[0,1)^m$ gives a fundamental domain.

If you choose some other rational basis $\mathcal{Y}$ not adapted to Γ, then $\psi_{\mathcal{Y}}^{-1}([0,1)^m)$ need not be a fundamental domain. In that case

$$\psi_{\mathcal{Y}}^*(dt)$$

is still Haar measure on G, but the covolume

$$\mu_{\mathcal{Y}}(G/\Gamma)$$

need not be 1. The normalized Haar measure is then

$$\mu_G^{\mathrm{norm}} = \frac{1}{\mathrm{covol}_{\mathcal{Y}}(\Gamma)}\psi_{\mathcal{Y}}^*(dt).$$

If $\mathcal{Y}$ is related to an adapted Mal'cev basis $\mathcal{X}$ by a rational change of basis of height $\leq M$, then this covolume factor is typically bounded above and below by $M^{O_m(1)}$. But that bound depends on the rational change of basis relative to the lattice, not merely on the Lie bracket structure constants.

So for your intended setting:

If $\mathcal{X}$ is the original Mal'cev basis for G/Γ, adapted to Γ, then the constant is exactly 1.

If $\mathcal{X}'$ is a Mal'cev basis for a rational subnilmanifold G'/Γ', adapted to $\Gamma' = G' \cap \Gamma$, then again the corresponding cube $[0,1)^d$ has Haar measure 1 on G'/Γ', so the constant is exactly 1.

Polynomial bounds only become necessary when comparing measures/coordinates coming from a non-adapted rational basis to those coming from an adapted Mal'cev basis.

---

64. User

2026-06-13 17:02:32 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T09:02:32.117000Z

is any connected simply connect rational subgroup of G a hyperplane in the malcev coordinates?

---

65. Assistant

2026-06-13 17:02:34 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T09:02:34.388046Z

No, not in general.

There are three related statements, and it is important to keep them separate.

First, in **Lie-algebra exponential coordinates of the first kind**, every connected subgroup is linear. Since G is connected, simply connected, and nilpotent, the Lie exponential

$$\exp : \mathfrak{g} \to G$$

is a global diffeomorphism. If $H \leq G$ is a connected subgroup, then

$$H = \exp(\mathfrak{h})$$

for some Lie subalgebra $\mathfrak{h} \leq \mathfrak{g}$. Therefore, in coordinates

$$g = \exp(t_1 X_1 + \cdots + t_m X_m),$$

the set H is just the linear subspace corresponding to $\mathfrak{h}$. If $\mathfrak{h}$ has codimension 1, then it is a hyperplane.

But this is **not usually what Mal'cev coordinates mean** in the Green–Tao/nilmanifold setting. There one usually uses coordinates of the second kind:

$$g = \exp(t_1 X_1)\exp(t_2 X_2)\cdots\exp(t_m X_m),$$

and defines

$$\psi_{\mathcal{X}}(g) = (t_1, \ldots, t_m).$$

In these coordinates, a connected subgroup $H = \exp(\mathfrak{h})$ need not be a linear subspace, and certainly need not be a hyperplane. It is usually a polynomial submanifold.

A simple example is the Heisenberg group with basis X, Y, Z and

$$[X, Y] = Z.$$

Use Mal'cev coordinates

$$g = \exp(xX)\exp(yY)\exp(zZ).$$

Consider the one-parameter subgroup

$$H = \{\exp(t(X + Y)) : t \in \mathbb{R}\}.$$

This is a connected rational subgroup, since $X + Y$ is a rational linear combination of the Mal'cev basis vectors. But in the above Mal'cev coordinates,

$$\exp(t(X + Y)) = \exp(tX)\exp(tY)\exp\left(-\frac{t^2}{2}Z\right).$$

So

$$\psi_{\mathcal{X}}(H) = \left\{(t, t, -t^2/2) : t \in \mathbb{R}\right\}.$$

That is a parabola, not a line or hyperplane.

What is true is the following.

If $H \leq G$ is a connected rational subgroup, then there exists a Mal'cev basis

$$\mathcal{Y} = (Y_1, \ldots, Y_d)$$

for $H/(H \cap \Gamma)$, and the Y_i's can be chosen to be rational combinations of the ambient basis $\mathcal{X}$, with quantitative control. In the $\mathcal{Y}$-coordinates intrinsic to H, the subgroup H is of course the whole space $\mathbb{R}^d$.

Also, one can often extend $\mathcal{Y}$ to a Mal'cev basis of G. In such an adapted basis, H can become a coordinate subgroup, something like

$$t_{d+1} = \cdots = t_m = 0,$$

depending on conventions. But that is after changing the Mal'cev basis.

So the answer is:

In logarithmic coordinates $g = \exp(\sum t_i X_i)$, connected rational subgroups are rational linear subspaces.

In ordinary Mal'cev product coordinates $\psi_{\mathcal{X}}$, connected rational subgroups are generally polynomial submanifolds, not hyperplanes.

Only special rational subgroups, for example certain kernels of horizontal characters or coordinate subgroups in an adapted Mal'cev basis, appear as literal hyperplanes in the original Mal'cev coordinates.

---

66. User

2026-06-13 20:08:21 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:08:21.641000Z

is there a codimension 1 rational subgroup that is not the kernel of a horizontal character?

---

67. Assistant

2026-06-13 20:08:24 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:08:24.948813Z

It depends on whether you mean the **full kernel** of a horizontal character

$$\eta : G \to \mathbb{R}/\mathbb{Z}$$

or the **identity component** of that kernel.

For a connected simply connected nilpotent Lie group G with lattice Γ, every connected codimension 1 rational subgroup $H \leq G$ is the **identity component of the kernel** of a horizontal character.

Indeed, let

$$\mathfrak{h} = \mathrm{Lie}(H) \le \mathfrak{g}.$$

Since $\mathfrak{h}$ has codimension 1, it is a maximal proper subalgebra. In a nilpotent Lie algebra, every maximal subalgebra is an ideal and contains the commutator algebra:

$$[\mathfrak{g}, \mathfrak{g}] \subseteq \mathfrak{h}.$$

So there exists a nonzero linear functional

$$\ell : \mathfrak{g} \to \mathbb{R}$$

such that

$$\ker \ell = \mathfrak{h},$$

and because $\mathfrak{h} \supseteq [\mathfrak{g}, \mathfrak{g}]$, the functional ℓ vanishes on brackets. Therefore ℓ integrates to a Lie-group homomorphism

$$\chi : G \to \mathbb{R}.$$

Concretely, since G is simply connected nilpotent, one can write

$$\chi(\exp X) = \ell(X).$$

The BCH formula causes no problem because ℓ kills commutators.

Now use rationality. Since H is rational with respect to Γ, the hyperplane $\mathfrak{h}$ is rational relative to the lattice structure. Therefore ℓ can be chosen rational on $\log \Gamma$, and after multiplying by an integer we may arrange

$$\chi(\Gamma) \subseteq \mathbb{Z}.$$

Then

$$\eta(g) = \chi(g) \bmod \mathbb{Z}$$

is a horizontal character $G \to \mathbb{R}/\mathbb{Z}$, and

$$H = \ker \chi = (\ker \eta)^0.$$

So in the connected codimension 1 rational case, there is no counterexample if “kernel” means **connected component of the identity of the kernel**.

But if you mean the **literal full kernel** of $\eta : G \to \mathbb{R}/\mathbb{Z}$, then yes, H is usually not the full kernel. The full kernel contains Γ, so it is typically disconnected.

For example, take

$$G = \mathbb{R}^2, \qquad \Gamma = \mathbb{Z}^2.$$

The connected rational codimension 1 subgroup

$$H = \{(0, y) : y \in \mathbb{R}\}$$

comes from the horizontal character

$$\eta(x, y) = x \bmod \mathbb{Z}.$$

But

$$\ker \eta = \mathbb{Z} \times \mathbb{R},$$

which is a union of parallel lines. Its identity component is

$$(\ker \eta)^0 = \{0\} \times \mathbb{R} = H.$$

So the clean answer is:

A connected codimension 1 rational subgroup of a nilpotent G is always the identity component of the kernel of a horizontal character. It is generally not the full kernel of the horizontal character $G \to \mathbb{R}/\mathbb{Z}$, because that full kernel usually has multiple connected components.

68. User

2026-06-13 20:10:03 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:10:03.915000Z

i did mean the connected component of the identity. Thanks for catching that. If we iterate this, won't this prove that the connected simply connected rational subgroups can be identified with hyperplanes?

---

69. Assistant

2026-06-13 20:10:06 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:10:06.066823Z

Not quite. Iterating proves something useful, but weaker than “all rational subgroups are hyperplanes in the original Mal’cev coordinates.”

The correct statement is more like this:

Given a connected rational subgroup $H \leq G$, one can often build a rational flag

$$G = G_0 \supset G_1 \supset \cdots \supset G_k = H$$

where each G_{j+1} has codimension 1 in G_j, and

$$G_{j+1} = (\ker \eta_j)^0$$

for some horizontal character η_j of G_j, not necessarily of the original ambient group G.

So H can be obtained by **successive horizontal cuts**, but the ambient group changes at every stage.

The distinction matters. A horizontal character of G must vanish on $[G,G]$. Therefore every identity-component kernel of an ambient horizontal character contains $[G,G]$. Hence any intersection of ambient horizontal kernels also contains $[G,G]$. So this can only describe subgroups H with

$$[G,G] \subseteq H.$$

Most rational subgroups do not have this property.

For example, take the Heisenberg group with Lie algebra

$$[X,Y] = Z.$$

The one-dimensional rational subgroup

$$H = \exp(\mathbb{R}X)$$

does not contain $\exp(\mathbb{R}Z) = [G,G]$. Therefore H cannot be an intersection of identity components of kernels of horizontal characters of G.

But it can be obtained by iteration. First take

$$G_1 = \exp(\operatorname{span}(X,Z)).$$

This is codimension 1 in G, and it is the identity component of the kernel of the horizontal character detecting the Y-coordinate. Then inside G_1, which is abelian, take the horizontal character detecting the Z-coordinate. Its identity-component kernel is

$$\exp(\mathbb{R}X) = H.$$

So the second character is horizontal on G_1, but it is not horizontal on G, because the Z-direction lies in the commutator direction of G.

Thus iteration gives:

$$H = \text{successive kernel inside successive subgroups.}$$

It does **not** give:

$$H = \text{intersection of ambient horizontal kernels in } G.$$

There is also a coordinate issue. In a Mal’cev basis adapted to the flag, H can be made into a coordinate subgroup, for example something like

$$t_{d+1} = \cdots = t_m = 0.$$

But in the original Mal'cev coordinates, H need not be a linear subspace or hyperplane intersection. It can be a polynomial submanifold. For instance in the Heisenberg group,

$$H = \{\exp(t(X+Y)) : t \in \mathbb{R}\}$$

is rational, but in product Mal'cev coordinates

$$\exp(xX)\exp(yY)\exp(zZ)$$

it appears as something like

$$(t, t, -t^2/2),$$

up to sign convention. That is a parabola, not a line.

So the clean conclusion is:

A connected rational subgroup can be obtained by successive codimension-one rational cuts, with each cut horizontal relative to the current subgroup. Equivalently, after choosing a Mal'cev basis adapted to the subgroup/flag, it becomes a coordinate subgroup. But it is generally not an intersection of hyperplanes coming from horizontal characters of the original ambient nilmanifold, and it need not look linear in the original Mal'cev coordinates.

---

70. User

2026-06-13 20:12:10 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:12:10.488000Z

i think your reasoning has a gap. When we apply a horizontal character on a subgroup G', the kernel contains [G',G'], not [G,G]. Since the change of basis from mathcal{X} to mathcal{X}' is linear, everything should still work out, right?

---

71. Assistant

2026-06-13 20:12:13 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-13T12:12:13.046318Z

Yes, you are right that my previous reasoning needed a sharper distinction. The iteration argument **does work**, but the conclusion is:

A connected rational subgroup can be obtained by successive codimension-one horizontal cuts inside successively smaller rational subgroups.

It does **not** imply that the subgroup is an intersection of hyperplanes in the original fixed Mal'cev coordinate chart $\psi_{\mathcal{X}} : G \to \mathbb{R}^m$.

The issue is exactly the one you point to: when we pass from G to G', a horizontal character on G' only has to vanish on

$$[G', G'],$$

not on

$$[G, G].$$

So it need not be the restriction of a horizontal character on G. That is why the iteration can reach subgroups that do not contain $[G, G]$.

Here is the clean algebraic statement.

Let G be connected simply connected nilpotent, Γ a lattice, and $H \le G$ a connected rational subgroup. Let

$$\mathfrak{h} = \mathrm{Lie}(H) \le \mathfrak{g}.$$

Then one can build a rational flag

$$\mathfrak{g} = \mathfrak{g}_0 \supset \mathfrak{g}_1 \supset \cdots \supset \mathfrak{g}_k = \mathfrak{h},$$

where each $\mathfrak{g}_{j+1}$ has codimension 1 in $\mathfrak{g}_j$, and each $\mathfrak{g}_{j+1}$ is the kernel of a rational linear functional

$$\ell_j : \mathfrak{g}_j \to \mathbb{R}$$

satisfying

$$\ell_j([\mathfrak{g}_j, \mathfrak{g}_j]) = 0.$$

Equivalently, after exponentiating,

$$G_j = \exp(\mathfrak{g}_j)$$

and

$$G_{j+1} = (\ker \eta_j)^0$$

for a horizontal character

$$\eta_j : G_j \to \mathbb{R}/\mathbb{Z}.$$

The character η_j is horizontal for G_j, not necessarily for the original G.

The proof is simple. Suppose $\mathfrak{h} \subsetneq \mathfrak{g}_j$. Consider the abelianization map

$$\pi_j : \mathfrak{g}_j \to \mathfrak{g}_j/[\mathfrak{g}_j, \mathfrak{g}_j].$$

Because $\mathfrak{g}_j$ is nilpotent, if

$$\mathfrak{h} + [\mathfrak{g}_j, \mathfrak{g}_j] = \mathfrak{g}_j,$$

then actually $\mathfrak{h} = \mathfrak{g}_j$. This is the Lie algebra version of the Frattini property for nilpotent Lie algebras. Therefore, when $\mathfrak{h} \neq \mathfrak{g}_j$, the image $\pi_j(\mathfrak{h})$ is a proper subspace of the abelianization.

Since $\mathfrak{h}$ is rational, $\pi_j(\mathfrak{h})$ is a rational subspace. Choose a rational hyperplane W_j in the abelianization containing $\pi_j(\mathfrak{h})$, and define

$$\mathfrak{g}_{j+1} := \pi_j^{-1}(W_j).$$

Then $\mathfrak{g}_{j+1}$ is a rational codimension-one subalgebra of $\mathfrak{g}_j$, it contains $\mathfrak{h}$, and it contains $[\mathfrak{g}_j, \mathfrak{g}_j]$. Thus it is the kernel of a rational horizontal functional on $\mathfrak{g}_j$. Iterating gives the flag.

So yes: every connected rational subgroup is obtained by iterated horizontal cuts.

The remaining subtlety is coordinate-theoretic.

The change from the ambient basis $\mathcal{X}$ to a subgroup basis $\mathcal{X}'$ is linear **at the Lie algebra level**. But the Mal'cev coordinate maps are usually product coordinates:

$$g = \exp(t_1 X_1) \cdots \exp(t_m X_m).$$

Changing the Lie algebra basis linearly does not induce a linear change of these product coordinates. Because of BCH, the coordinate change is generally polynomial.

That is why a subgroup may be a hyperplane in the current adapted Mal'cev chart, but a curved polynomial submanifold in the original $\psi_{\mathcal{X}}$-coordinates.

For example, in the Heisenberg Lie algebra

$$[X, Y] = Z,$$

take the one-dimensional rational subgroup

$$H = \{\exp(t(X + Y)) : t \in \mathbb{R}\}.$$

In Lie algebra coordinates of the first kind, this is just the line

$$t(X + Y),$$

so it is linear. But in product Mal'cev coordinates

$$g = \exp(xX) \exp(yY) \exp(zZ),$$

one has

$$\exp(t(X + Y)) = \exp(tX) \exp(tY) \exp\left(-\frac{t^2}{2} Z\right),$$

up to the sign convention for the bracket. Thus

$$\psi_{\mathcal{X}}(H) = \{(t, t, -t^2/2) : t \in \mathbb{R}\}.$$

This is a parabola in the original Mal'cev coordinates, not a linear hyperplane intersection.

However, the iterative hyperplane picture still holds. First cut G by the ambient horizontal character corresponding to $x - y$, giving the subgroup with Lie algebra

$$\mathrm{span}(X + Y, Z).$$

Inside that subgroup, which is abelian, apply a horizontal character detecting the Z-direction. Its connected kernel is

$$\mathrm{span}(X + Y).$$

So H is obtained by successive horizontal cuts, but the second cut is horizontal only inside the intermediate subgroup.

Therefore the corrected conclusion is:

Every connected rational subgroup $H \leq G$ can be made into a coordinate subspace after choosing an adapted Mal'cev basis, and it can be obtained by successive horizontal hyperplane cuts in successively smaller rational subgroups.

But H need not be an intersection of hyperplanes in the original fixed Mal'cev coordinates $\psi_{\mathcal{X}}$. In those coordinates, it is generally a polynomial submanifold, with polynomial complexity controlled by the rationality of the adapted basis.

---

72. User

2026-06-16 11:28:11 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-16T03:28:11.874000Z

can you explain to me one more time the difference between the lie bracket exponential and theh riemannian metric exponential assuming that we have a connected simply connected nilpotent lie group which admits both?

---

73. Assistant

2026-06-16 11:28:14 (Asia/Shanghai, UTC+08:00)
Source timestamp: 2026-06-16T03:28:14.977088Z

Yes. The cleanest way to say it is:

The **Lie exponential** is a map coming from the group structure.

The **Riemannian exponential** is a map coming from the metric and its geodesics.

They happen to have the same domain at the identity,

$$T_eG = \mathfrak{g},$$

but they are usually different maps.

1. The Lie exponential

Let G be a connected simply connected nilpotent Lie group with Lie algebra $\mathfrak{g}$. For $X \in \mathfrak{g}$, the Lie exponential is

$$\exp_{\mathrm{Lie}}(X) \in G.$$

It is defined by taking the unique one-parameter subgroup $\gamma_X : \mathbb{R} \to G$ with initial tangent vector X, and setting

$$\exp_{\mathrm{Lie}}(X) = \gamma_X(1).$$

Equivalently,

$$\gamma_X(t) = \exp_{\mathrm{Lie}}(tX).$$

This construction uses only the Lie group structure. In the nilpotent simply connected case, the Lie exponential is especially nice:

$$\exp_{\mathrm{Lie}} : \mathfrak{g} \to G$$

is a global analytic diffeomorphism.

The group law in $\mathfrak{g}$-coordinates is given by the BCH formula:

$$\exp_{\mathrm{Lie}}(X)\exp_{\mathrm{Lie}}(Y)=\exp_{\mathrm{Lie}}(\mathrm{BCH}(X,Y)).$$

Since G is nilpotent, BCH is a finite polynomial involving X, Y and their brackets.

So the Lie exponential is algebraic/group-theoretic. It does not know about lengths, angles, shortest paths, or curvature.

2. The Riemannian exponential

Now suppose G has a Riemannian metric g, for example the right-invariant metric obtained by declaring a Mal'cev basis orthonormal at the identity.

For $p \in G$, the Riemannian exponential is

$$\mathrm{Exp}_p^g : T_pG \to G.$$

Given $v \in T_pG$, let $\gamma_v(t)$ be the Riemannian geodesic satisfying

$$\gamma_v(0)=p, \qquad \dot\gamma_v(0)=v.$$

Then

$$\mathrm{Exp}_p^g(v)=\gamma_v(1).$$

This construction depends on the Riemannian metric. If you change the metric, you change the geodesics, hence you change Exp_p^g.

At the identity e, we can compare

$$\exp_{\mathrm{Lie}} : \mathfrak{g} \to G$$

and

$$\mathrm{Exp}_e^g : T_eG=\mathfrak{g} \to G.$$

They both start with the same tangent behavior:

$$d(\exp_{\mathrm{Lie}})_0=\mathrm{id}_{\mathfrak{g}}, \qquad d(\mathrm{Exp}_e^g)_0=\mathrm{id}_{\mathfrak{g}}\,.$$

So infinitesimally they agree to first order. But globally, and even at second order, they usually differ.

3. Why they differ

The curve

$$t \mapsto \exp_{\mathrm{Lie}}(tX)$$

is a one-parameter subgroup.

The curve

$$t \mapsto \mathrm{Exp}_e^g(tX)$$

is the Riemannian geodesic starting at e with initial velocity X.

These are the same only when the one-parameter subgroup is a geodesic. For a general right-invariant Riemannian metric on a nonabelian nilpotent group, one-parameter subgroups are not usually geodesics.

A useful test is this. Let the metric be right-invariant, and identify tangent vectors with $\mathfrak{g}$ using right translations. If a geodesic has body velocity $u(t) \in \mathfrak{g}$, then $u(t)$ satisfies an Euler-type equation

$$\dot u(t)=\pm\,\mathrm{ad}^*_{u(t)}\,u(t),$$

up to sign convention depending on left/right invariance. A one-parameter subgroup corresponds to constant body velocity $u(t)=X$. So it is a geodesic only if

$$\mathrm{ad}^*_X\,X=0.$$

That condition is not automatic.

4. Heisenberg example

Let $\mathfrak{g}$ be the Heisenberg Lie algebra with orthonormal basis

$$X, Y, Z$$

and bracket

$$[X, Y] = Z.$$

Use the right-invariant metric declaring X, Y, Z orthonormal at the identity.

Take

$$U = X + Z.$$

The Lie exponential curve is

$$t \mapsto \exp_{\mathrm{Lie}}(tU).$$

This is a one-parameter subgroup.

But it is not generally a Riemannian geodesic. Indeed, with respect to the orthonormal basis, one computes

$$\mathrm{ad}_U^* U = Y,$$

up to the relevant sign convention. This is nonzero. Therefore the one-parameter subgroup generated by U is not a geodesic for this metric.

So

$$\mathrm{Exp}_e^g(U) \neq \exp_{\mathrm{Lie}}(U)$$

in general.

5. Why nilpotence helps only for the Lie exponential

Nilpotence gives you a wonderful fact:

$$\exp_{\mathrm{Lie}} : \mathfrak{g} \to G$$

is a global diffeomorphism.

That means every element $g \in G$ has a unique logarithm

$$g = \exp_{\mathrm{Lie}}(X).$$

It also means Mal'cev coordinates built from Lie exponentials give global coordinates on G.

But this says nothing directly about the Riemannian exponential. The Riemannian exponential depends on geodesics of the chosen metric, and those geodesics are not generally one-parameter subgroups.

For a right-invariant Riemannian metric on a connected Lie group, geodesics are complete, so

$$\mathrm{Exp}_e^g$$

is defined on all of T_eG. But it need not be globally injective, and it need not equal the Lie exponential.

6. How this relates to Mal'cev coordinates

When you write

$$g = \exp_{\mathrm{Lie}}(t_1 X_1) \cdots \exp_{\mathrm{Lie}}(t_m X_m),$$

these are Mal'cev coordinates of the second kind. They are built using the Lie exponential.

They are not Riemannian normal coordinates.

Riemannian normal coordinates at e would instead be

$$g = \mathrm{Exp}_e^g(v), \qquad v \in T_eG.$$

Both coordinate systems are valid near e, and in the nilpotent case the Mal'cev coordinates are valid globally. But they are different coordinate systems.

For your purposes, this is why one can use Mal'cev coordinates to control the group law, lattice, Haar measure, and rationality, while using the right-invariant Riemannian metric to talk about distances and balls. The two structures are compatible locally because both maps have derivative equal to the identity at 0, but they are not the same map.