

Controller Synthesis through Riemannian Optimization

Spencer Kraisler

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Mehran Mesbahi, Chair

Behçet Açıkmese

Amir Taghvaei

Martin Nisser

Jing Wu

Program Authorized to Offer Degree:
William E. Boeing Department of Aeronautics & Astronautics

©Copyright 2026

Spencer Kraisler

University of Washington

Abstract

Controller Synthesis through Riemannian Optimization

Spencer Kraisler

Chair of the Supervisory Committee:

Mehran Mesbahi

William E. Boeing Department of Aeronautics & Astronautics

This dissertation studies controller synthesis for both open-loop and closed-loop systems through the usage of Riemannian optimization. In the closed-loop setting, the space of dynamic output-feedback controllers is formulated as a Riemannian orbit manifold, enabling fast and intrinsic optimization without reliance on extrinsic parameterizations. We establish topological characterizations of this space and provide convergence rate guarantees for the resulting algorithms. For open-loop control, we develop an intrinsic successive convexification framework that optimizes trajectories directly on the manifold of states and inputs, leading to improved numerical behavior in applications such as attitude guidance. This work demonstrates how exploiting the intrinsic manifold structure of control problems yields fast theoretically principled algorithms across open- and closed-loop control paradigms.

To my wife, parents, and mentors,

*The theory of convex functions reveals all its power and consequences only when it is
conceived on a Riemannian structure.*

Constantin Udriste [1]

TABLE OF CONTENTS

	Page
List of Figures	iv
Chapter 1: A Coordinate-Free Introduction	1
1.1 Outline of Dissertation	4
1.2 Basic Notation	5
Chapter 2: Background	6
2.1 What is a Smooth Manifold?	7
2.1.1 Defining Smoothness Rigorously	7
2.1.2 Tangent Spaces and Vector Fields	11
2.1.3 Derivatives	15
2.1.4 Parameterizations	17
2.1.5 Coordinates and Frames: Computation	19
2.2 What is a Riemannian Manifold?	22
2.2.1 Rank-2 Tensor Fields	22
2.2.2 Riemannian Metric	23
2.2.3 Levi-Civita Connection	24
2.2.4 Riemannian Gradient And Hessian	26
2.2.5 Geodesics	28
2.2.6 Retractions: Navigating Across the Manifold	31
2.2.7 Riemannian Gradient Descent	32
Chapter 3: The Space of Controllers is a Riemannian Manifold and How to Optimize Over It	33
3.1 LQG Policy Optimization	35
3.2 Geometry of Linear Systems	38
3.2.1 Spaces of Linear Systems	38

3.2.2	Spaces of Linear Controllers	40
3.2.3	Topology of Controller Spaces	41
3.3	First-Order Optimization over the Controller Orbit Space	45
3.3.1	Riemannian Gradient Descent	47
3.3.2	Computing the Riemannian Gradient	47
3.3.3	Why Gradient Descent does not Work	49
3.3.4	Convergence Analysis	49
3.3.5	Numerical Results	51
3.4	Second-Order Optimization over the Controller Orbit Space	53
3.4.1	Why Newton's Descent Does not Work	53
3.4.2	Riemannian Newton Descent Over the Controller Orbit Manifold	54
3.4.3	Convergence Analysis	55
3.4.4	Numerical Results	57
3.5	Concluding Remarks	60
Chapter 4:	Trajectory Optimization over Smooth Manifolds: Exploiting Ge-	
	ometry	61
4.1	Problem Statement	63
4.2	SCvx	64
4.2.1	The Main Algorithm	64
4.2.2	Issues with SCvx: non-Euclidean States and Inputs	68
4.3	Intrinsic SCvx	70
4.3.1	Linearizing Dynamics on Smooth Manifolds	70
4.3.2	Taylor Series on Riemannian Manifolds	74
4.3.3	Local Objective: Exploiting Geodesic Convexity	75
4.3.4	Slack Variables and Penalty	75
4.3.5	The Local Sub-problem	76
4.3.6	Coordinate parameterization: Computation	77
4.3.7	The i-SCvx Algorithm	80
4.4	Example: Constrained Attitude Path-planning via i-SCvx	83
4.4.1	Numerical Results	86
4.5	Concluding Remarks	89

Chapter 5: <i>Actual</i> Average Consensus over Lie Groups	91
5.1 Problem Statement	93
5.2 The Average: Riemannian Manifold Case	94
5.3 Consensus Reformulation and Distributed Optimization	97
5.4 Distributed Optimization: Euclidean Case	98
5.5 Distributed Optimization: Lie Group Case	102
5.6 Main Results	105
5.7 Simulations	111
5.7.1 Scenario 1 (Comparison of Consensus Algorithms)	111
5.7.2 Scenario 2 (Linear rate of convergence)	114
5.8 Concluding Remarks	117
Bibliography	118
Appendix A: Lie Theory	129
Appendix B: A Necessary Background on Topology	134
Appendix C: A Necessary Background on Orbit Geometry	138
C.1 Smooth Orbit Geometry	139
C.2 Riemannian Orbit Geometry	142
Appendix D: Smooth Manifolds: An Extrinsic Approach	146
D.1 Smooth Manifolds	147
D.2 Smooth Maps and Derivatives	149
D.3 Vector Fields	151
D.4 Retraction	152
D.5 Riemannian Manifolds	153
D.6 Riemannian Gradients and Hessians	154
D.7 Local Frames and Components	155
Appendix E: Fun Stuff 1: An Algebraic Definition of Smooth Manifolds	158
Appendix F: Fun Stuff 2: Motivation for the Levi–Civita Connection’s Definition	162

LIST OF FIGURES

Figure Number		Page
2.1	Visualization of RGD; here, $x_2 = R_{x_1}(-s_1 \nabla f(x_1))$	32
3.1	Example of $\tilde{\mathcal{C}}_1$ for the plant $(A, B, C) = (1.1, 1, 1)$	36
3.2	Illustration of $\tilde{\mathcal{C}}$ and its orbits. Descent directions must be orthogonal to the orbit through the current iterate.	37
3.3	Comparison of RGD vs. GD for LQG PO for four distinct systems. .	52
3.4	Comparison of RGD vs. orbit RND for LQG PO for four distinct systems.	59
4.1	Visualization of perturbing a trajectory along a tangent space.	77
4.2	The iterates of SCvx and i-SCvx for the constrained attitude guidance problem with parameters $N = 30$, $\Delta t = 0.05$, and $\theta_{\max} = 20^\circ$	88
5.1	The Consensus Error and the RCM Error at each iteration, comparing Algorithm 9 (in blue) with three other algorithms.	115
5.2	Statistical visualization of the Consensus error and RCM error of our proposed 9 over 100 randomly sampled problem instances.	116

Chapter 1

A COORDINATE-FREE INTRODUCTION

The design of control policies often reduces to solving smooth optimization problems whose feasible sets exhibit intrinsic geometric structures. When these sets are treated *extrinsically*, through specified coordinates or embeddings, one risks introducing extra constraints, extra variables, and extra numerical degeneracies; worst of all, one risks losing insight into the intrinsic geometric structures of the problem at hand. With that in mind, the central premise of this dissertation is that many spaces arising in control theory are fundamentally *Riemannian*: their structures are not artifacts of a chosen representation but rather they persist under any reparameterization of the problem. Moreover, working within this coordinate-free philosophy yields algorithms and frameworks whose numerical behavior, performance, and theoretical guarantees are invariant under reparameterization.

Let's walk through an example involving the set $\mathcal{H}^n$ of $n \times n$ Hurwitz stable real matrices. This set appears ubiquitously in control design since the stability of *some* closed-loop state matrix is imposed as a hard constraint. This is rather unfortunate as $\mathcal{H}^n$ is a nonconvex cone. However, this non-convexity is misleading because $\mathcal{H}^n$ admits a lossless transformation into a product space with convexity-like structure.

Fix $Q \succ 0$. An elementary control fact is $A \in \mathcal{H}^n$ if and only if there exists a unique $P \succ 0$ solving the *Lyapunov equation*

$$AP + PA^\top = -Q.$$

Define the maps

$$\Phi_{\text{Sym}}(A) := P, \quad \Phi_{\text{Skew}}(A) := AP + \frac{1}{2}Q.$$

Intuitively, the first map captures the symmetric portion of A through its Lyapunov solution, while the second captures the complementary skew-symmetric portion. Together they define a diffeomorphism

$$\Phi(A) := (\Phi_{\text{Sym}}(A), \Phi_{\text{Skew}}(A)), \quad \Phi : \mathcal{H}^n \rightarrow \text{Sym}_{++}(n) \times \text{Skew}(n).$$

The product manifold

$$\mathcal{M} := \text{Sym}_{++}(n) \times \text{Skew}(n)$$

of the positive definite and skew-symmetric matrices admits a canonical Riemannian structure. On $\text{Sym}_{++}(n)$, define the *affine-invariant metric*

$$g_P^{\text{Sym}}(V, W) := \text{tr}(P^{-1}VP^{-1}W), \quad P \succ 0, \quad V, W \in \text{Sym}(n);$$

we equip $\text{Skew}(n)$ with the *Frobenius inner product*

$$g^{\text{Skew}(n)}(V, W) := \text{tr}(V^\top W), \quad V, W \in \text{Skew}(n)$$

Finally, we equip $\mathcal{M}$ with the product metric $g := (g^{\text{Sym}}, g^{\text{Skew}})$. Since Φ is a diffeomorphism, g may be pulled back to $\mathcal{H}^n$, endowing it with a Riemannian structure. The resulting Riemannian manifold $(\mathcal{H}^n, \Phi^*g)$ admits several favorable geometric properties:

1. $(\mathcal{H}^n, \Phi^*g)$ is geodesically complete;
2. $(\mathcal{H}^n, \Phi^*g)$ is strongly geodesically convex;
3. geodesics on $(\mathcal{H}^n, \Phi^*g)$ have closed-form expressions.

This example motivates the epigraph of this dissertation, drawn from Udriste:

“The theory of convex functions reveals all its power and consequences only when it is conceived on a Riemannian structure.”

The observation raises the question of which other spaces arising in control theory admit similar geometric liftings; this is precisely the motivation behind *extended convex liftings* [2, 3], a line of work formalizing the lossless embedding of nonconvex control design spaces into convex extensions.

Motivated by this perspective, this dissertation develops geometry-aware methods for control synthesis in both the open-loop and closed-loop settings. We formulate control problems directly on the underlying manifolds that describe admissible policies and trajectories. This viewpoint enables the systematic use of geodesics, natural gradients and Hessians, orbit geometry, geodesic convexity, and intrinsic convexification, leading to algorithms that are numerically robust, theoretically principled, and fast.

1.1 *Outline of Dissertation*

This dissertation is organized as follows.

Chapter 2 provides the geometric and analytic background used throughout the dissertation, including smooth manifolds, Riemannian metrics, gradients, Hessians, geodesics, and optimization on manifolds.

Chapter 3 studies direct policy optimization for linear closed-loop controller synthesis. The space of dynamic output-feedback controllers with fixed state dimension is modeled as a Riemannian orbit manifold. By equipping this manifold with the Krishnaprasad–Martin metric, we derive intrinsic first- and second-order optimization algorithms and establish linear and quadratic convergence rates for LQG performance optimization.

Chapter 4 develops an intrinsic formulation of successive convexification for trajectory optimization problems whose states and inputs evolve on Riemannian manifolds, rather than linear parameter spaces. The proposed method performs linearization and convexification in an intrinsic, coordinate-free way, yielding an algorithm whose performance is invariant of parameterization and solves optimal control problems faster and more robustly than classical extrinsic SCvx schemes in such manifold set-ups.

Chapter 5 addresses distributed consensus for agents evolving on bi-invariant Lie groups. We formulate consensus as the distributed computation of the Frechet mean and derive a distributed algorithm based on distributed gradient descent with gradient tracking.

The appendix contains supplementary material used in this dissertation. In §A, we review Lie theory. In §B, we give a background in topology necessary for smooth manifolds. In §C, we go over smooth and Riemannian orbit geometry used in the analysis of controller orbit spaces. In §D, we provide an extrinsic definition of differential geometry. §E presents an algebraic definition of smooth manifolds. Finally, §F presents a motivation behind the definition of the Levi–Civita connection.

1.2 Basic Notation

The n -sphere is defined as $\mathcal{S}^n := \{p \in \mathbb{R}^{n+1} : \|p\|_2 = 1\}$. The set of $n \times n$ real Hurwitz stable matrices is $\mathcal{H}^n$. The Lie group of $n \times n$ invertible matrices is denoted $\text{GL}(n)$. The Lie group of proper orthogonal matrices is denoted $\text{SO}(n)$. The quaternion algebra is denoted $\mathbb{H}$. The set of unit quaternions will be denoted with $\mathcal{Q} \subset \mathbb{H}$. The vector space of pure quaternions are denoted with $\mathfrak{q}$. The vector space of symmetric and skew-symmetric matrices is denoted $\text{Sym}(n)$ and $\text{Skew}(n)$ respectively. The cone of positive definite and positive semi-definite matrices is denoted $\text{Sym}_{++}(n)$ and $\text{Sym}_+(n)$ respectively. Similar for negative definite and negative semi-definite matrices. The set of complex numbers with positive real part and non-negative real part will be denoted $\mathbb{C}_{++}$ and $\mathbb{C}_+$, respectively. Similar for negative and non-positive real part. The imaginary numbers are denoted $j\mathbb{R}$. The set of rational functions is denoted $\mathcal{R}$. The set of strictly proper rational functions with stable (i.e., $\mathbb{C}_{--}$) poles is denoted $\mathcal{RH}_2$. The set of proper rational functions with stable poles is denoted $\mathcal{RH}_\infty$.

Chapter 2

BACKGROUND

The abstraction [of charts and atlases] is necessary to fully appreciate the depth of optimization on manifolds as more than just a fancy tool for constrained optimization.

Nicholas Boumal [4]

For readers with an understanding of real analysis and topology, we recommend [5,6] as texts. For readers less familiar with those topics but comfortable with calculus and linear algebra, [4,7] provide accessible introductions. Those with prior exposure to differential geometry may appreciate texts that examine the philosophical and historical motivation behind the development of Riemannian geometry. For this, [8,9] are very insightful; in particular, Weyl's manuscript was especially influential in developing the coordinate-free viewpoint modern geometers use today. More advanced treatments can be found in [10–12].

2.1 What is a Smooth Manifold?

What is smoothness? A thing is smooth if it can be locally fitted with a linear thing. In other words, locally, the thing is seemingly linear even though globally it is not. A smooth function at any point can be approximated with a linear function. It follows that a smooth space at any point should be approximated with a linear space.

With that in mind, a *smooth manifold* is a space of points which admits, at any point, a *local coordinate system* defined on some neighborhood about that point; furthermore, for any sufficiently close pair of points, the translation of their respective coordinate systems is smooth (in the classical sense) on the intersection of their neighborhoods. The following example illustrates this.

Example 2.1.1. Consider the space $\mathcal{M} \subset \mathcal{RH}_2$ of stable transfer functions with 2 zeroes and 3 poles

$$P : s \mapsto \frac{(s - z_1)(s - z_2)}{(s - p_1)(s - p_2)(s - p_3)}, \quad s \in \mathbb{C}_{++}$$

with no pole-zero cancellations. This space can be locally parameterized by 5 coordinates: $x = (z_1, z_2, p_1, p_2, p_3) \in \mathbb{R}^5$, giving us $P(x) \in \mathcal{M}$. This parameterization is not global however because if we fix all parameters but z_1 and steer z_1 towards p_1 , then we approach a degeneracy in the form of a pole-zero cancellation. Given a specific non-degenerate $x \in \mathbb{R}^5$, there are 5 directions we can infinitesimally perturb $P(x)$ and still get another such transfer function. So, the space of such transfer functions can be locally parameterized by 5 coordinates, and is therefore a smooth manifold.

2.1.1 Defining Smoothness Rigorously

All of differential geometry stems from the small and simple concept known as a chart. Let $(\mathcal{M}, \tau)$ be a second-countable Hausdorff space¹. A *chart of dimension n*

¹See §B.

is a pairing (U, x) , where $U \subset \mathcal{M}$ is open and

$$x : U \rightarrow x(U) \subset \mathbb{R}^n$$

is a homeomorphism from U onto an open subset $x(U)$ of $\mathbb{R}^n$. We call U the *chart domain* and x the *local coordinates*. For $p \in U$, we call

$$x(p) = (x^1(p), \dots, x^n(p)) \in \mathbb{R}^n$$

the *coordinates* of p . We often write $x \equiv (x^i)$.

A collection of charts

$$\mathcal{A} := \{(U_\alpha, x_\alpha)\}$$

whose domains cover $\mathcal{M}$ is called an *atlas*. A space $(\mathcal{M}, \tau)$ admitting an atlas is called a *topological manifold of dimension n* . The existence of an atlas implies that the space is locally homeomorphic to Euclidean space.

Proposition 2.1.1 (Dimensional Invariance). *Let $\mathcal{M}$ be a second-countable Hausdorff space that admits an atlas $\mathcal{A}$. Then, all charts in $\mathcal{A}$ must have the same dimension. Furthermore, $\mathcal{M}$ cannot admit any chart of a different dimension.*

The following example shows that not every second-countable Hausdorff space admits an atlas.

Example 2.1.2 (Topological Manifold Non-example). *The subspace $\{(x, y) \in \mathbb{R}^2 : xy = 0\}$ equipped with the subspace topology does not admit an atlas and therefore is not a topological manifold.*

Let (U, x) and (V, y) be charts in an atlas $\mathcal{A}$. Their associated *coordinate transformation*

$$x \circ y^{-1} : y(U \cap V) \subset \mathbb{R}^n \rightarrow x(U \cap V) \subset \mathbb{R}^n$$

is defined wherever $U \cap V \neq \emptyset$. If this mapping is smooth (in the classical sense), we say the charts are *smoothly compatible*. Note that two charts are smoothly compatible

if and only if their coordinate transformation is a *diffeomorphism* (a smooth invertible map with smooth inverse).

An atlas is called *smooth* if all of its charts are pairwise smoothly compatible. This ensures that differentiability of any construction is independent of the choice of coordinates. Intuitively, the existence of a smooth atlas means that the space is locally *diffeomorphic* to Euclidean space.

Example 2.1.3 (Smooth Atlas on $\mathcal{S}^2$). *For $\mathcal{S}^2$, define*

$$U_i^+ := \{p \in \mathcal{S}^2 : p_i > 0\}, \quad U_i^- := \{p \in \mathcal{S}^2 : p_i < 0\},$$

and coordinate maps

$$x_1(p) := (p_2, p_3), \quad x_2(p) := (p_3, p_1), \quad x_3(p) := (p_1, p_2).$$

The six charts (U_i^+, x_i) and (U_i^-, x_i) form an atlas, and hence $\mathcal{S}^2$ is a topological manifold. Furthermore, this atlas is smooth.

Given a chart (U, x) , any restriction $(U', x|_{U'})$ with $U' \subset U$ open is also a chart. More generally, charts may be modified by composing with smooth maps or rescaling coordinates. This motivates the notion of *maximality*.

A smooth atlas $\mathcal{A}$ is called a *smooth structure* (or a *maximal smooth atlas*) if it is not properly contained in any other smooth atlas. That is, $\mathcal{A}$ is a smooth structure if for any other smooth atlas $\mathcal{B}$ whose charts are additionally pairwise smoothly compatible with the charts in $\mathcal{A}$, we necessarily have $\mathcal{B} \subset \mathcal{A}$. We have the following lemma on the maximality property of smooth atlases:

Lemma 2.1.1. *Let $\mathcal{A}$ be a smooth atlas. Then there exists a unique maximal smooth atlas $\mathcal{A}^+$ in which $\mathcal{A} \subset \mathcal{A}^+$.*

A smooth manifold is a triple $(\mathcal{M}, \tau, \mathcal{A}^+)$, consisting of a second-countable Hausdorff space $(\mathcal{M}, \tau)$ equipped with a smooth structure $\mathcal{A}^+$. By Lemma 2.1.1, a smooth

atlas $(\mathcal{M}, \mathcal{A})$ is also considered a smooth manifold, implying that we are essentially working with the uniquely induced maximal smooth atlas $\mathcal{A}^+$. The modern formulation of charts and atlases originates in the work of Weyl [13] on Riemann surfaces.

Henceforth, we assume $\mathcal{M}$ is a smooth manifold of dimension n unless otherwise stated.

Example 2.1.4 (Two Smooth Atlases on $\mathbb{R}$). *Let $\mathbb{R}$ be equipped with the Euclidean topology and the single-chart atlas $\{(\mathbb{R}, x)\}$, where $x(p) = p$. This induces the standard Euclidean smooth structure used in calculus. Under this smooth structure, smooth functions are precisely the ones considered smooth in the classical sense.*

Next, take the atlas $\{(\mathbb{R}, x^3)\}$, where $x^3(p) := p^3$. This is another smooth atlas on $\mathbb{R}$ distinct from the Euclidean one. Furthermore, they are sub-atlases of distinct smooth structures.

Smooth Maps Let $\mathcal{M}$ and $\mathcal{N}$ be smooth manifolds and let $f : \mathcal{M} \rightarrow \mathcal{N}$. The map f is said to be *smooth* at $p \in \mathcal{M}$ if there exist charts (U, x) about p and (V, y) about $f(p)$ with $f(U) \subset V$ such that

$$y \circ f \circ x^{-1} : x(U) \rightarrow y(V)$$

is smooth at $x(p)$ in the classical sense. If f is smooth at every point, it is called *smooth*. We wish to emphasize that the trick to defining *smooth constructions* in differential geometry is, essentially, mapping those constructions onto the chart co-domains and checking whether the mapped “coordinatized” constructions are smooth in the classical sense.

Smooth maps are the morphisms of the smooth manifold category. The set of *smooth scalar fields* $f : \mathcal{M} \rightarrow \mathbb{R}$ is denoted $\mathfrak{F}(\mathcal{M})$. The space $\mathfrak{F}(\mathcal{M})$ is an algebra under pointwise addition and multiplication.

A smooth map $F : \mathcal{M} \rightarrow \mathcal{N}$ that is bijective with smooth inverse is called a *diffeomorphism*. Diffeomorphisms are the isomorphisms of the smooth manifold cat-

egory; if one exists, we write $\mathcal{M} \cong \mathcal{N}$. We designate $\mathfrak{F}(\mathcal{M}, \mathcal{N})$ as the set of smooth mappings and $\text{Diff}(\mathcal{M}, \mathcal{N})$ as the group of diffeomorphisms, should any exist.

Remark 2.1.1. *If $\mathcal{A}$ and $\mathcal{B}$ are two smooth structures on $\mathcal{M}$, then $(\mathcal{M}, \mathcal{A})$ and $(\mathcal{M}, \mathcal{B})$ are homeomorphic but need not be diffeomorphic.*

If $F : \mathcal{M} \rightarrow \mathcal{N}$ is a diffeomorphism onto its image, it is called a *smooth embedding*. It is the smooth analog of a topological embedding. A particularly simple class of smooth maps are *curves*, which are smooth mappings

$$\gamma : I \subset \mathbb{R} \rightarrow \mathcal{M}$$

defined on intervals of $\mathbb{R}$. Henceforth, all curves are assumed smooth.

2.1.2 Tangent Spaces and Vector Fields

One of the most fundamental concepts in differential geometry is the tangent space. The tangent space is a vector space that encodes the set of *first-order variations* of smooth mappings, such as smooth scalar fields. While one may embed a smooth manifold into an ambient Euclidean space and define tangent spaces extrinsically as tangent subspaces — as justified by Thm. 2.1.4 — we instead adopt an intrinsic construction that makes no reference to an embedding.

Let $\mathcal{V}$ be a real finite-dimensional vector space. Let $f \in \mathfrak{F}(\mathcal{V})$. A first-order variation ξ of f at a point $p \in \mathcal{V}$ is the rate of change of f at p when p is perturbed along some vector $v \in \mathcal{V}$. That is,

$$\xi : f \mapsto \lim_{t \rightarrow 0} \frac{f(p + tv) - f(p)}{t}.$$

Along this line, we can think of a first-order variation as a mapping $\xi : \mathfrak{F}(\mathcal{V}) \rightarrow \mathbb{R}$ with the above definition. Such mappings have two identifying properties:

1. $\xi : \mathfrak{F}(\mathcal{V}) \rightarrow \mathbb{R}$ is $\mathbb{R}$ -linear,

2. ξ satisfies the product rule: $\xi[fg] = f(p)\xi[g] + g(p)\xi[f]$ for $f, g \in \mathfrak{F}(\mathcal{V})$.

Also remark that any such mapping with these two properties has a one-to-one correspondence to some vector in $\mathcal{V}$. As a result, we can take those two identifying properties and lift them to a smooth manifold setting. This brings us to the definition of a derivation.

A *p-derivation* is an $\mathbb{R}$ -linear mapping

$$\xi : \mathfrak{F}(\mathcal{M}) \rightarrow \mathbb{R}$$

satisfying the *product* rule

$$\xi[fg] = f(p)\xi[g] + g(p)\xi[f], \quad f, g \in \mathfrak{F}(\mathcal{M}).$$

The vector space of all p -derivations is called the *tangent space* of $\mathcal{M}$ at p and is denoted $T_p\mathcal{M}$. We will write $\xi f \equiv \xi[f]$ when convenient. Intuitively, ξf represents the directional derivative of f at p along a perturbation ξ .

Elements of $T_p\mathcal{M}$ are called *tangent vectors at p*. It is a standard result that $\dim T_p\mathcal{M} = \dim \mathcal{M}$ for all $p \in \mathcal{M}$. In contrast, a purely topological manifold does not admit tangent spaces, as the definition relies on the smooth structure.

Tangent Spaces of Vector Spaces If $\mathcal{V}$ is a finite-dimensional real vector space, then the tangent space at any point $p \in \mathcal{V}$ can be canonically identified with $\mathcal{V}$ itself. In this setting, it is useful to distinguish between *tangent vectors* (derivations) and *geometric vectors* (elements of $\mathcal{V}$).

Theorem 2.1.1 (Canonical identification). *Let $\mathcal{V}$ be a finite-dimensional real vector space, $p \in \mathcal{V}$, and $\xi \in T_p\mathcal{V}$. There exists a unique vector $v \in \mathcal{V}$ such that*

$$\xi f = \lim_{t \rightarrow 0} \frac{f(p + tv) - f(p)}{t}, \quad f \in \mathfrak{F}(\mathcal{V}).$$

The mapping $\xi \mapsto v$ is linear and invertible.

Tangent Bundle The *tangent bundle* of $\mathcal{M}$ is the disjoint union

$$T\mathcal{M} := \bigcup_{p \in \mathcal{M}} \{p\} \times T_p\mathcal{M}.$$

Strictly speaking, elements of $T\mathcal{M}$ are pairs (p, ξ) rather than tangent vectors alone. When no ambiguity arises, we will identify elements of $T\mathcal{M}$ with tangent vectors.

Theorem 2.1.2. *The smooth structure of $\mathcal{M}$ induces a smooth structure on $T\mathcal{M}$ with dimension $2n$.*

We will not present a proof of the above, but we will present a process for constructing a smooth atlas. First, remark $T\mathcal{M}$ is a second-countable Hausdorff space. Let $\pi : T\mathcal{M} \rightarrow \mathcal{M}$ be the left-projection mapping $\pi(p, \xi) := p$. For each chart (U, x) in $\mathcal{M}$, define

$$\tilde{x} : \pi^{-1}(U) \rightarrow x(U) \times \mathbb{R}^n, \quad \tilde{x}(p, \xi) := (x(p), (\xi[x^1] \ \xi[x^2], \dots, \xi[x^n])).$$

We claim $\tilde{x}$ is a homeomorphism. It follows $(\pi(U), \tilde{x})$ is a $2n$ -dimensional chart on $T\mathcal{M}$. Using this technique, we can construct an atlas on $T\mathcal{M}$. Unfortunately, proving that these charts are additionally pairwise smoothly compatible would require machinery from §2.1.3; nonetheless, this constructs of $2n$ -dimensional smooth atlas for $T\mathcal{M}$.

Vector Fields and Derivations A *vector field* on $\mathcal{M}$ is a smooth mapping

$$V : \mathcal{M} \rightarrow T\mathcal{M}$$

such that $V_p \in T_p\mathcal{M}$ for each $p \in \mathcal{M}$. We denote the space of all vector fields by $\mathfrak{X}(\mathcal{M})$ and write V_p instead of $V(p)$ to emphasize that the value is a tangent vector.

Vector fields may be equivalently defined through global derivations. A *global derivation* is an $\mathbb{R}$ -linear mapping between the algebra of smooth scalar fields

$$\tilde{V} : \mathfrak{F}(\mathcal{M}) \rightarrow \mathfrak{F}(\mathcal{M})$$

which satisfies the product rule

$$\tilde{V}(fg) = f \cdot \tilde{V}(g) + g \cdot \tilde{V}(f), \quad f, g \in \mathfrak{F}(\mathcal{M}).$$

Theorem 2.1.3. *There is a canonical linear isomorphism between the vector space of vector fields $\mathfrak{X}(\mathcal{M})$ and the vector space of global derivations. Given $V \in \mathfrak{X}(\mathcal{M})$, the associated derivation $\tilde{V}$ is defined by*

$$(\tilde{V}(f))(p) := V_p f.$$

In this work, we will identify vector fields with their global derivations; we write $Vf \equiv V[f] \equiv \tilde{V}(f)$.

Compositions of Vector Fields and the Lie Bracket Let $V, W \in \mathfrak{X}(\mathcal{M})$ with associated global derivations $\tilde{V}$ and $\tilde{W}$. Their composition is the linear mapping

$$VW := \tilde{V} \circ \tilde{W} : \mathfrak{F}(\mathcal{M}) \rightarrow \mathfrak{F}(\mathcal{M}),$$

that is, $(VW)(f) = V[Wf]$. Although VW is $\mathbb{R}$ -linear, it fails to be a global derivation because it does not satisfy the product rule. As a result, there is no vector field associated with VW . This holds in the classical Euclidean space too. Intuitively, VW represents the operator that applies the directional derivative along W followed by the directional derivative along V . In essence, VW is a second-order partial derivative.

So, while VW fails to be a vector field, the difference of VW and WV is a vector field:

$$[V, W] := VW - WV \in \mathfrak{X}(\mathcal{M}).$$

We call the operator

$$[\cdot, \cdot] : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M})$$

the *commutator* or *Lie Bracket*, and we call $[V, W]$ the *Lie derivative* of V along W .

Remark 2.1.2. The space $\mathfrak{X}(\mathcal{M})$ is an algebra under pointwise addition and the Lie bracket.

Remark 2.1.3. Let $\mathcal{V}$ be a real vector space, $V, W \in \mathfrak{X}(\mathcal{V})$, $f \in \mathfrak{F}(\mathcal{V})$ and $p \in \mathcal{V}$. Then

$$\begin{aligned}(VW)[f](p) &= D^2 f_p(V_p, W_p) + Df_p(DW_p(V_p)) \\ [V, W][f](p) &= Df_p(DW_p(V_p)) - Df_p(DV_p(W_p))\end{aligned}$$

where $D(\cdot)$ is the classical differential operator (also known as the Frechet derivative), and we are identifying V and W as geometric vector fields.

Vector fields V and W are said to *commute* if $[V, W] = 0$.

2.1.3 Derivatives

We now have the machinery to define notions of derivatives. Let $\mathcal{V}, \mathcal{W}$ be real vector spaces. Let $U \subset \mathcal{V}$ and $f : U \rightarrow \mathcal{W}$ be smooth. The differential of F at p is defined as

$$DF_p : \mathcal{V} \rightarrow \mathcal{W}, \quad Df_p(v) := \lim_{t \rightarrow 0} \frac{F(p + tv) - F(p)}{t}.$$

Intuitively, DF_p maps perturbations of p to perturbations of $F(p)$. This concept can be generalized to the smooth manifold case, however we will need to define the operation without relying on $p + tv$, as we cannot add tangent vectors to points on smooth manifolds. To work around this, we instead use curves $\gamma(\cdot)$ that satisfy $\gamma(0) = p$ and $\dot{\gamma}(0) = v$.

Given a smooth curve $\gamma : I \subset \mathbb{R} \rightarrow \mathcal{M}$, where I is an interval of $\mathbb{R}$, its *time derivative* or *velocity* is a tangent vector $\dot{\gamma}(t) \in T_{\gamma(t)}\mathcal{M}$ defined as

$$\dot{\gamma}(t)f := \lim_{\epsilon \rightarrow 0} \frac{f(\gamma(t + \epsilon)) - f(\gamma(t))}{\epsilon}, \quad f \in \mathfrak{F}(\mathcal{M}).$$

We say γ is a *regular curve* if $\dot{\gamma}(t) \neq 0$ for $t \in I$.

We have the following proposition on smooth curves:

Proposition 2.1.2. *Let $F : \mathcal{M} \rightarrow \mathcal{N}$ be smooth. Let $p \in \mathcal{M}$ and $\xi \in T_p\mathcal{M}$. Next, let $\gamma_i : (-\epsilon, \epsilon) \rightarrow \mathcal{M}$ be two curves with $\gamma_1(0) = \gamma_2(0) = p$ and $\dot{\gamma}_1(0) = \dot{\gamma}_2(0) = \xi$. Then*

$$(F \circ \gamma_1)'(0) = (F \circ \gamma_2)'(0) \in T_{F(p)}\mathcal{N}.$$

Intuitively, we are taking the function F at a point p and perturbing that point p along ξ . Consequently, the output of F is also perturbed and that perturbation remains invariant of any curve used to describe it, provided that the curve passes p with velocity ξ . We can now define the differential.

The *differential* of a smooth function $F : \mathcal{M} \rightarrow \mathcal{N}$ at $p \in \mathcal{M}$ is a linear mapping $dF_p : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{N}$ defined as

$$dF_p(\xi) := (F \circ \gamma)'(0) \quad (2.1)$$

where γ is any curve with $\gamma(0) = p$ and $\dot{\gamma}(0) = \xi$. With this in mind, dF_p can be intuited as the directional derivative of F at p along ξ .

The *global differential* of F is simply the assignment of the differential of F to each point:

$$dF : p \mapsto dF_p, \quad p \in \mathcal{M}.$$

Remark, the relationship between differentials and global differentials is analogous to the relationship between tangent vectors and vector fields.

Global differentials give us yet another way of interpreting tangent vectors. Let $(p, \xi) \in T\mathcal{M}$. We can write

$$\xi f = df_p(\xi), \quad \forall f \in \mathfrak{F}(\mathcal{M}).$$

Pushforwards and Pullbacks Let $F : \mathcal{M} \rightarrow \mathcal{N}$ is a diffeomorphism and $V \in \mathfrak{X}(\mathcal{M})$. The *pushforward* of V through F is a vector field on $\mathcal{N}$ defined as

$$W := F_*(V) \equiv F_*V \equiv dF(V), \quad W_p := dF_p(V_p)$$

When F is instead a diffeomorphism onto its image, we still write F_*V to imply a vector field on $F(\mathcal{M}) \subset \mathcal{N}$.

Similarly, since F is a diffeomorphism, the *pushback* of a vector field $W \in \mathfrak{X}(\mathcal{N})$ is a vector field $V \in \mathfrak{X}(\mathcal{M})$ defined as

$$V := F^*(W) \equiv F^*W, \quad V_p := (dF_p)^{-1}(W_{F(p)})$$

Final Remarks The differential $dF_p : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{N}$ is simply a linear operator between real finite-dimensional vector spaces. As a result, it carries the notions of rank and nullity, image and nullspace, injectivity and surjectivity and bijectivity, and eigenvalues, though not singular values.

When a smooth mapping $F : \mathcal{M} \rightarrow \mathcal{N}$ possesses an everywhere injective differential, we designate it an *immersion*. Immersions do not necessarily have to be injective. Furthermore, the existence of an immersion immediately implies $\dim(\mathcal{N}) \geq \dim(\mathcal{M})$. Dually, when F has an everywhere surjective differential, we label it a *submersion*. As before, submersions do not necessarily have to be surjective, and they immediately imply $\dim(\mathcal{N}) \leq \dim(\mathcal{M})$. Points for which $dF_p = 0$ are called *stationary*. Morse theory is the study of how the stationary points of a smooth scalar field $f \in \mathfrak{F}(\mathcal{M})$ give us information on global topological properties of $\mathcal{M}$ [14].

2.1.4 Parameterizations

Submanifolds Recall a smooth map $F : \mathcal{M} \rightarrow \mathcal{N}$ that is a diffeomorphism onto its image is called a smooth embedding. This is a smooth analog of a topological embedding, which is a mapping that is a homeomorphism onto its image. When such a map exists, we call $\mathcal{M}$ an *embedded submanifold* of $\mathcal{N}$. Remark $\mathcal{M} \cong F(\mathcal{M})$ inherits the same smooth manifold structure. We have the following propositions on smooth embeddings.

Proposition 2.1.3. *A mapping $F : \mathcal{M} \rightarrow \mathcal{N}$ is a smooth embedding if and only if it*

is an immersion and a topological embedding.

Proposition 2.1.4. *Let $\mathcal{N}$ be a smooth manifold and $\mathcal{M} \subset \mathcal{N}$ an arbitrary subset. We equip $\mathcal{M}$ with the subspace topology, making the inclusion mapping $\iota : \mathcal{M} \rightarrow \mathcal{N}$ a topological embedding. Then there exists at most one smooth structure on $\mathcal{M}$ for which ι is an immersion, making $\mathcal{M}$ an embedded submanifold of $\mathcal{N}$. We call this smooth structure the smooth submanifold structure. The existence of the smooth submanifold structure is not guaranteed; when it does exist, it is unique.*

Example 2.1.5. *Consider $\mathbb{R}^2$ and the 1-sphere $\mathcal{S}^1$. Then $\mathcal{S}^1$ admits the smooth embedded structure in the obvious sense.*

Example 2.1.6. *Consider $\mathbb{R}^2$ and the unit square $C := \{p = (x, y) \in \mathbb{R}^2 : \|p\| := \max(|x|, |y|) = 1\}$. Then C does not admit the smooth embedded structure.*

Embedded submanifolds of a vector space are called Euclidean submanifolds. The next theorem states all smooth manifolds are Euclidean submanifolds.

Theorem 2.1.4 (Whitney Embedding Theorem). *All smooth manifolds can be smoothly embedded in a vector space. In particular, a smooth manifold of dimension n can always be smoothly embedded in $\mathbb{R}^{2n}$.*

Smooth Covers Let $\mathcal{M}$ be a smooth manifold. A *smooth k -covering* of $\mathcal{M}$ is a pair $(\mathcal{E}, \pi)$, where $\mathcal{E}$ is a smooth manifold and $\pi : \mathcal{E} \rightarrow \mathcal{M}$ is a surjective smooth map with the following property. For every point $p \in \mathcal{M}$, there exists a neighborhood U with

$$\pi^{-1}(U) = \bigcup_{i=1}^k V_i,$$

where the $V_i \subset \mathcal{E}$ are pairwise disjoint and open. Furthermore, each restriction

$$\pi|_{V_i} : V_i \rightarrow U$$

is a diffeomorphism. As a side note, this implies $\dim(\mathcal{E}) = \dim(\mathcal{M})$.

When $\mathcal{E}$ is a Euclidean submanifold, we call $(\mathcal{E}, \pi)$ a *parameterization* of $\mathcal{M}$. The simplest type of parametrization is a 1-cover, implying $\mathcal{E} \cong \mathcal{M}$. For example, the unit quaternions form a 2-cover of $SO(3)$, whereas rotation matrices provide a 1-cover.

2.1.5 Coordinates and Frames: Computation

Let $\mathcal{V}$ be a real vector space of dimension n and let (e_i) be a basis. For any $v \in \mathcal{V}$, there exist unique scalars $v^i \in \mathbb{R}$ such that

$$v = v^i e_i.$$

The scalars v^i are called the *components* of v with respect to (e_i) . We write

$$[v] := (v^1, \dots, v^n) \in \mathbb{R}^n, \quad [v]_i := v^i.$$

A *covector* is a linear map $v^* : \mathcal{V} \rightarrow \mathbb{R}$. The set of all covectors forms the *dual space* $\mathcal{V}^*$. The basis (e_i) induces a unique dual basis $(\varepsilon^i) \subset \mathcal{V}^*$ characterized by

$$\varepsilon^i(e_j) = \delta_j^i.$$

For any $v = v^j e_j$, it follows that

$$\varepsilon^i(v) = v^j \varepsilon^i(e_j) = v^j \delta_j^i = v^i.$$

Thus, ε^i extracts the i th component of a vector with respect to the basis (e_i) .

Cotangent Spaces For a smooth manifold $\mathcal{M}$ and a point $p \in \mathcal{M}$, the *cotangent space* $T_p^* \mathcal{M}$ is defined as the dual space of $T_p \mathcal{M}$. The cotangent bundle $T^* \mathcal{M}$ and the space of covector fields $\mathfrak{X}^*(\mathcal{M})$ are defined analogously.

Example 2.1.7. Let $f \in \mathfrak{F}(\mathcal{M})$ be a smooth scalar field. Its global differential

$$df : p \mapsto df_p$$

is a covector field. So, $df \in \mathfrak{X}^*(\mathcal{M})$.

Frames and Coframes Let $U \subset \mathcal{M}$ be open. A *(local) frame* on U is a collection of n vector fields $(E_i) \subset \mathfrak{X}(\mathcal{M})$ such that $(E_i|_p)$ forms a basis of $T_p\mathcal{M}$ for each $p \in U$. If $U = \mathcal{M}$, the frame is called *global*. A *(local) coframe* is a collection $(E^i) \subset \mathfrak{X}^*(\mathcal{M})$ forming a basis of $T_p^*\mathcal{M}$ over U . The frame and coframe are dual if

$$E^i(E_j) = \delta_j^i.$$

Every frame admits a unique dual coframe, and vice versa.

When $\mathcal{M}$ admits a global frame, we say it is *parallelizable*. In general, parallelizable manifolds are rare.

Proposition 2.1.5. *A smooth manifold $\mathcal{M}$ of dimension n is parallelizable if and only if $T\mathcal{M} \cong \mathcal{M} \times \mathbb{R}^n$.*

Coordinate Frames Let (U, x) be a chart on $\mathcal{M}$. Each coordinate function $x^i : U \rightarrow \mathbb{R}$ is a smooth scalar field and its differential dx^i is a covector field. The collection (dx^i) forms a local coframe on U called the *coordinate coframe*. Its dual, called the *coordinate frame*, is denoted

$$\left(\frac{\partial}{\partial x^i} \right) \equiv (\partial_i).$$

The coordinate frame is defined pointwise as

$$\left. \frac{\partial}{\partial x^i} \right|_p = (dx_p)^{-1}(e_i),$$

where (e_i) is the standard basis of $\mathbb{R}^n$. For any $f \in \mathfrak{F}(\mathcal{M})$, we will write

$$\frac{\partial}{\partial x^i}[f] \equiv \frac{\partial}{\partial x^i}f \equiv \frac{\partial f}{\partial x^i} \equiv \partial_i f.$$

Components of Vector Fields Let $V \in \mathfrak{X}(\mathcal{M})$. On U , there exist unique smooth functions $V^i \in \mathfrak{F}(U)$ such that

$$V = V^i \frac{\partial}{\partial x^i}.$$

The tuple $[V] = (V^1, \dots, V^n)$ gives the components of V in the coordinate frame, where

$$V^i = dx^i(V) = V[x^i].$$

We also write $[V]_i := V^i$.

Components of Differentials Let $F : \mathcal{M} \rightarrow \mathcal{N}$ be smooth, and let (U, x) and (V, y) be charts on $\mathcal{M}$ and $\mathcal{N}$, respectively. We will assume $\mathcal{M}$ has dimension m and $\mathcal{N}$ dimension n . Let $\xi = \xi^i \frac{\partial}{\partial x^i} |_p \in T_p \mathcal{M}$. Then $\eta := dF_p(\xi) = \eta^i \frac{\partial}{\partial y^i} |_{F(p)}$. The components of $dF_p : T_p \mathcal{M} \rightarrow T_{F(p)} \mathcal{N}$ is a matrix $[dF_p] \in \mathbb{R}^{n \times m}$ that describes the linear relationship between the components (ξ^i) of ξ and the components (η^i) of $\eta = dF_p(\xi)$. The components of dF_p become

$$[dF_p]_{ij} := d(y^j)_{F(p)} \left(dF_p \left(\frac{\partial}{\partial x^i} \Big|_p \right) \right).$$

It follows

$$[dF_p]_{ij} \xi^j = \eta^i, \quad [dF_p][\xi] = [\eta].$$

2.2 What is a Riemannian Manifold?

The smooth structure of a smooth manifold $\mathcal{M}$ can only answer questions about local geometric and topological constructions. Questions relating to global geometric and topological properties will require machinery from differential and algebraic topology which pertain to global behaviors of smooth manifolds.

On the other hand, a Riemannian metric enables more global geometric constructions, such as lengths of curves, angles, distances between points, zero-acceleration curves, and even convexity. Moreover, a Riemannian metric can often provide ways of answering global topological questions with geometric tools. This is favorable as the geometric tools often come with numerical approximation methods. For example, suppose we wish to determine whether $\mathcal{M}$ is compact (a global topological property). If we design a geodesically complete Riemannian metric, then compactness is equivalent to whether $\mathcal{M}$ is bounded under that metric. We transformed a global topological question into a geometric one.

2.2.1 Rank-2 Tensor Fields

Let b be a smooth assignment of a bilinear form

$$b_p : T_p\mathcal{M} \times T_p\mathcal{M} \rightarrow \mathbb{R}$$

to each point p . Bilinear means that $b_p(\cdot, \cdot)$ is linear in each argument, and by smoothness we mean that $p \mapsto b_p(V_p, W_p)$ is a smooth scalar field for each $V, W \in \mathfrak{X}(\mathcal{M})$. We call b a $(0, 2)$ -tensor field. From here we can construct notions of symmetric, positive-definite, and skew-symmetric $(0, 2)$ -tensor fields.

Globally, we can also think of b as a mapping

$$b : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{F}(\mathcal{M}),$$

where $b(V, W) : p \mapsto B_p(V_p, W_p)$.

Let (E_i) be a frame. The components of b form a matrix $[b_p] \in \mathbb{R}^{n \times n}$ given by

$$[b_p]_{ij} := b_p(E_i|_p, E_j|_p).$$

Therefore $b(V, W) = [V]^\top [b][W]$.

Next, let h be a smooth assignment of a linear operator $h_p : T_p\mathcal{M} \rightarrow T_p\mathcal{M}$ to each point p . By smooth, we mean that for each $V \in \mathfrak{X}(\mathcal{M})$, the mapping $p \mapsto h_p(V_p)$ is a vector field. We call h a $(1, 1)$ -tensor field.

Globally, we can think of h as a mapping

$$h : \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}),$$

where $h(V) : p \mapsto h_p(V_p)$.

Let (E_i) be a frame with dual coframe (E^i) . The components of h form a matrix $[h_p] \in \mathbb{R}^{n \times n}$ given by

$$[h_p]_{ij} = E^i|_p(h_p(E_j|_p)).$$

Therefore, $[h(V)] = [h][V]$.

2.2.2 Riemannian Metric

A *Riemannian metric* g on $\mathcal{M}$ is a positive-definite symmetric $(0, 2)$ -tensor field. In other words, g assigns an inner product

$$g_p : T_p\mathcal{M} \times T_p\mathcal{M} \rightarrow \mathbb{R}$$

to each point p . The pairing $(\mathcal{M}, g)$ is called a *Riemannian manifold*.

With respect to a local frame (E_i) , the metric has components

$$g_{ij} \equiv [g]_{ij} := g(E_i, E_j).$$

Since g_p is positive-definite, its components will be a positive-definite matrix $[g_p] \in \text{Sym}_{++}(n)$. If the frame (E_i) is orthonormal at p , then $[g_p] = I_n$. A chart (U, x) is called *normal* at a point p if the coordinate frame $(\partial/\partial x^i)$ is orthonormal.

Mappings with Metrical Properties Let $(\mathcal{M}, g^{\mathcal{M}})$ and $(\mathcal{N}, g^{\mathcal{N}})$ be Riemannian manifolds and let $F : \mathcal{M} \rightarrow \mathcal{N}$ be smooth. The mapping F is called an *isometry* if it is a diffeomorphism and satisfies

$$g^{\mathcal{N}}(dF_p(\xi), dF_p(\eta)) = g^{\mathcal{M}}(\xi, \eta) \quad \forall p \in \mathcal{M}, \xi, \eta \in T_p\mathcal{M}.$$

The mapping F is called a *local isometry* if for every $p \in \mathcal{M}$ there exists a neighborhood U of p such that $F|_U : (U, g^{\mathcal{M}}) \rightarrow (F(U), g^{\mathcal{N}})$ is an isometry.

The mapping F is called a *Riemannian immersion* if it is an immersion and a local isometry onto its image. That is, dF_p is injective and satisfies

$$g^{\mathcal{N}}(dF_p(\xi), dF_p(\eta)) = g^{\mathcal{M}}(\xi, \eta) \quad \forall \xi, \eta \in T_p\mathcal{M}.$$

Dually, F is a *Riemannian submersion* if it is a smooth submersion and dF_p restricts to an isometry between the *horizontal space* $\mathcal{H}_p = (\ker dF_p)^\perp \subset T_p\mathcal{M}$ and $T_{F(p)}\mathcal{N}$ for all $p \in \mathcal{M}$.

2.2.3 Levi-Civita Connection

A Riemannian metric induces a canonical notion of directional differentiation of vector fields, called the *Levi-Civita connection*.

Theorem 2.2.1 (Fundamental Theorem of Riemannian Geometry). *Let $(\mathcal{M}, g)$ be a Riemannian manifold. There exists a unique mapping*

$$\nabla : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}), \quad (U, V) \mapsto \nabla_U V,$$

such that for all $U, V, W \in \mathfrak{X}(\mathcal{M})$ and all $f \in \mathfrak{F}(\mathcal{M})$,

∇ is $\mathbb{R}$ -bilinear,

$$\nabla_U(fV) = (Uf)V + f\nabla_U V,$$

$$Ug(V, W) = g(\nabla_U V, W) + g(V, \nabla_U W),$$

$$\nabla_U V - \nabla_V U = [U, V].$$

The unique connection satisfying these properties is called the Levi–Civita connection of $(\mathcal{M}, g)$.

Theorem 2.2.2. *Let $V, W \in \mathfrak{F}(\mathcal{M})$. Let $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M}$ be a curve with $\gamma(0) = p$ and $\dot{\gamma}(0) = V_p$. Then $\nabla_V W|_p$ depends only on V_p , and W along the curve $\gamma|_{[0, \epsilon)}$.*

Connection Coefficients Fix a chart (U, x) . The components of the connection are called the Christoffel symbols Γ_{ij}^k . They are computed as

$$\nabla_{\partial_i} \partial_j = \Gamma_{ij}^k \partial_k.$$

In these coordinates, the Christoffel symbols of the Levi-Civita connection are

$$\Gamma_{ij}^k = \frac{1}{2} g^{kl} (\partial_i g_{jl} + \partial_j g_{il} - \partial_l g_{ij}),$$

where $g^{ij} = ([g]^{-1})_{ij}$. For vector fields $V = V^i \partial_i$ and $W = W^j \partial_j$, we have

$$\nabla_V W = \left(V^i \partial_i [W^k] + \Gamma_{ij}^k V^i W^j \right) \partial_k.$$

Covariant Derivative Along A Curve Let $\gamma : I \rightarrow \mathcal{M}$ be a smooth curve and let $V \in \mathfrak{X}(\text{im}(\gamma))$ be a vector field along γ . We can define

$$V(t) := V_{\gamma(t)},$$

where $V(\cdot)$ is to be understood as a curve in $T\mathcal{M}$. We identify $V(\cdot)$ with V :

$$V(t) \equiv V_{\gamma(t)}.$$

The *covariant derivative of V along γ* is a curve $\frac{D}{dt}V : I \rightarrow T\mathcal{M}$ defined as

$$\frac{D}{dt}V := \nabla_{\dot{\gamma}} V \in T_{\gamma}\mathcal{M}.$$

In coordinates, if $\dot{\gamma} = \dot{\gamma}^i \partial_i|_{\gamma}$ and $V = V^j \partial_j$, then

$$\frac{D}{dt}V(t) = \left(\dot{V}^k(t) + \Gamma_{ij}^k(\gamma(t)) \dot{\gamma}^i(t) V^j(t) \right) \partial_k \Big|_{\gamma(t)}.$$

We say that V is *parallel along* γ if $\frac{D}{dt}V(t) \equiv 0$, equivalently,

$$\dot{V}^k(t) = -\Gamma_{ij}^k(\gamma(t)) \dot{\gamma}^i(t) V^j(t).$$

We have the following theorem which follows from the fact that parallel vector fields follow an LTV system:

Theorem 2.2.3. *Let $(\mathcal{M}, g)$ be a Riemannian manifold and $\gamma : [a, b] \rightarrow \mathcal{M}$ a smooth curve. Let $\xi \in T_a\mathcal{M}$. Then there exists a unique parallel vector field $V \in \mathfrak{X}(\text{im}(\gamma))$ along γ .*

Parallel Transport Given the previous theorem, we can now transport tangent vectors along curves in a unique and meaningful way. Let $\gamma : [a, b] \rightarrow \mathcal{M}$ be a smooth curve and $\xi \in T_{\gamma(a)}\mathcal{M}$. Define

$$P_\gamma : T_{\gamma(a)}\mathcal{M} \rightarrow T_{\gamma(b)}\mathcal{M}, \quad P_\gamma(\xi) := V(b),$$

where $V \in \mathfrak{X}(\text{im}(\gamma))$ is the unique parallel vector field along γ . This operator is linear.

The following lemma shows that the Levi-Civita connection can be obtained uniquely from parallel transport.

Lemma 2.2.1. *Let $V \in \mathfrak{X}(\mathcal{M})$, $\xi \in T_p\mathcal{M}$, and let $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{M}$ satisfy $\gamma(0) = p$ and $\dot{\gamma}(0) = \xi$. Then*

$$\nabla_\xi V = \lim_{t \rightarrow 0} \frac{P_{t \rightarrow 0}^\gamma V_{\gamma(t)} - V_p}{t} = \frac{D}{dt} \left(P_{t \rightarrow 0}^\gamma V_{\gamma(t)} \right) \Big|_{t=0}.$$

2.2.4 Riemannian Gradient And Hessian

Gradient Let $f \in \mathfrak{F}(\mathcal{M})$. The *Riemannian gradient* of f is the unique vector field $\nabla f \in \mathfrak{X}(\mathcal{M})$ satisfying

$$g(V, \nabla f) = V f \quad \forall V \in \mathfrak{X}(\mathcal{M}). \quad (2.2)$$

In local coordinates (x^i) , the gradient has the components

$$\nabla f = [\nabla f]_i \partial_i, \quad [\nabla f]_i = [g]^{ij} \partial_j f.$$

Operator Form Of The Hessian The *Riemannian Hessian* of f is a $(1, 1)$ -tensor field

$$\nabla^2 f : \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}), \quad \nabla^2 f(V) := \nabla_V(\nabla f).$$

For $V \in \mathfrak{X}(\mathcal{M})$, the components of $\nabla^2 f(V)$ become

$$[\nabla^2 f(V)]_k = V^i \left(\partial_i [(\nabla f)^k] + \Gamma_{ij}^k (\nabla f)^j \right).$$

Equivalently, $\nabla^2 f_p$ admits a matrix representation $[\nabla^2 f]_{ki}$ defined by

$$[\nabla^2 f(V)]_k = [\nabla^2 f]_{ki} V^i,$$

where

$$[\nabla^2 f]_{ki} = \partial_i [(\nabla f)^k] + \Gamma_{ij}^k (\nabla f)^j$$

is the coordinate representation of the Hessian as a linear map. Here, it follows

$$[\nabla^2 f][V] = [W]$$

where $V, W \in \mathfrak{X}(\mathcal{M})$ and $W_p := \nabla^2 f_p(V_p)$.

Bilinear Form of the Hessian Using the Riemannian metric, the Hessian can equivalently be viewed as a $(0, 2)$ -tensor (a symmetric bilinear form)

$$\text{Hess} f : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{F}(\mathcal{M}), \quad \text{Hess} f(V, W) := g(\nabla_V(\nabla f), W).$$

Remark we use Hess and ∇^2 to differentiate between the two concepts.

The two definitions are related by

$$\text{Hess} f(V, W) = g(\nabla^2 f(V), W).$$

In coordinates, the components of the bilinear Hessian are

$$\text{Hess} f_{ij} := \text{Hess} f(\partial_i, \partial_j) = g_{jk} [\nabla^2 f]_{ki}.$$

Substituting the expression above yields

$$\text{Hess}f_{ij} = \partial_i \partial_j f - \Gamma_{ij}^k \partial_k f \quad (2.3)$$

which is symmetric in i, j . We have $[\text{Hess}f] = [g][\nabla^2 f]$. This implies for any $V, W \in \mathfrak{X}(\mathcal{M})$,

$$\text{Hess}f(V, W) = [W]^\top [\text{Hess}f][V] = [W]^\top [g][\nabla^2 f][V].$$

2.2.5 Geodesics

Let $\gamma : I \rightarrow \mathcal{M}$ be a regular curve. Its velocity $\dot{\gamma}(t)$ is a tangent vector in $T_{\gamma(t)}\mathcal{M}$, and can be viewed as a vector field along γ . The *intrinsic acceleration* of γ is defined by

$$\frac{D}{dt}\dot{\gamma}(t) := \nabla_{\dot{\gamma}(t)}\dot{\gamma}(t).$$

In coordinates, writing $\dot{\gamma}(t) = \dot{\gamma}^i(t)\partial_i$, we have

$$\nabla_{\dot{\gamma}}\dot{\gamma} = \left(\ddot{\gamma}^k + \Gamma_{ij}^k(\gamma(t))\dot{\gamma}^i\dot{\gamma}^j \right) \partial_k.$$

A curve γ is a *geodesic* if its intrinsic acceleration vanishes,

$$\nabla_{\dot{\gamma}}\dot{\gamma} = 0,$$

equivalently, if it satisfies the *geodesic equation*

$$\ddot{\gamma}^k(t) = -\Gamma_{ij}^k(\gamma(t))\dot{\gamma}^i(t)\dot{\gamma}^j(t).$$

This is a second-order ODE, hence locally admits a unique solution given initial conditions $\gamma(0) = p$ and $\dot{\gamma}(0) = v$.

The manifold $(\mathcal{M}, g)$ is called *geodesically complete* if every geodesic extends to a solution defined for all $t \in \mathbb{R}$.

Geodesics As Critical Curves Of Length Let $\gamma : [a, b] \rightarrow \mathcal{M}$ be a piecewise smooth curve. The *length* of γ is defined as

$$\Lambda(\gamma) := \int_a^b \sqrt{g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))} dt.$$

Remark the length of γ is invariant of its velocity profile, as one would expect. Let's fix $p, q \in \mathcal{M}$. We are interested in curves that connect p to q which serve as stationary points of Λ . Due to the reparameterization-invariance of the length operator Λ , we will define the family $\Gamma(p, q)$ of piecewise smooth curves $\gamma : [a, b] \rightarrow \mathcal{M}$ with $\gamma(a) = p$, $\gamma(b) = q$, and $\frac{d}{dt}\|\dot{\gamma}(t)\|_{\gamma(t)} = 0$. We have the following theorem:

Theorem 2.2.4. *A curve in $\Gamma(p, q)$ is a geodesic if and only if it is a stationary point of Λ are geodesic curves.*

This implies that geodesics may also be characterized as curves which locally extremize length. They may locally minimize length or may act as saddle points; the length functional admits no local maxima. For example, on $\mathcal{S}^2$, both the short and long great-circle arcs between nearby points are geodesics, but only the short arc is minimizing.

Metrical Structure The *geodesic distance* between $p, q \in \mathcal{M}$ is defined as

$$d_g(p, q) := \inf_{\gamma \in \Gamma(p, q)} \Lambda(\gamma).$$

The function $d_g(\cdot, \cdot)$ defines a metric and therefore induces a topology τ_g on $\mathcal{M}$.

Proposition 2.2.1. *Let $(\mathcal{M}, g)$ be a connected Riemannian manifold with underlying topology τ . Let τ_g denotes the metric topology induced by $d_g(\cdot, \cdot)$. Then*

$$\tau_g = \tau.$$

This equivalence shows that a smooth manifold can be studied as a metric space once a Riemannian metric is chosen, allowing topological properties such as compactness and connectedness to be approached through geometric tools.

We end this section with the celebrated Hopf–Rinow theorem, which relates geodesics to metric completeness of $d_g(\cdot, \cdot)$:

Theorem 2.2.5 (Hopf–Rinow). *Let $(\mathcal{M}, g)$ be a connected Riemannian manifold and $d_g(\cdot, \cdot)$ the geodesic distance. Then $(\mathcal{M}, d_g)$ is a complete metric space if and only if all geodesic curves $\gamma : [a, b] \rightarrow \mathcal{M}$ can be extended over all $\mathbb{R}$.*

Geodesic Convexity What is convexity? Convexity is a notion that allows one can use local properties of a set or function to conclude global properties of a set or function. A subset $C \subset \mathbb{R}^n$ is convex if and only if for any pair $p, q \in C$, the line connecting them is contained in C . It follows to ask what is a line in Riemannian manifolds. Geodesics are precisely that, as they are the zero-acceleration curves which extremize length.

However, we only want to focus on geodesics which globally minimize length. We say a geodesics $\gamma \in \Gamma(p, q)$ is minimizing if γ is a global minimum of Λ over $\Gamma(p, q)$. We will denote $\Gamma^{\min}(p, q)$ as the family of minimizing geodesics connecting p to q . Without loss of generality, we restrict to curves defined over $\gamma : [a, b] \rightarrow \mathcal{M}$.

A subset $A \subset \mathcal{M}$ is called *geodesically convex* if for any $p, q \in A$, there exists a unique minimizing geodesic γ connecting p to q (that is, $\Gamma^{\min}(p, q)$ is a singleton set), and that geodesic is contained in A . Some authors refer to this as strong geodesic convexity.

Let $A \subset \mathcal{M}$ be geodesically convex. A function $f \in \mathfrak{F}(A)$ is called (strongly) geodesically convex if for all $p, q \in \mathcal{M}$,

$$f \circ \gamma : [a, b] \rightarrow \mathbb{R}$$

is (strongly) convex in the classical sense.

Lemma 2.2.2 (First-order characterization of strong geodesic convexity). *Let $A \subset \mathcal{M}$ be geodesically convex and $f \in \mathfrak{F}(A)$. Then f is geodesically convex if and only if*

for all $p, q \in A$,

$$f(q) \geq f(p) + \langle \text{grad} f_p, \log_p(q) \rangle,$$

Lemma 2.2.3 (Second-order characterization via the Hessian). *Let $A \subset \mathcal{M}$ be geodesically convex and $f \in \mathfrak{F}(A)$. Then f is strongly geodesically convex if and only if $\text{Hess} f_p$ is positive definite at all $p \in A$.*

2.2.6 Retractions: Navigating Across the Manifold

Let $\mathcal{M}$ be a smooth manifold. A retraction is a smooth mechanism for moving from a point p along a tangent direction in an intrinsic and principled way. Retractions are fundamental in optimization on manifolds, where updates repeatedly step from the current iterate along a chosen tangent direction [4, 7, 15].

A *retraction* is a smooth map $R : \mathcal{S} \subset T\mathcal{M} \rightarrow \mathcal{M}$ defined on an open subset $\mathcal{S}$ containing $(p, 0) \in T_p\mathcal{M}$ for every $p \in \mathcal{M}$, such that for all $(p, \xi) \in \mathcal{S}$ the curve $\gamma(t) := R(p, t\xi) \equiv R_p(t\xi)$ satisfies

$$\gamma(0) = p, \quad \dot{\gamma}(0) = \xi,$$

The second condition is equivalent to requiring $\text{d}R_p|_0 : T_0(T_p\mathcal{M}) \equiv T_p\mathcal{M} \rightarrow T_p\mathcal{M}$ is the identity mapping.

By the inverse function theorem, for each p there exists a neighborhood $U \subset T_p\mathcal{M}$ of 0_p such that R_p restricts to a diffeomorphism $U \rightarrow R_p(U)$. The resulting local inverse is written as

$$R_p^{-1} : R_p(U) \subset \mathcal{M} \rightarrow U \subset T_p\mathcal{M}.$$

Remark that $\text{d}R_p^{-1}|_p : T_p\mathcal{M} \rightarrow T_0(T_p\mathcal{M}) \equiv T_p\mathcal{M}$ must also be the identity mapping.

2.2.7 Riemannian Gradient Descent

Given a smooth function $f \in \mathfrak{F}(\mathcal{M})$ and a choice of retraction R , the Riemannian Gradient Descent iteration is

$$p_{t+1} := R_{p_t}(-s_t \nabla f_{p_t}), \quad (2.4)$$

where $s_t \geq 0$ is a step size. Figure 2.1 provides a visual depiction of the RGD update. We recommend [4, 7] for optimization-oriented references on Riemannian manifolds; see also [5, 6].




Figure 2.1: Visualization of RGD; here, $x_2 = R_{x_1}(-s_1 \nabla f(x_1))$.

Chapter 3

THE SPACE OF CONTROLLERS IS A RIEMANNIAN MANIFOLD AND HOW TO OPTIMIZE OVER IT

There is an intimate relationship between global topological properties of spaces of linear systems and basic control-theoretic questions in the theory of linear systems.

Uwe Helmke [16]

Direct policy optimization (DPO) is an emerging paradigm in control theory that synthesizes feedback controllers by directly optimizing controller parameters with gradient-based methods, rather than relying exclusively on algebraic synthesis pipelines based on Lyapunov arguments or linear matrix inequalities (LMIs) [17]. This viewpoint is attractive in settings where one works directly with data and structural constraints on the controller class.

Motivated by reinforcement learning (RL) [18–20], DPO aims to solve control design problems using policy gradient methods and an RL-friendly framework. DPO is different from RL because of the focus on obtaining rigorous control-theoretic guarantees. This includes convergence guarantees, stability guarantees, and performance guarantees [21].

In the linear–quadratic–Gaussian (LQG) setting, systematic study of these ideas began with the LQG policy optimization framework [22] as well as LQR policy optimization [21, 23, 24]. A central difficulty in LQG policy optimization is that the objective depends only on the closed-loop input–output behavior and is therefore invariant under changes of the controller’s internal state coordinates. For example, if $\mathbf{K} = (A_c, B_c, C_c)$ is a state-space realization of a dynamic controller of order q , then for any $S \in \text{GL}(q)$ the coordinate change $x'_c = Sx_c$ induces the transforma-

tion $\mathcal{T}_S(\mathbf{K}) = (SA_cS^{-1}, SB_c, C_cS^{-1})$. The transformation preserves the closed-loop transfer map and hence leaves the LQG cost unchanged. As a consequence, stationary points in realization space are degenerate; each stationary point corresponds to an entire $\text{GL}(q)$ -orbit. As a result, the Hessian is singular along the orbit directions. This is the reason behind the ill-conditioning and sublinear behavior observed for classical gradient descent [22, 25]. This degeneracy motivates optimization procedures that explicitly avoid the redundant coordinate directions.

We show that the set of stabilizing *minimal* dynamic linear controllers admits a smooth orbit manifold structure under the $\text{GL}(q)$ action. Establishing this requires combining classical linear systems theory with smooth orbit geometry [16, 26–51].

With this manifold structure in hand, we equip the controller orbit manifold with a coordinate-invariant Riemannian metric and carrying out optimization intrinsically, both in a first-order and second-order way. This removes the q^2 coordinate degrees of freedom from the search space, reducing the effective dimension from $q^2 + qp + qm$ to $qp + qm$ [52]. Finally, we present theoretical convergence guarantees and guarantees of a linear and quadratic rate for the first- and second-order optimization methods, respectively.

This geometric viewpoint can extend beyond LQG to other coordinate-invariant control synthesis problems [2, 3, 23, 25, 53–57].

3.1 LQG Policy Optimization

In this section, we present the continuous-time formulation of direct policy optimization (DPO) for linear-quadratic-Gaussian (LQG) control using dynamic output-feedback controllers.

Plant Model Consider a continuous-time LTI system

$$\dot{x}(t) = Ax(t) + Bu(t) + w(t), \quad y(t) = Cx(t) + v(t), \quad (3.1)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $C \in \mathbb{R}^{p \times n}$. The process noise $w(\cdot)$ and measurement noise $v(\cdot)$ are zero-mean Gaussian processes with covariance matrices $W \in \mathbb{S}_+^n$ and $V \in \mathbb{S}_+^p$, respectively. We assume that (A, B) and $(A, W^{1/2})$ are controllable and that (A, C) is observable.

Controller Model We stabilize (3.1) using a continuous-time LTI controller

$$\dot{x}_c(t) = A_c x_c(t) + B_c y(t), \quad u(t) = C_c x_c(t), \quad (3.2)$$

where $A_c \in \mathbb{R}^{q \times q}$, $B_c \in \mathbb{R}^{q \times p}$, and $C_c \in \mathbb{R}^{m \times q}$. We assume that (A_c, B_c) is controllable and (A_c, C_c) is observable. We represent a controller realization of order q by

$$\mathbf{K} := (A_c, B_c, C_c) \in \mathcal{T}_q, \quad (3.3a)$$

where

$$\mathcal{T}_q := \mathbb{R}^{q \times q} \times \mathbb{R}^{q \times p} \times \mathbb{R}^{m \times q}.$$

In this context, $\mathbf{P} := (A, B, C) \in \mathcal{T}$.

Closed-Loop Realization The combined plant-controller dynamics are given by

$$\begin{bmatrix} \dot{x} \\ \dot{x}_c \end{bmatrix} = \begin{bmatrix} A & BC_c \\ B_c C & A_c \end{bmatrix} \begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} I_n & 0_{n \times p} \\ 0_{q \times n} & B_c \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix}, \quad (3.4a)$$

$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} C & 0_{p \times q} \\ 0_{m \times n} & C_c \end{bmatrix} \begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} 0_{p \times n} & I_p \\ 0_{m \times n} & 0_{m \times p} \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix}. \quad (3.4b)$$

The inputs to the closed-loop system are (w, v) and the outputs are (y, u) .

We denote the resulting closed-loop realization matrices by

$$A_{\text{cl}}(\mathbf{K}) \in \mathbb{R}^{(n+q) \times (n+q)}, \quad B_{\text{cl}}(\mathbf{K}) \in \mathbb{R}^{(n+q) \times (n+p)}, \quad (3.5a)$$

$$C_{\text{cl}}(\mathbf{K}) \in \mathbb{R}^{(m+p) \times (n+q)}, \quad D_{\text{cl}}(\mathbf{K}) \in \mathbb{R}^{(m+p) \times (n+p)}. \quad (3.5b)$$

The controller (3.2) is stabilizing if and only if $A_{\text{cl}}(\mathbf{K}) \in \mathcal{H}_{n+q}$, where $\mathcal{H}_k$.

Controller Space Let $\tilde{\mathcal{C}}_q$ denote the set of all order- q controllers that are minimal (controllable and observable) and stabilize the plant (3.1). When the context is clear, we write $\tilde{\mathcal{C}} \equiv \tilde{\mathcal{C}}_q$ and $\mathcal{T} \equiv \mathcal{T}_q$.




Figure 3.1: Example of $\tilde{\mathcal{C}}_1$ for the plant $(A, B, C) = (1.1, 1, 1)$.

Coordinate Transformations And Orbits For any $S \in \text{GL}(q)$, define the coordinate transformation $\mathcal{T}_S : \tilde{\mathcal{C}}_q \rightarrow \tilde{\mathcal{C}}_q$ by

$$\mathcal{T}_S(\mathbf{K}) = (SA_cS^{-1}, SB_c, C_cS^{-1}). \quad (3.6)$$

A property is said to be *coordinate-invariant* if it is invariant under $\mathcal{T}_S$ for all $S \in \text{GL}(q)$.

Let $\tilde{J} : \tilde{\mathcal{C}} \rightarrow \mathbb{R}$ be a coordinate-invariant cost function. A canonical example is the LQG objective

$$\tilde{J}(\mathbf{K}) := \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T (x(t)^\top Q x(t) + u(t)^\top R u(t)) dt \right], \quad (3.7)$$

where $Q \in \text{Sym}_{++}(n)$ and $R \in \text{Sym}_{++}(m)$. The resulting optimization problem is

$$\min_{\mathbf{K} \in \tilde{\mathcal{C}}} \tilde{J}(\mathbf{K}). \quad (3.8)$$

Any memoryless iterative method for solving (3.8) has the form

$$\mathbf{K}_{t+1} = \mathbf{K}_t + \mathbf{V}(t, \mathbf{K}_t).$$

Since $\tilde{J}$ is coordinate-invariant (and analytic in our setting), it is constant along the orbits

$$\pi(\mathbf{K}) := \{ \mathcal{T}_S(\mathbf{K}) : S \in \text{GL}(q) \}. \quad (3.9)$$

Consequently, meaningful descent directions must be orthogonal to the tangent space of the orbit at $\mathbf{K}$, i.e.,

$$\mathbf{V}_t \perp T_{\mathbf{K}_t} \pi(\mathbf{K}_t).$$




Figure 3.2: Illustration of $\tilde{\mathcal{C}}$ and its orbits. Descent directions must be orthogonal to the orbit through the current iterate.

3.2 Geometry of Linear Systems

In this section, we apply smooth orbit manifold theory to spaces of linear systems and linear controllers. See §C.1 for a primer on smooth orbit geometry.

Spaces of controllers arise as coordinate-invariant subsets of linear system spaces, so structural results established for linear systems transfer directly to controller spaces.

3.2.1 Spaces of Linear Systems

Let

$$\tilde{\Sigma}_{m,p}^{n,*} := \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times m} \times \mathbb{R}^{p \times n}$$

denote the vector space of state–space realizations (A, B, C) of continuous-time LTI systems with m inputs, n states, and p outputs. We refer to $\tilde{\Sigma}^* \equiv \tilde{\Sigma}_{m,p}^{n,*}$ as the *total linear system space*.

Let $\tilde{\Sigma} \subset \tilde{\Sigma}^*$ denote the subset of minimal realizations. We call $\tilde{\Sigma}$ the *linear system space*.

Coordinate Transformations Coordinate transformations act on $\tilde{\Sigma}^*$ via

$$\mathcal{T}_S(A, B, C) := (SAS^{-1}, SB, CS^{-1}), \quad S \in \text{GL}(n).$$

This defines a smooth and free group action of $\text{GL}(n)$.

The corresponding orbit spaces are

$$\begin{aligned} \Sigma^* &:= \tilde{\Sigma}^* / \text{GL}(n), \\ \Sigma &:= \tilde{\Sigma} / \text{GL}(n), \end{aligned}$$

which we call the *total linear system orbit space* and the *linear system orbit space*, respectively.

Proposition 3.2.1. (*[35]*) *The space Σ^* is not Hausdorff.*

Proof. Let $P = (A, B, C) \in \tilde{\Sigma}^*$ with $A, B, C \neq 0$, and define $P_t := (tA, tB, tC)$. For every $t > 0$,

$$P_t = \mathcal{T}_{tI}(P),$$

so P_t lies in the same orbit as P . The sequence $P_{1/k}$ converges in $\tilde{\Sigma}^*$ to $P_0 = (0, 0, 0)$. By continuity of the orbit map π , the sequence $\pi(P_{1/k})$ converges to both $\pi(P)$ and $\pi(P_0)$, which are distinct. Hence Σ^* is not Hausdorff. $\square$

Smooth Orbit Submanifolds Any $\mathrm{GL}(n)$ -invariant embedded submanifold $\tilde{\Sigma}' \subset \tilde{\Sigma}^*$ on which the group action is free and proper admits a smooth orbit manifold structure.

Theorem 3.2.1. *Let $\tilde{\Sigma}^c \subset \tilde{\Sigma}$ and $\tilde{\Sigma}^o \subset \tilde{\Sigma}$ denote the sets of controllable and observable systems, respectively. Then*

$$\Sigma^c := \tilde{\Sigma}^c / \mathrm{GL}(n), \quad \Sigma^o := \tilde{\Sigma}^o / \mathrm{GL}(n)$$

are smooth orbit manifolds. Moreover, the orbit space of any $\mathrm{GL}(n)$ -invariant embedded submanifold of these spaces is a smooth orbit manifold, including controller spaces defined by fixing a plant and restricting to stabilizing realizations.

Proof. Let $\mathbf{C}(A, B)$ and $\mathbf{O}(A, C)$ denote the controllability and observability matrices. The sets

$$\tilde{\Sigma}^c = \{\det(\mathbf{C}\mathbf{C}^\top) \neq 0\}, \quad \tilde{\Sigma}^o = \{\det(\mathbf{O}^\top \mathbf{O}) \neq 0\}$$

are open subsets of $\tilde{\Sigma}^*$ and therefore smooth manifolds. Since the $\mathrm{GL}(n)$ -action is smooth, free, and proper [58], their orbit spaces are smooth orbit manifolds. $\square$

Laurent Isomorphism We conclude with a classical result relating three equivalent realizations of linear systems.

Let $\mathcal{H}_{m,p}^n$ denote the set of $p \times m$ block Hankel matrices of order n , and let $\mathrm{Rat}_{m,p}^n$ denote the set of $p \times m$ transfer matrices of McMillan degree n .

Theorem 3.2.2. (*[29, Col. A.2]*) *The following hold:*

- $\mathcal{H}_{m,p}^n$ is a smooth quasi-affine variety and an embedded submanifold of $\mathbb{R}^{(n+1)p \times (n+1)m}$.
- The orbit space $\tilde{\Sigma}_{m,p}^n / \text{GL}(n)$ is both an orbit variety and a smooth orbit manifold.
- $\text{Rat}_{m,p}^n$ is a quasi-projective variety.
- These three spaces are related by bi-regular mappings. In particular, $\tilde{\Sigma}_{m,p}^n / \text{GL}(n)$ is diffeomorphic to $\mathcal{H}_{m,p}^n$, and both induce the same smooth structure on $\text{Rat}_{m,p}^n$.

3.2.2 Spaces of Linear Controllers

Fix a plant $P \in \tilde{\Sigma}^*$. Let $\tilde{\mathcal{C}}_q^*$ denote the space of all stabilizing linear controllers of order q for P , and let $\tilde{\mathcal{C}}_q \subset \tilde{\mathcal{C}}_q^*$ denote the subset of minimal controllers. We refer to these as the *total controller space* and the *controller space*, respectively. Coordinate transformations of controller realizations act by $\text{GL}(q)$, and the associated orbit spaces are

$$\mathcal{C}^* := \tilde{\mathcal{C}}^* / \text{GL}(q), \quad \mathcal{C} := \tilde{\mathcal{C}} / \text{GL}(q).$$

which we will call the *total controller orbit space* and *controller orbit space*, respectively.

Lemma 3.2.1. *The following hold:*

- $\tilde{\mathcal{C}}^*$ and $\tilde{\mathcal{C}}$ are open, coordinate-invariant subsets of $\tilde{\Sigma}_{p,m}^{*,q}$.
- $\mathcal{C}$ is the smooth orbit manifold of $\tilde{\mathcal{C}}$.
- $\mathcal{C}^*$ is non-Hausdorff.

Proof. Define

$$\mathbf{I}(\mathbf{K}) := \det(\mathbf{C}\mathbf{C}^\top) \det(\mathbf{O}^\top \mathbf{O}),$$

where (A_c, B_c, C_c) is the controller realization. Then $\tilde{\mathcal{C}} = \mathbf{I}^{-1}(\mathbb{R} \setminus \{0\})$, which is open in $\tilde{\mathcal{C}}^*$. Openness of $\tilde{\mathcal{C}}^*$ in $\tilde{\Sigma}^*$ is shown in [22]. The remaining claims follow from Theorem 3.2.1. $\square$

Vertical Space And Retraction The vertical space at $\mathbf{K} \in \tilde{\mathcal{C}}$ is

$$\mathcal{V}_{\mathbf{K}} = \{(HA_c - A_cH, HB_c, -C_cH) : H \in \mathbb{R}^{q \times q}\}. \quad (3.10)$$

Since $\tilde{\mathcal{C}}$ is open, a coordinate-equivariant retraction is given by simple addition:

$$\tilde{R}_{\mathbf{K}}(\mathbf{V}) = \mathbf{K} + \mathbf{V}, \quad (3.11)$$

for sufficiently small $\mathbf{V}$.

3.2.3 Topology of Controller Spaces

In this section, we study the topology and dimension-theoretic structure of various controller spaces. In particular, we analyze the total controller space $\tilde{\mathcal{C}}^*$, the controller space $\tilde{\mathcal{C}}$, and the set of non-minimal stabilizing controllers

$$\mathcal{X} := \tilde{\mathcal{C}}^* \setminus \tilde{\mathcal{C}};$$

and their corresponding orbit spaces.

Theorem 3.2.3. *Suppose $\min(m, p) \geq 2$. Then the following hold:*

1. *The set $\mathcal{X}$ of non-minimal stabilizing controllers is a finite union of embedded submanifolds of $\tilde{\Sigma}^*$, is a closed semi-algebraic subset of $\tilde{\mathcal{C}}^*$, and has codimension $\min(m, p)$.*
2. *The controller space $\tilde{\mathcal{C}}$ has the same number of path-connected components as the total controller space $\tilde{\mathcal{C}}^*$.*
3. *The controller orbit space $\mathcal{C}$ has at most the same number of path-connected components as the controller space $\tilde{\mathcal{C}}$.*

The proof relies on the semi-algebraic structure of $\tilde{\mathcal{C}}^*$ and $\tilde{\mathcal{C}}$. In particular, we show that the set $\mathcal{X}$ is semi-algebraic and therefore admits a well-defined notion of dimension. A Whitney stratification of $\mathcal{X}$ yields a lower bound on its codimension, allowing us to conclude that $\mathcal{X}$ has insufficient dimension to separate $\tilde{\mathcal{C}}^*$. Consequently, the complement $\tilde{\mathcal{C}} = \tilde{\mathcal{C}}^* \setminus \mathcal{X}$ has the same number of path-connected components as $\tilde{\mathcal{C}}^*$. Our arguments build on the homological analysis of linear system spaces developed in [41].

Semi-Algebraic Preliminaries We briefly review the relevant notions from real algebraic geometry; see [41, 59, 60] for details.

Let $\mathbb{R}[x_1, \dots, x_n]$ denote the ring of real polynomials in n variables. Given a finite set $p = \{p_1, \dots, p_r\} \subset \mathbb{R}[x_1, \dots, x_n]$, a *basic closed semi-algebraic set* is

$$\{x \in \mathbb{R}^n : p_i(x) \leq 0, i = 1, \dots, r\}.$$

A general semi-algebraic set is any set obtained from basic closed semi-algebraic sets by a finite number of unions, intersections, and complements. Semi-algebraic sets are closed under projection, closure, interior, unions, and intersections.

Let $M : \mathbb{R}^n \rightarrow \text{Sym}(s)$ be a symmetric matrix polynomial. The associated *polynomial matrix inequality* is

$$\mathcal{S}(M) := \{x \in \mathbb{R}^n : M(x) \preceq 0\}.$$

If M is affine-linear, $\mathcal{S}(M)$ is a linear matrix inequality (LMI), also known as a spectrahedron.

Lemma 3.2.2. [61, §II.D], [59, Prop. 2.20] *Every polynomial matrix inequality defines a basic closed semi-algebraic set.*

Proof. Let $M : \mathbb{R}^n \rightarrow \text{Sym}(s)$ be symmetric. The condition $M(x) \preceq 0$ holds if and only if all eigenvalues of $M(x)$ are non-positive. Since $M(x)$ is symmetric, its

eigenvalues are real and coincide with the roots of the characteristic polynomial

$$\chi_{M(x)}(t) := \det(tI_s - M(x)).$$

Define $p_x(t) := (-1)^s \chi_{M(x)}(-t)$, which is a monic polynomial. Then $M(x) \preceq 0$ if and only if p_x has no positive real roots. By [59, Lemma 2.21], this is equivalent to non-negativity of its coefficients. Since these coefficients are polynomial functions of x , the claim follows. $\square$

Semi-Algebraicity of Controller Spaces We now show that $\tilde{\mathcal{C}}^*$, $\tilde{\mathcal{C}}$, and $\mathcal{X}$ are semi-algebraic. Define the matrix polynomial

$$F(P, \mathbf{K}) = \begin{bmatrix} -P & 0 \\ 0 & A_{\text{cl}}(\mathbf{K})P + PA_{\text{cl}}(\mathbf{K})^\top \end{bmatrix},$$

with $P \in \text{Sym}(n+q)$ and $\mathbf{K}$ a controller realization. The set $\mathcal{S}(F \prec 0) \subset \text{Sym}(n+q) \times \mathcal{T}$ is semi-algebraic, and its projection onto the controller variable yields

$$\tilde{\mathcal{C}}^* = \{\mathbf{K} \in \mathcal{T} : \exists P \succ 0 \text{ such that } F(P, \mathbf{K}) \prec 0\},$$

which is therefore semi-algebraic.

The set of minimal systems

$$\tilde{\Sigma} = \{(A, B, C) : f(A, B, C) \neq 0\}, \quad f = \det(\mathbf{C}\mathbf{C}^\top) \det(\mathbf{O}^\top \mathbf{O}),$$

is semi-algebraic, hence so are

$$\tilde{\mathcal{C}} = \tilde{\mathcal{C}}^* \cap \tilde{\Sigma}, \quad \mathcal{X} = \tilde{\mathcal{C}}^* \setminus \tilde{\mathcal{C}}.$$

Dimension And Codimension Of Non-Minimal Controllers For $1 \leq r \leq q$, let $\mathcal{R}'_r \subset \tilde{\Sigma}$ denote the set of systems with controllability rank r . By [41, Thm. 5.1], $\mathcal{R}'_r$ is a semi-algebraic embedded submanifold of dimension $q^2 + rp + qm$. The set

$$\mathcal{R}_r := \tilde{\mathcal{C}}^* \cap \mathcal{R}'_r$$

is open in $\mathcal{R}'_r$ and hence an analytic submanifold of the same dimension. Define

$$\mathcal{R}^\circ := \bigcup_{r=1}^{q-1} \mathcal{R}_r,$$

the set of non-controllable controllers. Then

$$\dim(\mathcal{R}^\circ) = q^2 + qm + (q-1)p.$$

An analogous construction for observability yields the semi-algebraic submanifold $\mathcal{O}^\circ$ of non-observable controllers with

$$\dim(\mathcal{O}^\circ) = q^2 + qp + (q-1)m.$$

Since

$$\mathcal{X} = \mathcal{R}^\circ \cup \mathcal{O}^\circ,$$

we obtain

$$\dim(\mathcal{X}) = q^2 + qp + qm - \min(m, p), \quad \text{codim}(\mathcal{X}) = \min(m, p).$$

Since $\mathcal{X}$ is a closed semi-algebraic subset of $\tilde{\mathcal{C}}^*$, is a finite union of embedded submanifolds, and has codimension at least 2, we may apply [41, Lemma 5.3] to conclude that $\tilde{\mathcal{C}}$ and $\tilde{\mathcal{C}}^*$ have the same number of path-connected components. The final claim of Theorem 3.2.3 follows from continuity of the orbit map $\pi : \tilde{\mathcal{C}} \rightarrow \mathcal{C}$.

Corollary 3.2.1. *In the case of full-order controllers ($q = n$), the controller space $\tilde{\mathcal{C}}_n$ has at most two path-connected components and the controller orbit space $\mathcal{C}_n$ is path-connected.*

Proof. It is shown in [22, Thm. 3.1] that $\tilde{\mathcal{C}}_n^*$ has at most two path-connected components. Since the orbit map $\pi : \tilde{\mathcal{C}}_n \rightarrow \mathcal{C}_n$ is a smooth surjection, the same holds for $\tilde{\mathcal{C}}_n$. Moreover, [22, Thm. 3.2] shows that these two components are related by a coordinate transformation $\mathcal{T}_S$ with $\det S < 0$, implying they lie in the same orbit. Hence $\mathcal{C}_n$ is path-connected. $\square$

3.3 First-Order Optimization over the Controller Orbit Space

The Krishnaprasad-Martin Metric Herein, we will introduce the particular Riemannian metric used for optimizing over $\tilde{\mathcal{C}}$. See §C.2 for a primer on Riemannian orbit geometry.

The metric, known as the Krishnaprasad-Martin (KM) metric, was introduced in [52]. Let $\mathbf{K} \in \tilde{\mathcal{C}}$ and $\mathbf{V} = (\mathbf{E}, \mathbf{F}, \mathbf{G}) \in \mathcal{T}$ be a tangent vector. Define the mappings

$$\begin{aligned} E(\mathbf{V}) &:= dA_{\text{cl}}|_{\mathbf{K}}(\mathbf{V}) = \begin{bmatrix} 0_{n \times n} & B\mathbf{G} \\ \mathbf{F}C & \mathbf{E} \end{bmatrix}, \\ F(\mathbf{V}) &:= dB_{\text{cl}}(\mathbf{V}) = \begin{bmatrix} 0_{n \times n} & 0_{n \times p} \\ 0_{n \times n} & \mathbf{F} \end{bmatrix}, \\ G(\mathbf{V}) &:= dC_{\text{cl}}(\mathbf{V}) = \begin{bmatrix} 0_{p \times n} & 0_{p \times n} \\ 0_{m \times n} & \mathbf{G} \end{bmatrix}. \end{aligned}$$

Consider now the following Riemannian metric,

$$g_{\mathbf{K}}^{\text{KM}}(\mathbf{V}, \mathbf{W}) := c_1 \text{tr}[W_o(\mathbf{K}) E(\mathbf{V}) W_c(\mathbf{K}) E(\mathbf{W})^\top] \quad (3.12a)$$

$$+ c_2 \text{tr}[F(\mathbf{V})^\top W_o(\mathbf{K}) F(\mathbf{W})] \quad (3.12b)$$

$$+ c_3 \text{tr}[G(\mathbf{V}) W_c(\mathbf{K}) G(\mathbf{W})^\top], \quad (3.12c)$$

where $c_1, c_2, c_3 > 0$ are constants. Here, $W_o(\mathbf{K})$ and $W_c(\mathbf{K})$ are, respectively, the observability and controllability Gramians of the closed-loop system:

$$W_c(\mathbf{K}) := L(A_{\text{cl}}(\mathbf{K}), B_{\text{cl}}(\mathbf{K})B_{\text{cl}}(\mathbf{K})^\top),$$

$$W_o(\mathbf{K}) := L(A_{\text{cl}}(\mathbf{K})^\top, C_{\text{cl}}(\mathbf{K})^\top C_{\text{cl}}(\mathbf{K})).$$

Remark these exist since the closed-loop system must be stable by definition.

Lemma 3.3.1. *For the system (A, B, C) in (3.1) and $\mathbf{K} \in \tilde{\mathcal{C}}_q$, the triplet*

$$(A_{\text{cl}}(\mathbf{K}), B_{\text{cl}}(\mathbf{K}), C_{\text{cl}}(\mathbf{K}))$$

is minimal.

Proof. This follows from the Popov–Belevitch–Hautus test and is omitted for brevity; see also [22, Lem. 4.5]. $\square$

Theorem 3.3.1. *The mapping defined in (3.12) is a Riemannian metric and coordinate-invariant.*

Proof. Note that for any $\mathbf{K}$, the mapping $g_{\mathbf{K}}^{\text{KM}}$ is smooth, bi-linear, symmetric, and hence it suffices to show positive-definiteness. By Lemma 3.3.1 and [62, Thm. 12.4], the Gramians satisfy,

$$\begin{aligned} W_c(\mathbf{K}) &\in \text{Sym}_{++}(n+q), \\ W_o(\mathbf{K}) &\in \text{Sym}_{++}(n+q). \end{aligned}$$

Hence, (3.12) is positive-definite.

Now, we show coordinate-invariance. Let $S \in \text{GL}(q)$ and set $\mathbf{L} := \mathcal{T}_S(\mathbf{K})$. Then

$$(A_{\text{cl}}(\mathbf{L}), B_{\text{cl}}(\mathbf{L}), C_{\text{cl}}(\mathbf{L})) = \mathcal{T}_{\hat{S}}(A_{\text{cl}}(\mathbf{K}), B_{\text{cl}}(\mathbf{K}), C_{\text{cl}}(\mathbf{K})), \quad \hat{S} := \text{diag}(I_n, S).$$

It follows that

$$\begin{aligned} W_c(\mathbf{L}) &= \hat{S} W_c(\mathbf{K}) \hat{S}^\top, \\ W_o(\mathbf{L}) &= \hat{S}^{-\top} W_o(\mathbf{K}) \hat{S}^{-1}, \end{aligned}$$

and similarly the differentials satisfy the corresponding equivariance relations. Plugging these identities into (3.12) yields the claim. $\square$

The horizontal space $\mathcal{H}_{\mathbf{K}}$ is the orthogonal complement of $\mathcal{V}_{\mathbf{K}}$ with respect to the KM metric. Constructing a basis for $\mathcal{H}_{\mathbf{K}}$ is detailed in Algorithm 1.

A few notes on Algorithm 1. First, see §3.3.2 on the construction of a global frame for $\tilde{\mathcal{C}}$. Second, $(\mathbf{V}_i)$ is a basis for the vertical space $\mathcal{V}_{\mathbf{K}}$, obtained by plugging in ∂_i^q into the expression in (3.10). Third, V is the corresponding *vectorized* basis for $\mathcal{V}_{\mathbf{K}}$. Fourth, H is a vectorized basis for the horizontal space $\mathcal{H}_{\mathbf{K}}$. Last, $(\mathbf{H}_i)$ is a basis for $\mathcal{H}_{\mathbf{K}}$.

Algorithm 1 Constructing the Basis for $\mathcal{H}_K$

- 1: **Input:** $K \in \tilde{\mathcal{C}}$
 - 2: $\mathbf{V}_i \leftarrow (\partial_i^q A_c - A_c \partial_i^q, \partial_i^q B_c, -C_c \partial_i^q)$ for $i = 1, \dots, q^2$
 - 3: $V \leftarrow \begin{bmatrix} [\mathbf{V}_1] & \dots & [\mathbf{V}_{q^2}] \end{bmatrix} \in \mathbb{R}^{N \times q^2}$
 - 4: Let $H = (H_1, \dots, H_{N-q^2}) \in \mathbb{R}^{N \times (N-q^2)}$ be any basis for $\ker(V^\top G(K))$
 - 5: $\mathbf{H}_i \leftarrow H_i^\wedge$ for $i = 1, \dots, N - q^2$
 - 6: **Return** $H, (\mathbf{H}_i)$
-

3.3.1 Riemannian Gradient Descent

In this section, we present a 1st-order Riemannian gradient method for solving (3.8). Algorithm 2 is RGD over the domain $\tilde{\mathcal{C}}$. We use the KM metric defined in §3.3. Since $\tilde{\mathcal{C}}$ is open, the optimizer uses the Euclidean retraction (3.11).

Algorithm 2 Riemannian Gradient Descent

- 1: **Input:** $K_0 \in \tilde{\mathcal{C}}, \epsilon > 0, \alpha_t \geq 0$
 - 2: $K \leftarrow K_0, t \leftarrow 0$
 - 3: **while** $\|\nabla \tilde{J}_K\|_K \geq \epsilon$ **do**
 - 4: $K \leftarrow K - \alpha_t \nabla \tilde{J}_K$
 - 5: $t \leftarrow t + 1$
 - 6: **end while**
 - 7: **Return:** K
-

Since $\tilde{\mathcal{C}}$ is open, one must choose a sufficiently small step size $\alpha_t > 0$ to ensure $K^+ := K - \alpha_t \nabla \tilde{J}_K$ remains stabilizing and minimal.

3.3.2 Computing the Riemannian Gradient

To compute the Riemannian gradient and other geometric quantities, we fix a global frame on $\tilde{\mathcal{C}}$. Since $\tilde{\mathcal{C}}$ is an open subset of the Euclidean space $\mathcal{T}$, any basis of $\mathcal{T}$ induces

a global frame on $\tilde{\mathcal{C}}$. We adopt the standard basis $(\boldsymbol{\partial}_i)_{i=1}^N \subset \mathcal{T}$, ordered arbitrarily, where

$$N := q^2 + qm + qp.$$

Here, $\boldsymbol{\partial}_i$ denotes the i th standard basis matrix of $\mathcal{T}$. With this choice of frame, the i th component of any tangent vector $\mathbf{V} \in \mathcal{T}$ is

$$[\mathbf{V}]_i := \langle \mathbf{V}, \boldsymbol{\partial}_i \rangle, \quad \mathbf{V} = \sum_{i=1}^N [\mathbf{V}]_i \boldsymbol{\partial}_i,$$

where $\langle \cdot, \cdot \rangle$ is the Frobenius inner product *mutatis mutandis*. We define the vectorization operator

$$[\mathbf{V}] := [[\mathbf{V}]_1, \dots, [\mathbf{V}]_N] \in \mathbb{R}^N, \quad \mathbf{V} \in \mathcal{T}$$

and its inverse mapping

$$w^\wedge := \sum_{i=1}^N \langle w, \boldsymbol{\partial}_i \rangle \boldsymbol{\partial}_i \in \mathcal{T}, \quad v \in \mathbb{R}^N.$$

Let $\partial_i \tilde{J}_{\mathbf{K}}$ denote the i th partial derivative of $\tilde{J}$ at $\mathbf{K}$ in these global coordinates. The Euclidean gradient is

$$\bar{\nabla} \tilde{J}_{\mathbf{K}} = [\partial_i \tilde{J}_{\mathbf{K}}]^\wedge = \sum_{i=1}^N (\partial_i \tilde{J}_{\mathbf{K}}) \boldsymbol{\partial}_i,$$

where $\bar{\nabla}$ is the classical Euclidean gradient.

Let $G(\mathbf{K}) := [g_{\mathbf{K}}^{\text{KM}}] \in \mathbb{R}^{N \times N}$ denote the components of the KM metric in the chosen frame. Then the Riemannian gradient satisfies

$$[\nabla \tilde{J}_{\mathbf{K}}] = G(\mathbf{K})^{-1} [\bar{\nabla} \tilde{J}_{\mathbf{K}}], \quad \nabla \tilde{J}_{\mathbf{K}} = (G(\mathbf{K})^{-1} [\bar{\nabla} \tilde{J}_{\mathbf{K}}])^\wedge.$$

There is a faster method involving Cholesky decomposition, since $G(\mathbf{K})$ is always positive definite. This is expressed in Algorithm 3.

Algorithm 3 Computing $\nabla \tilde{J}_{\mathbf{K}}$

```

1: Input:  $\mathbf{K} \in \tilde{\mathcal{C}}$ 
2:  $G \leftarrow 0_{N \times N}$ ,  $w \leftarrow 0_N$ 
3: for  $i = 1, \dots, N$  do
4:    $w_i \leftarrow \partial_i \tilde{J}(\mathbf{K})$ 
5:   for  $j = i, i + 1, \dots, N$  do
6:      $G_{ij}, G_{ji} \leftarrow g_{\mathbf{K}}^{\text{KM}}(\partial_i, \partial_j)$ 
7:   end for
8: end for
9:  $L \leftarrow \text{Cholesky}(G)$ 
10: Solve  $Lv = w$ 
11: Solve  $L^\top u = v$ 
12: Return:  $u^\wedge \in \mathcal{T}$ 

```

3.3.3 Why Gradient Descent does not Work

Let $\bar{\nabla} \tilde{J}_{\mathbf{K}} \in \mathcal{T}$ denote the Euclidean gradient of $\tilde{J}$ at $\mathbf{K}$. Gradient descent with the Euclidean metric is not coordinate-equivariant (Lemma C.2.1), that is,

$$\mathcal{T}_S(\mathbf{K} - s \cdot \bar{\nabla} \tilde{J}_{\mathbf{K}}) \neq \mathcal{T}_S(\mathbf{K}) - s \cdot \bar{\nabla} \tilde{J}_{\mathcal{T}_S(\mathbf{K})}.$$

As a result, Euclidean GD generally expends effort in the redundant coordinate directions, and it can be numerically ill-conditioned when the realization coordinates are poorly scaled. These issues are resolved when the metric is coordinate-invariant and the retraction is coordinate-equivariant.

3.3.4 Convergence Analysis

We establish a local convergence guarantee for Riemannian gradient descent (RGD). We begin by recalling a standard fixed-point result.

Theorem 3.3.2. [4, Theorem 4.19] Let $\mathcal{M}$ be a Riemannian manifold and let $\mathcal{U} \subset \mathcal{M}$ be open. Suppose $F : \mathcal{U} \rightarrow \mathcal{M}$ is smooth and admits a fixed point $p^* \in \mathcal{U}$ satisfying $\|dF_{p^*}\|_{p^*} < 1$. Then there exists a neighborhood $\mathcal{U}' \subset \mathcal{U}$ such that $F(\mathcal{U}') \subset \mathcal{U}'$ and, for any $p_0 \in \mathcal{U}'$, the iterates $p_{k+1} = F(p_k)$ converge to p^* at least linearly. If $\|dF_{p^*}\|_{p^*} = 0$, then the convergence is at least quadratic.

Theorem 3.3.3. Let $\tilde{J} \in \mathfrak{F}(\tilde{\mathcal{C}})$ be a coordinate-invariant and let $J \in \mathfrak{F}(\mathcal{C})$ denote the induced orbit function. That is, $J \circ \pi = \tilde{J}$. Suppose $\mathbf{K}^* \in \tilde{\mathcal{C}}$ is a minimum such that $\nabla^2 J_{\pi(\mathbf{K}^*)} > 0$. Then there exist $L > 0$ and a neighborhood $\tilde{\mathcal{U}}$ of $\mathbf{K}^*$ such that the sequence $\{\mathbf{K}_t\}$ defined by

$$\mathbf{K}_{t+1} = \tilde{F}(\mathbf{K}_t) := \mathbf{K}_t - \frac{1}{L} \nabla \tilde{J}_{\mathbf{K}_t}$$

satisfies

$$\lim_{t \rightarrow \infty} \pi(\mathbf{K}_t) = \pi(\mathbf{K}^*)$$

with at least a linear rate of convergence.

Proof. Since $\tilde{\mathcal{C}}$ is open, for sufficiently large $L > 0$ define

$$\tilde{\mathcal{U}} := \{\mathbf{K} \in \tilde{\mathcal{C}} : \mathbf{K} - \frac{1}{L} \nabla \tilde{J}_{\mathbf{K}} \in \tilde{\mathcal{C}}\}.$$

Then $\tilde{\mathcal{U}}$ is open, contains $\mathbf{K}^*$, and is coordinate-invariant. Setting $\mathcal{U} := \pi(\tilde{\mathcal{U}})$, define

$$F : \mathcal{U} \rightarrow \mathcal{C}, \quad F(\mathbf{k}) := R_{\mathbf{k}}\left(-\frac{1}{L} \nabla J_{\mathbf{k}}\right).$$

By coordinate invariance of $\tilde{J}$ and equivariance of the retraction, F is well-defined and satisfies $\pi(\tilde{F}(\mathbf{K})) = F(\pi(\mathbf{K}))$ for all $\mathbf{K} \in \tilde{\mathcal{U}}$. Smoothness of F follows from [4, Thm. 9.21].

Since R is a retraction, $dR_{\pi(\mathbf{K}^*)}(0) = \text{Id}$, and hence

$$dF_{\pi(\mathbf{K}^*)} = \text{Id} - \frac{1}{L} \nabla^2 J_{\pi(\mathbf{K}^*)}.$$

Therefore,

$$\|dF_{\pi(\mathbf{K}^*)}\| = \max\left(\left|1 - \frac{1}{L} \lambda_{\max}(\nabla^2 J_{\pi(\mathbf{K}^*)})\right|, \left|1 - \frac{1}{L} \lambda_{\min}(\nabla^2 J_{\pi(\mathbf{K}^*)})\right|\right).$$

For sufficiently large L , this quantity is strictly less than 1. The claim then follows from Theorem 3.3.2. $\square$

The condition $\nabla^2 J_{\pi(K^*)} > 0$ holds if and only if

$$\ker(\nabla^2 \tilde{J}_{K^*}) = \mathcal{V}_{K^*}.$$

3.3.5 Numerical Results

We now compare RGD with classical gradient descent (Figure 3.3). Our step size procedure is Algorithm 4.

Algorithm 4 Backtracking Line-Search

Require: $K \in \tilde{\mathcal{C}}$, $\gamma \in (0, 1)$, $\beta \in (0, 1)$, $\bar{s} > 0$

```

 $s \leftarrow \bar{s}$ 
 $K^+ \leftarrow K - s \nabla \tilde{J}_K$ 
while  $K^+ \notin \tilde{\mathcal{C}}$  or  $\tilde{J}(K) - \tilde{J}(K^+) < \gamma s \|\nabla \tilde{J}_K\|_K^2$  do
     $s \leftarrow \beta s$ 
     $K^+ \leftarrow K - s \nabla \tilde{J}_K$ 
end while return  $s$ 

```

The parameters in our algorithm were chosen as $T = 10^4$, $\gamma = 0.01$, $\beta = 0.5$, $\epsilon = 10^{-6}$, and $\bar{s} = 1$. We halted the simulation when $\tilde{J}(K) - \tilde{J}^* < 10^{-5}$. We initialized K_0 by generating a gain and observer with random pole placement in $(-2, -1)$. For GD, we used the same parameters and starting point. We compared GD against two KM metrics: **(1)** $c_1 = c_2 = c_3 = 1$ and **(2)** with $c_1 = 1$, $c_2 = c_3 = 0$.

We ran our numerical experiments against four representative systems.¹ The first system is Doyle's counterexample [63]. The second system is a plant whose LQG controller is non-minimal, and the third system admits saddle points with vanishing Hessians; these systems are found in [22]. The fourth system has dimensions

¹Our code is available at github.com/rainlabuw/riemannian-PO-for-LQG.

$(n, m, p) = (4, 3, 3)$ and entries either set to zero or sampled from the standard Gaussian distribution with probability 0.8 to promote sparsity.




Figure 3.3: Comparison of RGD vs. GD for LQG PO for four distinct systems.

As we observe, in all four cases, Algorithm 2 significantly outperforms GD. In fact, for the vanishing Hessian system, GD gets stuck in the non-strict saddle point.

3.4 Second-Order Optimization over the Controller Orbit Space

Let $\tilde{J} \in \mathfrak{F}(\tilde{\mathcal{C}})$ be coordinate-invariant. As reviewed in §2.2.4, we distinguish between the bilinear Hessian

$$\text{Hess}\tilde{J}_{\mathbf{K}} : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$$

and the Hessian operator

$$\nabla^2 \tilde{J}_{\mathbf{K}} : \mathcal{T} \rightarrow \mathcal{T},$$

which is the linear map used in second-order optimization methods. In components, the two are related by

$$[\nabla^2 \tilde{J}_{\mathbf{K}}]_{ij} = G^{ik}(\mathbf{K}) [\text{Hess}\tilde{J}_{\mathbf{K}}]_{kj}.$$

Bilinear Hessian The bilinear Hessian accounts for the curvature induced by the Riemannian metric. Its components are

$$[\text{Hess}\tilde{J}_{\mathbf{K}}]_{ij} = \partial_i \partial_j \tilde{J}(\mathbf{K}) - \Gamma_{ij}^k(\mathbf{K}) \partial_k \tilde{J}(\mathbf{K}),$$

where the Christoffel symbols are given by

$$\Gamma_{ij}^k(\mathbf{K}) = \frac{1}{2} G^{k\ell}(\mathbf{K}) (G_{\ell j;i}(\mathbf{K}) + G_{i\ell;j}(\mathbf{K}) - G_{ij;\ell}(\mathbf{K})), \quad G_{ij;k}(\mathbf{K}) := \partial_k G_{ij}(\mathbf{K}).$$

Hessian Operator The Hessian operator is represented in matrix form as

$$[\nabla^2 \tilde{J}_{\mathbf{K}}] = G^{-1}(\mathbf{K}) [\text{Hess}\tilde{J}_{\mathbf{K}}],$$

and for any $\mathbf{V} \in \mathcal{T}$ satisfies

$$[\nabla^2 \tilde{J}_{\mathbf{K}}(\mathbf{V})]_i = G^{ik}(\mathbf{K}) [\text{Hess}\tilde{J}_{\mathbf{K}}]_{kj} [\mathbf{V}]_j.$$

3.4.1 Why Newton's Descent Does not Work

The Riemannian Newton method updates the iterates as

$$\mathbf{K}_{t+1} = \mathbf{K}_t - \mathbf{V}_t,$$

where $\mathbf{V}_t$ solves

$$\nabla^2 \tilde{J}_{\mathbf{K}_t}(\mathbf{V}_t) = \nabla \tilde{J}_{\mathbf{K}_t}. \quad (3.13)$$

Here we assume that a solution $\mathbf{V}_t$ exists.

When this method admits a trajectory, we have experimentally found that it is not nearly as fast compared with RGD. The issue is that, while $\nabla \tilde{J}_{\mathbf{K}}$ is horizontal (and hence orthogonal to the orbit directions), the Newton direction $\mathbf{V}_t$ generally acquires vertical components, and these components do not contribute to decreasing a coordinate-invariant cost. Additionally, saddle points are attractive in the Newton landscape, which is problematic in the presence of degenerate saddles.

There is a way to recover a superlinear rate using Newton methodology. Since the KM metric is coordinate-invariant, the orbit space $\mathcal{C} = \tilde{\mathcal{C}}/\text{GL}(q)$ admits a Riemannian orbit manifold structure with fewer dimensions. As such, instead of performing Newton descent on $\tilde{\mathcal{C}}$, we perform Newton descent on the orbit manifold $\mathcal{C}$, and lift the resulting direction back to $\tilde{\mathcal{C}}$.

3.4.2 Riemannian Newton Descent Over the Controller Orbit Manifold

We now describe how to perform Newton descent directly on the controller orbit manifold $\mathcal{C} = \tilde{\mathcal{C}}/\text{GL}(q)$, rather than on the total controller space $\tilde{\mathcal{C}}$. We refer to this procedure as *orbit Newton descent* (orbit RND).

Let $\text{Proj}_{\mathcal{H}_{\mathbf{K}}} : \mathcal{T} \rightarrow \mathcal{H}_{\mathbf{K}}$ denote the orthogonal projection onto the horizontal space. Define

$$\mathcal{M} := \{\mathbf{k} \in \mathcal{C} : \nabla^2 J_{\mathbf{k}} \text{ is invertible}\}, \quad \tilde{\mathcal{M}} := \pi^{-1}(\mathcal{M}).$$

So, $\mathcal{M}$ is the subset of the orbit manifold on which Riemannian Newton descent is well-defined.

For $\mathbf{k}_t \in \mathcal{M}$, one step of Newton descent on the orbit manifold is given by

$$\mathbf{k}_{t+1} = R_{\mathbf{k}_t}(-\mathbf{v}_t),$$

where $\mathbf{v}_t \in T_{\mathbf{k}}\mathcal{C}$ is the unique solution of

$$\nabla^2 J_{\mathbf{k}_t}(\mathbf{v}_t) = \nabla J_{\mathbf{k}_t}. \quad (3.14)$$

Equation (3.14) defines the intrinsic Newton direction on the orbit manifold.

In practice, however, (3.14) is posed on $\mathcal{C}$, whereas optimization is performed on the total space $\tilde{\mathcal{C}}$. The following lemma shows how to recover the orbit Newton direction by solving a projected equation on $\tilde{\mathcal{C}}$.

Lemma 3.4.1. *[4, §9.12] Let $\mathbf{K} \in \tilde{\mathcal{M}}$, and let $\mathbf{v} \in T_{\pi(\mathbf{K})}\mathcal{C}$ be the unique solution of*

$$\nabla^2 J_{\pi(\mathbf{K})}(\mathbf{v}) = \nabla J_{\pi(\mathbf{K})}.$$

Suppose $\mathbf{W} \in T_{\mathbf{K}}\tilde{\mathcal{C}}$ satisfies

$$(\text{Proj}_{\mathcal{H}_{\mathbf{K}}} \circ \nabla^2 \tilde{J}_{\mathbf{K}} \circ \text{Proj}_{\mathcal{H}_{\mathbf{K}}})(\mathbf{W}) = \nabla \tilde{J}_{\mathbf{K}}. \quad (3.15)$$

Then there exists a unique vertical vector $\mathbf{V} \in \mathcal{V}_{\mathbf{K}}$ such that

$$\mathbf{W} = \text{lift}_{\mathbf{K}}(\mathbf{v}) + \mathbf{V}.$$

Moreover, the horizontal component satisfies

$$\text{Proj}_{\mathcal{H}_{\mathbf{K}}}(\mathbf{W}) = \text{lift}_{\mathbf{K}}(\mathbf{v}) \in \mathcal{H}_{\mathbf{K}}.$$

Lemma 3.4.1 shows that the Newton direction on the orbit manifold can be computed by solving the projected Newton equation (3.15) on the total controller space and retaining only its horizontal component. This procedure removes all vertical (coordinate) degrees of freedom and yields the intrinsic Newton direction on $\mathcal{C}$.

We summarize this procedure in Algorithm 5.

3.4.3 Convergence Analysis

Here we present a local quadratic convergence guarantee for orbit RND.

Algorithm 5 Orbit Riemannian Newton Descent

```

1: Input:  $\mathbf{K} \in \tilde{\mathcal{C}}$  such that  $\nabla^2 J_{\pi(\mathbf{K})}$  is invertible, tolerance  $\epsilon > 0$ 
2: while  $\|\nabla \tilde{J}_{\mathbf{K}}\|_{\mathbf{K}} \geq \epsilon$  do
3:   Construct a basis  $H$  for the horizontal space  $\mathcal{H}_{\mathbf{K}}$ 
4:    $P \leftarrow H(H^\top G(\mathbf{K})H)^{-1}H^\top G(\mathbf{K})$ 
5:    $H' \leftarrow P[\nabla^2 \tilde{J}_{\mathbf{K}}]P$ 
6:   Solve  $H'w = [\nabla \tilde{J}_{\mathbf{K}}]$ 
7:    $\mathbf{V} \leftarrow (Pw)^\wedge$ 
8:    $\mathbf{K} \leftarrow \mathbf{K} - \mathbf{V}$ 
9: end while
10: Return:  $\mathbf{K}$ 

```

Theorem 3.4.1. For $\mathbf{K} \in \tilde{\mathcal{C}}$, let

$$\tilde{H}_{\mathbf{K}} := \text{Proj}_{\mathcal{H}_{\mathbf{K}}}(\mathbf{v}) \in \mathcal{H}_{\mathbf{K}},$$

where $\mathbf{v}$ is any solution to

$$(\text{Proj}_{\mathcal{H}_{\mathbf{K}}} \circ \nabla^2 \tilde{J}_{\mathbf{K}} \circ \text{Proj}_{\mathcal{H}_{\mathbf{K}}})(\mathbf{v}) = \nabla \tilde{J}_{\mathbf{K}}. \quad (3.16)$$

Assume $\pi(\mathbf{K}) \in \mathcal{C}$ is such that $\nabla^2 J_{\pi(\mathbf{K})}$ is invertible. Then $\tilde{H}_{\mathbf{K}}$ is well-defined over a coordinate-invariant subset $\tilde{\mathcal{W}} \subset \tilde{\mathcal{C}}$, and there exists a neighborhood $\tilde{\mathcal{U}} \subset \tilde{\mathcal{W}}$ of $\mathbf{K}^*$ such that the map

$$\tilde{F} : \tilde{\mathcal{U}} \rightarrow \tilde{\mathcal{C}}, \quad \tilde{F}(\mathbf{K}) := \mathbf{K} - \tilde{H}_{\mathbf{K}}$$

is well-defined. Moreover, for any $\mathbf{K}_0 \in \tilde{\mathcal{U}}$ the iterates $\mathbf{K}_{t+1} = \tilde{F}(\mathbf{K}_t)$ satisfy

$$\lim_{t \rightarrow \infty} \pi(\mathbf{K}_t) = \pi(\mathbf{K}^*)$$

with quadratic rate.

Proof. Let

$$\mathcal{W} := \{\mathbf{k} \in \mathcal{C} : \nabla^2 J_{\mathbf{k}} \text{ is invertible}\}, \quad \tilde{\mathcal{W}} := \pi^{-1}(\mathcal{W}).$$

For $\mathbf{k} \in \mathcal{W}$, let $\mathbf{h}_{\mathbf{k}} \in T_{\mathbf{k}}\mathcal{C}$ denote the unique solution of the Newton equation

$$\nabla^2 J_{\mathbf{k}}(\mathbf{h}_{\mathbf{k}}) = \nabla J_{\mathbf{k}}.$$

By smoothness (in particular, analyticity) of J , the mapping $\mathbf{k} \mapsto \mathbf{h}_{\mathbf{k}}$ is smooth on $\mathcal{W}$.

Fix $\mathbf{K} \in \widetilde{\mathcal{W}}$ and set $\mathbf{k} := \pi(\mathbf{K})$. By Lemma 3.4.1, the projected equation (3.16) admits a solution $\mathbf{v}$ and the horizontal component $\text{Proj}_{\mathcal{H}_{\mathbf{k}}}(\mathbf{v})$ is unique. Moreover,

$$\widetilde{H}_{\mathbf{K}} = \text{lift}_{\mathbf{K}}(\mathbf{h}_{\mathbf{k}}) \in \mathcal{H}_{\mathbf{K}},$$

so $\widetilde{H}_{\mathbf{K}}$ is well-defined on $\widetilde{\mathcal{W}}$ and depends smoothly on $\mathbf{K}$.

Since $\widetilde{\mathcal{C}}$ is open, there exists a neighborhood $\widetilde{\mathcal{U}} \subset \widetilde{\mathcal{W}}$ of $\mathbf{K}^*$ such that $\mathbf{K} - \widetilde{H}_{\mathbf{K}} \in \widetilde{\mathcal{C}}$ for all $\mathbf{K} \in \widetilde{\mathcal{U}}$. Thus $\widetilde{F}(\mathbf{K}) = \mathbf{K} - \widetilde{H}_{\mathbf{K}}$ is well-defined on $\widetilde{\mathcal{U}}$. Setting $\mathcal{U} := \pi(\widetilde{\mathcal{U}})$, define the Newton map on the orbit manifold

$$F : \mathcal{U} \rightarrow \mathcal{C}, \quad F(\mathbf{k}) := R_{\mathbf{k}}(-\mathbf{h}_{\mathbf{k}}).$$

By construction and equivariance of the lift,

$$\pi(\widetilde{F}(\mathbf{K})) = F(\pi(\mathbf{K})) \quad \text{for all } \mathbf{K} \in \widetilde{\mathcal{U}}.$$

Finally, the standard Newton argument on manifolds yields

$$\|dF_{\pi(\mathbf{K}^*)}\|_{\pi(\mathbf{K}^*)} = 0$$

[64, Thm. 6.7]. Applying Theorem 3.3.2 yields a neighborhood $\mathcal{U}' \subset \mathcal{U}$ on which the orbit iterates converge quadratically to $\pi(\mathbf{K}^*)$. Lifting back to $\widetilde{\mathcal{C}}$ yields the stated claim. $\square$

3.4.4 Numerical Results

In this section, we compare orbit RND with RGD, shown in Figure 3.4. Each iteration of orbit RND is exactly Algorithm 5 with $\epsilon = 10^{-6}$. However, in the instance when

$\mathbf{K} - \mathbf{H}$ is either non-minimal, non-stabilizing, or fails to decrease the cost, we instead execute an iteration of RGD with backtracking, as described in [52, Alg. 2]. When this degeneracy flag is triggered, the algorithm parameters are set as $\gamma = .01$ and $\beta = 0.5$ while keeping $\epsilon = 10^{-6}$. We illustrate degeneracy events using orange and green dots in Figure 3.4.

We halt the simulation when $J(\mathbf{K}) - J^* < \epsilon$. The modified orbit RND procedure is compared against RGD with backtracking, with the same algorithm parameters as above. For both algorithms, we initialize $\mathbf{K}_0$ by generating a gain and observer with random pole placement in $(-2, -1)$. For both algorithms, the KM metric uses coefficients $c_1 = c_2 = c_3 = 1$.

We run numerical experiments against four representative systems. The first system is Doyle’s counterexample [63]. The second system is a plant whose LQG controller is non-minimal, and the third system admits saddle points with vanishing Hessians; these systems are found in [22]. The fourth system has dimensions $(n, m, p) = (4, 3, 3)$ with entries sampled from a standard Gaussian distribution.

We observe that for plants relatively free of degenerate saddles and admitting minimal LQG gains, there is an eventual quadratic convergence rate. The system with the vanishing Hessian saddle point remains a fatal flaw of vanilla Newton descent, since saddle points are attractive, and this is inherited by orbit RND unless one uses a quasi-Newton modification. In conclusion, while orbit RND is significantly faster than RGD in favorable instances, there is a lack of robustness for more degenerate plants.




Figure 3.4: Comparison of RGD vs. orbit RND for LQG PO for four distinct systems.

3.5 Concluding Remarks

This chapter developed a geometric framework for direct policy optimization over dynamic output-feedback controllers in the continuous-time LQG setting. The starting point was the observation that the objective is invariant under coordinate transformations of the controller state. This invariance induces degeneracy in classical optimization methods: stationary points appear as entire orbits, and the Hessian is necessarily singular along the corresponding coordinate directions.

Our first contribution was showing that the set of stabilizing *minimal* order- q controllers admits a smooth orbit manifold structure under the $GL(q)$ action. This provides a principled, intrinsic search space in which the redundant coordinate degrees of freedom are removed. Our second contribution was to equip this orbit manifold with a coordinate-invariant Riemannian metric and develop an intrinsic first-order optimization method, yielding a locally linear convergence rate guarantee.

Finally, we introduced orbit Riemannian Newton descent (orbit RND), which computes Newton steps intrinsically on the orbit manifold while implementing all computations on the total space. This method yields a local quadratic convergence rate.

Beyond LQG, the same viewpoint

Chapter 4

TRAJECTORY OPTIMIZATION OVER SMOOTH MANIFOLDS: EXPLOITING GEOMETRY

The introduction of coordinates is an act of violence.

Hermann Weyl [9]

A fundamental challenge in trajectory optimization on smooth manifolds is the choice of state-space parameterizations. For example, common parameterizations of rigid body systems include dual quaternions and homogeneous transformation matrices. Although these parameterizations describe the same underlying configuration space, for non-linear optimal control problems many cost objectives, linearization procedures, and algorithms depend directly on the selected parameterization. For example, the widespread use of dual quaternions in this setting is largely due to their unique and favorable numerical and algebraic properties [65–67].

The idea of intrinsic constrained optimization on manifolds is a mature area with especially renewed interest [68–71]. For intrinsic *trajectory* optimization, prior work has optimized directly over manifolds of trajectories on Lie groups [72, 73]. We develop this viewpoint within the framework of *successive convexification* (SCvx), a popular method for fast trajectory optimization [74, 75].

Standard SCvx implementations typically enforce manifold-valued dynamics *extrinsically*, either by enforcing the dynamics (which characterize the implicit manifold constraint), relying on a projection procedure, or by directly enforcing the manifold constraint. This can introduce redundant search directions and increase memory usage and computation time [76].

The main contribution of this chapter is an SCvx-like framework that is invari-

ant to parameterization, which we call *intrinsic SCvx* (i-SCvx) [77]. Consequently, algorithm performance is also parameterization-invariant, provided the objective, dynamics, and constraints depend only on coordinate-free quantities defined on the state and input manifolds. We develop i-SCvx for discrete-time systems over a finite horizon and show how it resolves the redundancy and parameterization dependence inherent to extrinsic SCvx formulations.

The chapter is organized as follows. §4.1 formulates the trajectory optimization problem on smooth manifolds. §4.2 reviews SCvx and its local convex subproblem. §4.3 introduces i-SCvx, including intrinsic linearization of dynamics and constraints, construction of the local objective, and the resulting convex subproblem. §4.4 presents the spacecraft attitude guidance example and numerical comparisons between SCvx and i-SCvx. Finally, §4.5 concludes the chapter and discusses limitations and future directions.

4.1 Problem Statement

Let $\mathcal{M}$ and $\mathcal{U}$ be Riemannian manifolds, characterizing the system and input manifolds. For example, $\mathcal{M} = \text{SO}(3)$ and $\mathcal{U} = \mathbb{R}^3$. Let $f: \mathcal{M} \times \mathcal{U} \rightarrow \mathcal{M}$ be the system dynamics. The stage and final stage costs are denoted $\phi: \mathcal{M} \times \mathcal{U} \rightarrow \mathbb{R}$ and $\phi_f: \mathcal{M} \rightarrow \mathbb{R}$. Last, we define the constraints function $g: \mathcal{M} \times \mathcal{U} \rightarrow \mathbb{R}^{n_g}$. We assume all function are twice continuously differentiable, with the stage and final stage costs also geodesically convex.

Let $N > 0$ be the time horizon of the problem, $\mathbf{x} = (x_0, \dots, x_N) \in \mathcal{M}^{N+1}$ the state trajectory, and $\mathbf{u} = (u_0, \dots, u_{N-1}) \in \mathcal{U}^N$ the control sequence. The trajectory optimization problem of interest is solving

$$\min_{\mathbf{x}, \mathbf{u}} \quad C(\mathbf{x}, \mathbf{u}) := \sum_{k=0}^{N-1} \phi(x_k, u_k) + \phi_f(x_N) \quad (4.1a)$$

$$\text{s.t.} \quad x_{k+1} = f(x_k, u_k), \quad (4.1b)$$

$$g(x_k, u_k) \leq 0, \quad (4.1c)$$

$$x_0 \text{ given.} \quad (4.1d)$$

The key idea pursued in this work is solving (4.1) in an “intrinsic” manner. That is, the framework must be completely independent of how $\mathcal{M}$ and $\mathcal{U}$ are embedded in their ambient vector spaces.

4.2 SCvx

In this section, we introduce and derive the SCvx framework, and highlight the limitations of SCvx in the case the states and inputs lie on smooth manifolds rather than linear parameter spaces. Remark SCvx descends from the Sequential Convex Programming (SCP) paradigm [78].

4.2.1 The Main Algorithm

Focus on the main problem (4.1) with $\mathcal{M} = \mathbb{R}^n$ and $\mathcal{U} = \mathbb{R}^m$. Consider a possibly infeasible trajectory $(\mathbf{x}, \mathbf{u})$. A convex local subproblem is constructed to compute the optimal perturbation $(\delta\mathbf{x}, \delta\mathbf{u})$ around $(\mathbf{x}, \mathbf{u})$, minimizing the objective while satisfying the relaxed constraints:

$$\min_{\delta\mathbf{x}, \delta\mathbf{u}, \mathbf{v}, \mathbf{s}} \quad C(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}) + \sum_{k=0}^{N-1} \lambda_k P(v_k, s_k) \quad (4.2a)$$

$$\text{s.t.} \quad g(x_k, u_k) + S_k \delta x_k + T_k \delta u_k \leq s_k, \quad (4.2b)$$

$$f(x_k, u_k) - x_{k+1} + A_k \delta x_k + B_k \delta u_k = v_k, \quad (4.2c)$$

$$s_k \geq 0, \quad \|\delta x_k\|_2 \leq r, \quad \|\delta u_k\|_2 \leq r. \quad (4.2d)$$

Here, we define

$$\begin{aligned} A_k &:= \left. \frac{\partial f}{\partial x} \right|_{(x_k, u_k)}, & B_k &:= \left. \frac{\partial f}{\partial u} \right|_{(x_k, u_k)}, \\ S_k &:= \left. \frac{\partial g}{\partial x} \right|_{(x_k, u_k)}, & T_k &:= \left. \frac{\partial g}{\partial u} \right|_{(x_k, u_k)}. \end{aligned}$$

as the Jacobians. We will walk through the development of this subproblem.

In order to “localize” Problem 4.1 about $(\mathbf{x}, \mathbf{u})$, it makes sense to linearize the constraints about $(\mathbf{x}, \mathbf{u})$, resulting in

$$\min_{\delta \mathbf{x}, \delta \mathbf{u}} C(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u}) \quad (4.3a)$$

$$\text{s.t.} \quad g(x_k, u_k) + S_k \delta x_k + T_k \delta u_k \leq 0, \quad (4.3b)$$

$$f(x_k, u_k) - x_{k+1} + A_k \delta x_k + B_k \delta u_k = 0. \quad (4.3c)$$

The issue is while (4.3) is convex, it is often infeasible. As a result, we must relax the linearized constraints with slack variables, as well as append a penalty function onto the objective for the slack variables. Let $P(v, s) := \|v\|_1 + \|s\|_1$ and $\lambda_k \geq 0$. Focus on the problem

$$\min_{\delta \mathbf{x}, \delta \mathbf{u}, \mathbf{v}, \mathbf{s}} C(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u}) + \sum_{k=0}^{N-1} \lambda_k P(v_k, s_k) \quad (4.4a)$$

$$\text{s.t.} \quad g(x_k, u_k) + S_k \delta x_k + T_k \delta u_k \leq s_k, \quad (4.4b)$$

$$s_k \geq 0, \quad (4.4c)$$

$$f(x_k, u_k) - x_{k+1} + A_k \delta x_k + B_k \delta u_k = v_k. \quad (4.4d)$$

Take note that while (4.4) is feasible, it may possibly be unbounded. To ensure (4.4) is bounded, we bound the perturbations $\|\delta x_k\|_2, \|\delta u_k\|_2 \leq r$. This motivates Problem 4.2.

We solve Problem 4.2, which is convex, obtaining optimal perturbations. Then, we set the new anchor trajectory as the perturbed trajectory $(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u})$. The process repeats.

Trust Region Update We next must discuss how the trust region radius is updated. Details are elaborated in [79]. Let the penalized objective be defined as

$$J(\mathbf{x}, \mathbf{u}) := C(\mathbf{x}, \mathbf{u}) + \sum_{k=1}^{N-1} \lambda_k P(x_{k+1} - f(x_k, u_k), [g(x_k, u_k)]_+) \quad (4.5)$$

with $[v]_+$ is the element-wise application of $\max(v, 0)$. Intuitively, (4.5) captures both the optimality of $(\mathbf{x}, \mathbf{u})$ and how much the reference trajectory violates the dynamic

and state-input constraints. If $(\mathbf{x}, \mathbf{u})$ were admissible, then $P(x_{k+1} - f(x_k, u_k), [g(x_k, u_k)]_+) = 0$. We will also call (4.2a) the *localized objective*:

$$L(\delta\mathbf{x}, \delta\mathbf{u}, \mathbf{v}, \mathbf{s}) = C(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}) + \sum_{k=0}^{N-1} \lambda_k P(v_k, s_k).$$

Next, let $(\delta\mathbf{x}^*, \delta\mathbf{u}^*, \mathbf{v}^*, \mathbf{s}^*)$ be the solution of the localized subproblem (4.2) about $(\mathbf{x}, \mathbf{u})$. We compare the reduction in the true penalized objective to the reduction predicted by the localized objective.

Define the actual and predicted reductions as

$$\Delta J := J(\mathbf{x}, \mathbf{u}) - J(\mathbf{x} + \delta\mathbf{x}^*, \mathbf{u} + \delta\mathbf{u}^*), \quad \Delta L := J(\mathbf{x}, \mathbf{u}) - L(\delta\mathbf{x}^*, \delta\mathbf{u}^*, \mathbf{v}^*, \mathbf{s}^*).$$

If $\Delta L > 0$, the trust-region ratio is defined as

$$\rho := \frac{\Delta J}{\Delta L}.$$

The ratio ρ compares the actual reduction in the penalized objective to the reduction predicted by the localized model. Values $\rho \approx 1$ indicate close agreement between the localized model and the true model. Small positive values of ρ indicate that the localized model is overly optimistic, predicting a larger decrease than is actually realized.

If $\rho < 0$, the step fails to reduce the penalized objective and is rejected. If $0 < \rho < 1$, the step produces a decrease, but the model is unreliable and the trust-region radius should be reduced. If ρ is sufficiently close to one, the model is deemed accurate and the trust region is expanded. An intermediate range of ρ values is used to accept the step and leave the trust-region radius unchanged.

Following the above intuition, the trust region radius is updated according to

$$r^+ = \begin{cases} \alpha r & \text{if } \rho < \rho_1, \\ r & \text{if } \rho_1 \leq \rho < \rho_2, \\ \beta r & \text{if } \rho \geq \rho_2 \end{cases}$$

where $0 \leq \rho_0 < \rho_1 < \rho_2 < 1$ and $0 < \alpha < 1 < \beta$ are the algorithm parameters.

The complete SCvx methodology is presented in Algorithm 6.

Algorithm 6 (Discrete) SCvx Algorithm

Input: Choose initial trajectory $(\mathbf{x}, \mathbf{u})$. Select tolerance ϵ_{tol} , trust region $r > 0$, penalty weight $\lambda > 0$, and parameters $0 \leq \rho_0 < \rho_1 < \rho_2 < 1$ and $0 < \alpha < 1 < \beta$, and sufficiently large $\Delta J > 0$.

while $|\Delta J| > \epsilon_{\text{tol}}$ **do**

Solve Problem (4.2) to get $\delta \mathbf{x}, \delta \mathbf{u}, \mathbf{v}, \mathbf{s}$.

Compute

$$\Delta J \leftarrow J(\mathbf{x}, \mathbf{u}) - J(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u})$$

$$\Delta L \leftarrow J(\mathbf{x}, \mathbf{u}) - L(\delta \mathbf{x}, \delta \mathbf{u}, \mathbf{v}, \mathbf{s})$$

Compute $\rho \leftarrow \Delta J / \Delta L$.

if $\rho < \rho_0$ **then**

$$\text{padding-left: 4em;} r \leftarrow \alpha r$$

else

Update $\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u}$

$$\text{padding-left: 4em;} \text{Update } r \leftarrow \begin{cases} \alpha r, & \text{if } \rho < \rho_1; \\ r, & \text{if } \rho_1 \leq \rho < \rho_2; \\ \beta r, & \text{if } \rho \geq \rho_2. \end{cases}$$

end if

end while

return $\mathbf{x}, \mathbf{u}$

4.2.2 Issues with SCvx: non-Euclidean States and Inputs

A significant issue arises when the states and inputs are points on smooth manifolds rather than vectors in linear parameter spaces. In particular, for dynamics of the form $x_{k+1} = f(x_k, u_k)$, there exists a Euclidean submanifold $\mathcal{M} \subset \mathbb{R}^d$ such that $f(\mathcal{M}, \mathcal{U}) \subset \mathcal{M}$ and $\dim(\mathcal{M}) \leq d$.

In the resulting subproblem, searching over all directions in $\mathbb{R}^d$ for the optimal perturbation is inefficient. Only directions tangential to $\mathcal{M}$ at the current point are meaningful, as non-tangential directions primarily exacerbate violations of the dynamical constraints.

A related challenge is manifold reparameterization. The subproblem and objective are strictly parameterization-dependent, whereas the underlying dynamics are often intrinsic, as is typical for physical systems. Consequently, the same system can be represented in many equivalent parameterizations of $\mathcal{M}$, motivating a methodology whose performance is invariant under changes of representation.

Another critical aspect concerns how the reference trajectory is updated after solving the subproblem. One option is to project the updated trajectory onto $\mathcal{M}$, but such projections are not modeled in the subproblem itself. Alternatively, avoiding projection preserves the subproblem solution and can improve algorithmic robustness, but it risks violating the manifold constraint, requiring additional iterations to return the trajectory to the state-space manifold.

Finally, an additional issue arises from the use of *virtual control* terms in SCvx. When the state evolves on a manifold embedded in an ambient Euclidean space, deviations of a trajectory can be decomposed into three components: (1) error in dynamic feasibility along tangential directions of the manifold, (2) violation of explicit state-control constraints, and (3) violation of the implicit manifold constraint induced by orthogonal perturbations. In extrinsic SCvx formulations, the virtual control term absorbs all three effects simultaneously. As a result, feasibility of the localized sub-

problem is achieved partly by compensating for parameterization-induced artifacts, rather than exclusively for approximation error in the dynamics and constraints.

4.3 Intrinsic SCvx

Intrinsic SCvx (i-SCvx) is a natural generalization of SCvx to Riemannian manifolds. Rather than linearizing dynamics and constraints in an ambient Euclidean space, all linearizations are performed with respect to the smooth manifold structure, so perturbations live in the appropriate tangent spaces.

Because perturbations cannot be added directly to the state variables as in Problem 4.2, state updates are carried out via a retraction. That is, after solving the local subproblem, the new reference trajectory is formed by mapping the optimal tangent perturbations back to the manifolds.

Finally, in Euclidean SCvx the objective appearing in the local subproblem is convex, and that convexity is exploited directly in the convex program. In i-SCvx, the localized objective is instead replaced by its second-order Taylor approximation on the manifold, which is intrinsic. If the stage cost is geodesically convex (over the chosen domain), then this second-order expansion is a convex quadratic, and the resulting local subproblem is a convex quadratic program.

4.3.1 Linearizing Dynamics on Smooth Manifolds

In this section, we explore the linearization of dynamical systems on a smooth manifold. Linearizing dynamics on smooth manifolds has applications in nonlinear consensus, distributed optimization, and stability analysis of nonlinear systems [80–82].

Linearization About a Trajectory: Euclidean Case Consider a discrete-time system

$$x_{k+1} = f(x_k, u_k), \quad f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n,$$

with a given initial state x_0 . We assume f is continuously differentiable. Let $(\mathbf{x}, \mathbf{u})$ be a state-input sequence that is not assumed to satisfy the dynamics, where $\mathbf{x} =$

$(x_0, x_1, \dots, x_N)$ and $\mathbf{u} = (u_0, u_1, \dots, u_{N-1})$. Define the *dynamics defect*

$$r_k := f(x_k, u_k) - x_{k+1}.$$

Clearly, the trajectory $(\mathbf{x}, \mathbf{u})$ is dynamically feasible if and only if $r_k = 0$ for all k .

Now consider perturbations $(\delta\mathbf{x}, \delta\mathbf{u})$ and define the *perturbed trajectory*

$$\bar{x}_k := x_k + \delta x_k, \quad \bar{u}_k := u_k + \delta u_k.$$

We wish to characterize what perturbations $(\delta\mathbf{x}, \delta\mathbf{u})$ yield a dynamically feasible perturbed trajectory $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$. Assuming this, we can write

$$\delta x_{k+1} = f(x_k + \delta x_k, u_k + \delta u_k) - x_{k+1}.$$

A first-order Taylor expansion about (x_k, u_k) gives us

$$\delta x_{k+1} \approx r_k + A_k \delta x_k + B_k \delta u_k, \tag{4.6}$$

where

$$A_k := \left. \frac{\partial f}{\partial x} \right|_{(x_k, u_k)}, \quad B_k := \left. \frac{\partial f}{\partial u} \right|_{(x_k, u_k)}.$$

This defines a linear time-varying system with an affine forcing term given by the defect r_k . The linearized system describes how small perturbations must evolve in order to correct the dynamics violation, to first order.

The same can be said for the constraints. Suppose we also require $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$ to satisfy the constraints $g(x_k, u_k) \leq 0$. This yields

$$g(x_k + \delta x_k, u_k + \delta u_k) \leq 0.$$

A first-order Taylor expansion on the right-hand side about (x_k, u_k) gives us

$$g_k + S_k \delta x_k + T_k \delta u_k \leq 0 \tag{4.7}$$

where $g_k := g(x_k, u_k)$ is the *constraints defect* and

$$S_k := \left. \frac{\partial g}{\partial x} \right|_{(x_k, u_k)}, \quad T_k := \left. \frac{\partial g}{\partial u} \right|_{(x_k, u_k)}.$$

In summary, if we find a perturbation trajectory $(\delta \mathbf{x}, \delta \mathbf{u})$ from (4.6) and subject it to the linear constraints (4.7), then $(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u})$ will be approximately feasible in terms of dynamics and constraints, to first-order. We can repeat this process again, over and over, and the resulting trajectory will become more and more feasible in dynamics and constraints. This is a basic concept that is utilized in SCvx. Equivalently, the LTV system (4.6) subject to (4.7) describes how perturbations must evolve to correct the dynamics and constraints violations, to first order.

Linearization About a Trajectory: Smooth Manifold Case Let $\mathcal{M}$ and $\mathcal{U}$ be smooth manifolds. Let $R^{\mathcal{M}} : T\mathcal{M} \rightarrow \mathcal{M}$ and $R^{\mathcal{U}} : T\mathcal{U} \rightarrow \mathcal{U}$ be retractions. For simplicity of notation, we will refer to both retractions as R .

Consider a discrete-time system

$$x_{k+1} = f(x_k, u_k), \quad f : \mathcal{M} \times \mathcal{U} \rightarrow \mathcal{M}.$$

We assume f is continuously differentiable. As before, we will select a state-input trajectory $(\mathbf{x}, \mathbf{u})$ that is possibly infeasible. The dynamics defect becomes

$$r_k := R_{x_{k+1}}^{-1}(f(x_k, u_k)) \in T_{x_{k+1}}\mathcal{M}.$$

If x_{k+1} and $f(x_k, u_k)$ are too far apart, then the dynamics defect is not well-defined, so we will assume they are sufficiently close. Remark, $(\mathbf{x}, \mathbf{u})$ is dynamically feasible if and only if $r_k = 0$ for all k .

Next, consider perturbations $\delta x_k \in T_{x_k}\mathcal{M}$ and $\delta u_k \in T_{u_k}\mathcal{U}$. We will write the perturbed trajectory as

$$\bar{x}_k := R_{x_k}(\delta x_k), \quad \bar{u}_k := R_{u_k}(\delta u_k)$$

As before, we wish to characterize what perturbations $(\delta \mathbf{x}, \delta \mathbf{u})$ yield a dynamically feasible perturbed trajectory $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$. By requiring $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$ to satisfy the dynamics, we can write

$$\delta x_{k+1} = R_{x_{k+1}}^{-1}(f(R_{x_k}(\delta x_k), R_{u_k}(\delta u_k))).$$

Of course, this assumes δx_k and δu_k are sufficiently small to make the above equation well-defined.

A first-order Taylor expansion about (x_k, u_k) gives us

$$\delta x_{k+1} \approx r_k + \mathbf{D}_k(\mathbf{A}_k(\delta x_k) + \mathbf{B}_k(\delta u_k)), \quad (4.8)$$

where

$$\begin{aligned} \mathbf{A}_k &:= d_x f_{(x_k, u_k)} : T_{x_k} \mathcal{M} \rightarrow T_{f(x_k, u_k)} \mathcal{M}, \\ \mathbf{B}_k &:= d_u f_{(x_k, u_k)} : T_{u_k} \mathcal{U} \rightarrow T_{f(x_k, u_k)} \mathcal{M}, \\ \mathbf{D}_k &:= dR_{x_{k+1}}^{-1}|_{f(x_k, u_k)} : T_{f(x_k, u_k)} \mathcal{M} \rightarrow T_{x_{k+1}} \mathcal{M}. \end{aligned}$$

Although we are not on a single vector space anymore, this still defines a linear time-varying system, except with each state and control term lying on different vector spaces! But it is still nonetheless a linear time-varying system with an affine forcing term given by the defect r_k .

The operator $\mathbf{D}_k$ accounts for the change in base point induced by the retraction inverse and serves to transport variations from $T_{f(x_k, u_k)} \mathcal{M}$ to $T_{x_{k+1}} \mathcal{M}$. This is similar to, but not the same, as parallel transport; it is simply the Jacobian of the chosen coordinate mapping induced by the retraction. We can simply think of $\mathbf{D}_k$ as describing how R_x^{-1} changes, as a mapping, as we vary x .

If the nominal trajectory is dynamically feasible, then $r_k = 0$ and the standard variational dynamics are recovered; furthermore, since we would have $f(x_k, u_k) = x_{k+1}$ and consequently $\mathbf{D}_k = \text{id}$. If the trajectory is infeasible, the defect term quantifies the local violation of the dynamics, and the linearized system describes how small perturbations must evolve in order to correct this violation to first order.

The same can be said for the constraints. Suppose we also require $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$ to satisfy the constraints $g(x_k, u_k) \leq 0$. This yields

$$g(R_{x_k}(\delta x_k), R_{u_k}(\delta u_k)) \leq 0.$$

A first-order Taylor expansion on the right-hand side about (x_k, u_k) gives us

$$g_k + \mathbf{S}_k(\delta x_k) + \mathbf{T}_k(\delta u_k) \leq 0 \quad (4.9)$$

where $g_k := g(x_k, u_k)$ is the constraints defect and

$$\mathbf{S}_k := d_x g(x_k, u_k) : T_{x_k} \mathcal{M} \rightarrow \mathbb{R}^{n_g}, \quad \mathbf{T}_k := d_u g(x_k, u_k) : T_{u_k} \mathcal{U} \rightarrow \mathbb{R}^{n_g}.$$

In summary, if we solve for some perturbation trajectory $(\delta \mathbf{x}, \delta \mathbf{u})$ from (4.8) and subject to the linear constraints (4.9), then $(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u})$ will be approximately feasible in dynamics and constraints. We can repeat this process again, over and over, and the resulting trajectory will become more and more feasible in dynamics and constraints. This is a basic concept that is utilized in SCvx.

4.3.2 Taylor Series on Riemannian Manifolds

Let $(\mathcal{M}, g)$ be a Riemannian manifold. Let $f \in \mathfrak{F}(\mathcal{M})$ and $p \in \mathcal{M}$. The first-order $f_p^{(1)} : T_p \mathcal{M} \rightarrow \mathbb{R}$ and second-order $f_p^{(2)} : T_p \mathcal{M} \rightarrow \mathbb{R}$ expansions of f at p are defined as

$$\begin{aligned} f_p^{(1)}(\eta) &:= f(p) + df_p(\eta) = f(p) + g_p(\nabla f_p, \eta), \\ f_p^{(2)}(\eta) &:= f_p^{(1)}(\eta) + \frac{1}{2} \text{Hess} f_p(\eta, \eta) = f_p^{(1)}(\eta) + \frac{1}{2} g_p(\nabla^2 f_p(\eta), \eta), \end{aligned}$$

respectively. Remark the first-order expansion of f is an affine-linear mapping, and the second-order expansion of f is a quadratic mapping. If f were geodesically convex, then $f_p^{(2)}$ would be a positive semi-definite quadratic mapping. Remark we make use of both the Hessian operator and bilinear Hessian of f (see §2.2.4).

Proposition 4.3.1. *Let $f \in \mathfrak{F}(\mathcal{M})$ and $p \in \mathcal{M}$. Let $\gamma(\cdot)$ be a curve with $\gamma(0) = p$ and $\dot{\gamma}(0) = \eta$. Then*

$$f(\gamma(t)) = f_p^{(1)}(t\eta) + o(t) = f_p^{(2)}(t\eta) + o(t^2)$$

4.3.3 Local Objective: Exploiting Geodesic Convexity

We will describe how to construct the localized trajectory cost for the subproblem. We require that our stage and final stage costs are geodesically convex with respect to the Riemannian structures of $\mathcal{M}$ and $\mathcal{U}$. This will imply that the second-order approximation is a positive semi-definite quadratic function, thereby making the subproblem convex.

Suppose the stage costs $\phi(\cdot, \cdot), \phi_f(\cdot)$ are geodesically convex. We also assume that the partial differential

$$\mathrm{d}_x \phi_{(x,u)} : T_x \mathcal{M} \rightarrow \mathbb{R}$$

is independent of the input variable u . An analogous assumption is made for $\mathrm{d}_u \phi_{(x,u)}$ with respect to x .

With this in mind, their second-order approximations are convex quadratics:

$$\begin{aligned} \hat{\phi}_{(x,u)}^{(2)}(\eta, \xi) &= \phi(x, u) + \mathrm{d}_x \phi_{(x,u)}(\eta) + \mathrm{d}_u \phi_{(x,u)}(\xi) \\ &\quad + \frac{1}{2} \mathrm{Hess}_{xx} \phi_{(x,u)}(\eta, \eta) + \frac{1}{2} \mathrm{Hess}_{uu} \phi_{(x,u)}(\xi, \xi), \\ \hat{\phi}_f^{(2)}|_x(\eta) &= \phi_f(x) + \mathrm{d}_x \phi_f|_x(\eta) + \frac{1}{2} \mathrm{Hess} \phi_f|_x(\eta, \eta). \end{aligned}$$

This implies that the trajectory cost is geodesically convex, and so its second-order approximation will be a convex quadratic:

$$\hat{C}_{(\mathbf{x}, \mathbf{u})}^{(2)}(\delta \mathbf{x}, \delta \mathbf{u}) = \sum_{k=0}^{N-1} \hat{\phi}_{(x_k, u_k)}^{(2)}(\delta x_k, \delta u_k) + \hat{\phi}_f^{(2)}|_{x_N}(\delta x_N). \quad (4.10)$$

4.3.4 Slack Variables and Penalty

Like SCvx, we use slack variables $s_k \geq 0$, $v_k \in T_{x_{k+1}} \mathcal{M}$ to ensure that the subproblem is feasible:

$$\begin{aligned} g(x_k, u_k) + \mathbf{S}_k(\delta x_k) + \mathbf{T}_k(\delta u_k) &\leq s_k, \\ R_{x_{k+1}}^{-1}(f(x_k, u_k)) + \mathbf{D}_k \circ (\mathbf{A}_k(\delta x_k) + \mathbf{B}_k(\delta u_k)) + v_k &= \delta x_{k+1}. \end{aligned}$$

We also require an exact penalty term in order to minimize usage of these slack variables. Define a function $P : T\mathcal{M} \times \mathbb{R}^{n_g} \rightarrow \mathbb{R}$ using only coordinate-free quantities. For instances, $P(v, s) = \|v\|_{x_{k+1}} + \|s\|_1$. From here, we can define the penalized objective as

$$J(\mathbf{x}, \mathbf{u}) := C(\mathbf{x}, \mathbf{u}) + \sum_{k=0}^{N-1} \lambda_k P(R_{x_{k+1}}^{-1}(f(x_k, u_k)), [g_k]_+) \quad (4.11)$$

4.3.5 The Local Sub-problem

The local subproblem about $(\mathbf{x}, \mathbf{u})$ is

$$\min_{\delta\mathbf{x}, \delta\mathbf{u}, \mathbf{v}, \mathbf{s}} \hat{C}_{(\mathbf{x}, \mathbf{u})}^{(2)}(\delta\mathbf{x}, \delta\mathbf{u}) + \sum_{k=0}^{N-1} \lambda_k P(v_k, s_k) \quad (4.12a)$$

$$\text{s.t. } \delta x_{k+1} = R_{x_{k+1}}^{-1}(f(x_k, u_k)) + \tilde{\mathbf{A}}_k(\delta x_k) + \tilde{\mathbf{B}}_k(\delta u_k) + v_k, \quad (4.12b)$$

$$g(x_k, u_k) + \mathbf{S}_k(\delta x_k) + \mathbf{T}_k(\delta u_k) \leq s_k, \quad (4.12c)$$

$$s_k \geq 0, \quad \|\delta x_k\|_{x_k} \leq r, \quad \|\delta u_k\|_{u_k} \leq r. \quad (4.12d)$$

where $\tilde{\mathbf{A}}_k := \mathbf{D}_k \circ \mathbf{A}_k$ and $\tilde{\mathbf{B}}_k := \mathbf{D}_k \circ \mathbf{B}_k$. We call (4.12a) the localized objective and denote it $L(\delta\mathbf{x}, \delta\mathbf{u}, \mathbf{v}, \mathbf{s})$. It generalizes the objective (4.2a).

Lemma 4.3.1. *If the stage costs are continuously differentiable and geodesically convex over a chosen domain, then a solution exists for (4.12).*

Proof. Recall that the subproblem is feasible if the feasibility set is non-empty. Due to the virtual control term v_k , any state in the feasible region of the convex subproblem is reachable in finite time. Furthermore, the virtual buffer zone variables s_k ensure that the feasible set defined by the linearized state and control constraints is non-empty. Since the perturbations are bounded by the trust-region constraints and the objective is continuous, an optimal solution exists. $\square$




Figure 4.1: Visualization of perturbing a trajectory along a tangent space.

4.3.6 Coordinate parameterization: Computation

Problem 4.12 is formulated intrinsically in terms of tangent vectors and linear operators between tangent spaces. To solve it numerically, we must choose local frames and express all quantities in coordinates. The optimization variables then become the coordinate representations of the tangent perturbations $\delta x_k \in T_{x_k}\mathcal{M}$ and $\delta u_k \in T_{u_k}\mathcal{U}$.

With a choice of frames along the reference trajectory, the coordinatization of Problem 4.12 becomes

$$\min_{[\delta \mathbf{x}], [\delta \mathbf{u}], [\mathbf{v}], \mathbf{s}} \hat{C}_{(\mathbf{x}, \mathbf{u})}^{(2)}([\delta \mathbf{x}], [\delta \mathbf{u}]) + \sum_{k=0}^{N-1} \lambda_k P([v_k], s_k) \quad (4.13a)$$

$$\text{s.t.} \quad [\delta x_{k+1}] = [R_{x_{k+1}}^{-1}(f(x_k, u_k))] + [\mathbf{D}_k]([\mathbf{A}_k][\delta x_k] + [\mathbf{B}_k][\delta u_k]) + [v_k], \quad (4.13b)$$

$$g(x_k, u_k) + [\mathbf{S}_k][\delta x_k] + [\mathbf{T}_k][\delta u_k] \leq s_k, \quad (4.13c)$$

$$s_k \geq 0, \quad [\delta x_k]^\top [g_{x_k}^{\mathcal{M}}][\delta x_k] \leq r^2, \quad [\delta u_k]^\top [g_{u_k}^{\mathcal{U}}][\delta u_k] \leq r^2. \quad (4.13d)$$

See §2.1.5 and §2.2.1 for a general discussion of covectors and tensor components. We collect here the coordinate representations of the intrinsic objects appearing in Problem 4.13. Throughout, let (E_i) be a frame on $\mathcal{M}$ and (F_i) a frame on $\mathcal{U}$.

Vector Components Any tangent vector $\eta \in T_x\mathcal{M}$ admits a coordinate representation $[\eta] = (\eta^1, \dots, \eta^n) \in \mathbb{R}^n$ such that

$$\eta = \sum_{i=1}^n \eta^i E_i|_x.$$

In the case that (E_i) is orthonormal, we conveniently have

$$\eta^i = g_x^{\mathcal{M}}(\eta, E_i|_x).$$

Covector (Differential) Components Let $f \in \mathfrak{F}(\mathcal{M})$. Its differential $df_x \in T_x^*\mathcal{M}$ has components represented by a row vector $[df_x] \in \mathbb{R}^{1 \times n}$ with entries

$$[df_x]_i := df_x(E_i|_x).$$

For $\eta \in T_x\mathcal{M}$, we can write

$$df_x(\eta) = [df_x][\eta].$$

Linear Operator Components The linear operators $\mathbf{A}_k$, $\mathbf{B}_k$, $\mathbf{S}_k$, $\mathbf{T}_k$, and $\mathbf{D}_k$ appearing in the i-SCvx subproblem are linear operators between tangent spaces. To describe their coordinate representations in a unified way, let $\mathcal{X}$ and $\mathcal{Y}$ be smooth manifolds of dimensions n and m , respectively. Fix points $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, and let

$$\mathbf{F} : T_x\mathcal{X} \rightarrow T_y\mathcal{Y}$$

be a linear operator. Let (E_i) be a local frame on $\mathcal{X}$ and let (F_i) be a local frame on $\mathcal{Y}$. Then $[\mathbf{F}] \in \mathbb{R}^{m \times n}$ is a matrix that satisfies, for all $\eta \in T_x\mathcal{X}$ and $\xi = \mathbf{F}(\eta) \in T_y\mathcal{Y}$,

$$[\mathbf{F}][\eta] = [\xi] \in \mathbb{R}^m.$$

The components of $[\mathbf{F}]$ can be computed column-wise as

$$[\mathbf{F}]_i := [\mathbf{F}(E_i|_x)] \in \mathbb{R}^m, \quad i = 1, \dots, n,$$

so that

$$[\mathbf{F}] = \begin{bmatrix} [\mathbf{F}]_1 & [\mathbf{F}]_2 & \cdots & [\mathbf{F}]_n \end{bmatrix} \in \mathbb{R}^{m \times n}.$$

In the case that (E_i) and (F_i) are orthonormal, we conveniently have

$$[\mathbf{F}]_{ij} = g_y^{\mathcal{Y}}(F_i|_y, \mathbf{F}(E_j|_x)).$$

Metric Components The Riemannian metric on $\mathcal{M}$ is represented in coordinates by a symmetric positive definite matrix $[g_x^{\mathcal{M}}] \in \text{Sym}_{++}(n)$ with entries

$$[g_x^{\mathcal{M}}]_{ij} := g_x^{\mathcal{M}}(E_i|_x, E_j|_x).$$

With this convention, the inner product of any two tangent vectors $\eta, \xi \in T_x \mathcal{M}$ is

$$g_x^{\mathcal{M}}(\eta, \xi) = [\eta]^\top [g_x^{\mathcal{M}}] [\xi].$$

The Riemannian metric on $\mathcal{U}$ is represented analogously by a matrix $[g_u^{\mathcal{U}}]$.

In the case that (E_i) is orthonormal, we conveniently have

$$[g_x^{\mathcal{M}}] = I_n.$$

Gradient and Hessian Components Associated with $f \in \mathfrak{F}(\mathcal{M})$ are its gradient and Hessian. The Riemannian gradient $\nabla f_x \in T_x \mathcal{M}$ is defined as the unique vector satisfying

$$g_x^{\mathcal{M}}(\nabla f_x, \eta) = \text{d}f_x(\eta) \quad \forall \eta \in T_x \mathcal{M}.$$

In coordinates, this relation yields

$$[\nabla f_x] = [g_x^{\mathcal{M}}]^{-1} [\text{d}f_x]^\top \in \mathbb{R}^{n \times 1}.$$

The bilinear Hessian $\text{Hess } f_x : T_x \mathcal{M} \times T_x \mathcal{M} \rightarrow \mathbb{R}$, associated with the Levi-Civita connection, admits a coordinate representation $[\text{Hess } f_x] \in \mathbb{R}^{n \times n}$ with entries

$$[\text{Hess } f_x]_{ij} := \text{Hess } f_x(E_i|_x, E_j|_x).$$

Accordingly, for any $\eta \in T_x \mathcal{M}$,

$$\text{Hess } f_x(\eta, \eta) = [\eta]^\top [\text{Hess } f_x] [\eta].$$

In the case that (E_i) is orthonormal, we conveniently have

$$[\nabla f_x] = [\text{d}f_x]^\top \in \mathbb{R}^{n \times 1}, \quad [\text{Hess } f_x] = [\nabla^2 f_x] \in \mathbb{R}^{n \times n}.$$

Taylor Expansions in Coordinates With these conventions, the Taylor expansions of $f \in \mathfrak{F}(\mathcal{M})$ at x admit simple coordinate expressions. The first-order expansion is

$$f_x^{(1)}(\eta) = f(x) + [\text{d}f_x][\eta] = f(x) + [\nabla f_x]^\top [g_x^{\mathcal{M}}][\eta],$$

and the second-order expansion is

$$f_x^{(2)}(\eta) = f_x^{(1)}(\eta) + \frac{1}{2}[\eta]^\top [\text{Hess } f_x] [\eta].$$

If f is geodesically convex on the chosen domain, then $[\text{Hess } f_x] \succeq 0$ and $f_x^{(2)}$ is a convex quadratic function of the coordinates $[\eta]$.

Applying this construction to the stage and terminal costs yields

$$\begin{aligned} \hat{\phi}_{(x,u)}^{(2)}([\eta], [\xi]) &= \phi(x, u) + [\text{d}_x \phi_{(x,u)}][\eta] + [\text{d}_u \phi_{(x,u)}][\xi] \\ &\quad + \frac{1}{2}[\eta]^\top [\text{Hess}_{xx} \phi_{(x,u)}][\eta] + \frac{1}{2}[\xi]^\top [\text{Hess}_{uu} \phi_{(x,u)}][\xi], \\ \hat{\phi}_f^{(2)}|_x([\eta]) &= \phi_f(x) + [\text{d}_x \phi_f|_x][\eta] + \frac{1}{2}[\eta]^\top [\text{Hess } \phi_f|_x][\eta]. \end{aligned}$$

4.3.7 The *i*-SCvx Algorithm

We present Algorithm 7 with the following remarks. First, it adheres to the SCvx protocol laid out in §4.2 and further detailed in [74, Alg. 2.1]. The quantities J and L are defined in (4.11) and (4.12a), respectively. Second, since Problem 4.12 optimizes over tangent vectors coordinatized by a chosen frame, the subproblem involves fewer optimization variables. This results in faster computation. Because the

perturbations δx_k , δu_k , v_k are tangent vectors, we are effectively optimizing over the necessary directions, reducing the magnitude and need to satisfy the implicit manifold constraint. Third, since Problem 4.12 is entirely written in intrinsic manifold quantities, the computational outcome is independent of parameterization.

Algorithm 7 i-SCvx Algorithm

- 1: **Input:** Choose initial trajectory $(\mathbf{x}, \mathbf{u}) \in \mathcal{M}^{N+1} \times \mathcal{U}^N$. Select tolerance ϵ_{tol} , trust region $r > 0$, penalty weight $\lambda > 0$, and parameters $0 \leq \rho_0 < \rho_1 < \rho_2 < 1$ and $0 < \alpha < 1 < \beta$, and sufficiently large $\Delta J > 0$.
 - 2: **while** $|\Delta J| > \epsilon_{\text{tol}}$ **do**
 - 3: Solve Problem (4.13) to get $[\delta \mathbf{x}], [\delta \mathbf{u}], [\mathbf{v}], \mathbf{s}$.
 - 4: Compute

$$\Delta J \leftarrow J(\mathbf{x}, \mathbf{u}) - J(R_{\mathbf{x}}([\delta \mathbf{x}]), R_{\mathbf{u}}([\delta \mathbf{u}]))$$

$$\Delta L \leftarrow J(\mathbf{x}, \mathbf{u}) - L([\delta \mathbf{x}], [\delta \mathbf{u}], [\mathbf{v}], \mathbf{s})$$
 - 5: Compute $\rho \leftarrow \Delta J / \Delta L$.
 - 6: **if** $\rho < \rho_0$ **then**
 - 7: $r \leftarrow \alpha r$
 - 8: **else**
 - 9: Update $\mathbf{x} \leftarrow R_{\mathbf{x}}([\delta \mathbf{x}])$, $\mathbf{u} \leftarrow R_{\mathbf{u}}([\delta \mathbf{u}])$
 - 10: Update $r \leftarrow \begin{cases} \alpha r, & \text{if } \rho < \rho_1; \\ r, & \text{if } \rho_1 \leq \rho < \rho_2; \\ \beta r, & \text{if } \rho \geq \rho_2. \end{cases}$
 - 11: **end if**
 - 12: **end while**
 - 13: **return** $\mathbf{x}, \mathbf{u}$
-

4.4 Example: Constrained Attitude Path-planning via i-SCvx

In this section, we compare i-SCvx with its Euclidean SCvx counterpart on a constrained attitude guidance problem.

Dynamics and Retraction The system dynamics are given by discrete-time quaternion kinematics

$$q_{k+1} = f(q_k, \omega_k) = q_k \odot \exp(\Delta t \omega_k),$$

where $q_k \in \mathcal{Q} \subset \mathbb{R}^4$ is a unit quaternion representing attitude, $\omega_k \in \mathbb{R}^3$ is the body-frame angular velocity, and $\Delta t > 0$ is the discretization step.

Perturbations are modeled intrinsically using a left-invariant retraction. Given a tangent vector $v \in T_q \mathcal{Q}$, we define

$$R_q(v) := q \odot \exp(\omega), \quad \omega := q^{-1} \odot v \in \mathbb{R}^3. \quad (4.14)$$

That is, perturbations are expressed in body coordinates and lifted to the manifold via left multiplication. This choice ensures that perturbations evolve consistently with the Lie group structure of $\mathcal{Q}$.

A global orthonormal frame for $\mathcal{Q}$ is given by

$$E_i(q) := q \odot e_i,$$

where $\{e_i\}_{i=1}^3$ is the standard basis of $\mathbb{R}^3$. While the resulting coordinate expressions depend on this frame, the underlying operators are intrinsic; choosing a different quaternion parameterization induces the same frame via left translation.

Keep-Out Constraint Let $y_B \in \mathbb{R}^3$ denote a fixed boresight direction expressed in the body frame, and let $t_I \in \mathbb{R}^3$ denote a fixed inertial-frame keep-out direction. The inertial-frame boresight direction is

$$y_I(q) := q \odot y_B \odot q^{-1}.$$

Given a maximum allowable exclusion angle $0 \leq \theta_{\max} < \pi$, the keep-out constraint is

$$g(q) := t_I^\top y_I(q) - \cos \theta_{\max},$$

so that $g(q) \leq 0$ enforces a minimum angular separation of $\theta_{\max}$ between t_I and the boresight.

For both SCvx and i-SCvx we employ the exact penalty

$$P(v, s) := \|q^{-1} \odot v\|_{1_q} + \|s\|_1,$$

where the norm on v is taken in body coordinates. This penalty cost involves measuring tangent vectors in body coordinates (left-trivialization), which is independent of the ambient embedding. For instance, using rotation matrices instead, while choosing the same equivalent global frame, results in the same penalty.

Euclidean SCvx Cost For Euclidean SCvx, we use the standard quadratic cost

$$\sum_{k=0}^{N-1} (\lambda_q \|q_k - q_f\|_2^2 + \lambda_\omega \|\omega_k\|_2^2) + \lambda_f \|q_N - q_f\|_2^2,$$

where $\lambda_q, \lambda_\omega, \lambda_f \geq 0$ are weighting coefficients. This cost is intentionally non-invariant: it depends on the ambient embedding $\mathcal{Q} \subset \mathbb{R}^4$, highlighting the parameterization dependence of Euclidean SCvx.

Geodesically Convex i-SCvx Cost For i-SCvx, we construct geodesically convex costs using the half squared geodesic distance to the target attitude:

$$\ell(q) := \frac{1}{2} \|\log(q_f^{-1} \odot q)\|_2^2.$$

This function is geodesically convex on the domain

$$D := \left\{ q \in \mathcal{Q} : \|\log(q_f^{-1} \odot q)\|_2 < \frac{\pi}{2} \right\},$$

see [83]. The stage and terminal costs are defined as

$$\begin{aligned}\phi(q, \omega) &:= \lambda_q \ell(q) + \lambda_\omega \|\omega\|_2^2, \\ \phi_f(q) &:= \lambda_f \ell(q),\end{aligned}$$

where $\lambda_q, \lambda_\omega, \lambda_f \geq 0$ are weighting coefficients.

Their differentials are

$$\begin{aligned}\mathrm{d}\phi_f|_q(v) &= \lambda_f \mathrm{d}\ell_q(v), & \mathrm{d}_q\phi_{(q,\omega)}(v) &= \lambda_q \mathrm{d}\ell_q(v), \\ \mathrm{d}_\omega\phi_{(q,\omega)}(\xi) &= 2\lambda_\omega \omega^\top \xi,\end{aligned}$$

and their bilinear Hessians satisfy

$$\begin{aligned}\mathrm{Hess}\phi_f|_q(v, v) &= \lambda_f \mathrm{Hess}\ell_q(v, v), \\ \mathrm{Hess}_{qq}\phi_{(q,\omega)}(v, v) &= \lambda_q \mathrm{Hess}\ell_q(v, v), \\ \mathrm{Hess}_{\omega\omega}\phi_{(q,\omega)}(\xi, \xi) &= 2\lambda_\omega \|\xi\|_2^2,\end{aligned}$$

with $\mathrm{Hess}_{q\omega}\phi = 0$.

The differential and bilinear Hessian of ℓ are obtained from those of the squared distance by scaling:

$$\begin{aligned}\mathrm{d}\ell_q(v) &= \frac{1}{2} \mathrm{d}(\|\log(q_f^{-1} \odot q)\|_2^2)_q(v) = v^\top (q \odot \log(q_f^{-1} \odot q)), \\ \mathrm{Hess}\ell_q(v, v) &= \frac{1}{2} \mathrm{Hess}(\|\log(q_f^{-1} \odot q)\|_2^2)_q(v, v) \\ &= v^\top \left(uu^\top + \theta \cot \theta (I - qq^\top - uu^\top) \right) v,\end{aligned}$$

where

$$\theta := \|\log(q_f^{-1} \odot q)\|_2, \quad u := \frac{(I - qq^\top)}{\sin \theta} q_f.$$

Linearized Dynamics and Constraints The intrinsic differentials of the dynamics and constraint are

$$\begin{aligned}\mathbf{A}_k(\delta x_k) &:= d_q f_{(q_k, \omega_k)}(\delta x_k) = \delta x_k \odot \exp(\Delta t \omega_k), \\ \mathbf{B}_k(\delta u_k) &:= d_\omega f_{(q_k, \omega_k)}(\delta u_k) = \Delta t q_k \odot d \exp_{\Delta t \omega_k}(\delta u_k), \\ \mathbf{S}_k(\delta x_k) &:= dg_{q_k}(\delta x_k) = t_I^\top (q_k \odot [q_k^{-1} \odot \delta x_k, y_B] \odot q_k^{-1}), \\ \mathbf{D}_k(v) &:= q_{k+1} \odot d \log_{q_{k+1}^{-1} \odot z_k}(q_{k+1}^{-1} \odot v),\end{aligned}$$

where $z_k := f(q_k, \omega_k)$ and $[\cdot, \cdot]$ denotes the Lie bracket on $\mathfrak{so}(3)$, which coincides with the vector cross product.

The derivatives of the quaternion exponential and logarithm used above admit the closed-form expressions [84–86]

$$\begin{aligned}d \exp_\omega(\eta) &= \begin{bmatrix} -\text{sinc}(\omega) \omega^\top \\ \text{sinc}(\omega) I_3 + \nabla \text{sinc}(\omega) \omega^\top \end{bmatrix} \eta, \\ d \log_q(v) &= \begin{bmatrix} -\|q_v\|_2^{-2} q_v & \alpha \|q_v\|_2^{-1} I_3 - \alpha \|q_v\|_2^{-3} q_v q_v^\top \end{bmatrix} v,\end{aligned}$$

where $\alpha := \text{atan2}(\|q_v\|_2, q_0)$.

4.4.1 Numerical Results

The numerical simulation¹ begins with randomly chosen initial q_0 and desired q_f attitudes. The initial trajectory was derived through spherical linear interpolation. We implemented both SCvx and i-SCvx with algorithm parameters $\epsilon_{\text{tol}} = 10^{-5}$, $r = 1$, $\alpha = .5$, $\beta = 3.2$, $\rho_0 = 0$, $\rho_1 = .25$, and $\rho_2 = .7$. The objective coefficients are $\lambda_q = 1$, $\lambda_f = 10$, $\lambda_\omega = .1$, and $\lambda = 10^5$. We repeated this process 100 times and recorded the average and standard deviation of the number of iterations for each algorithm, as well as the average clock time, geodesic, and Euclidean trajectory costs. We also stopped the algorithms if the iteration count exceeded $M = 100$. The corresponding results are presented in Tables 4.1 and 4.2.

¹Code available at github.com/Rainlabuw/intrinsic-scvx

In Table 4.1, i-SCvx outperforms SCvx in every category for $\theta_{\max} = 30^\circ$ and $\theta_{\max} = 10^\circ$. Paired with the fact that i-SCvx achieves significantly lower average iteration count, standard deviation of iteration counts, and wall clock time, this experiment shows i-SCvx has a natural advantage over SCvx for this particular optimal control problem. In Table 4.2, while we see SCvx achieving a lower Euclidean trajectory cost, the difference is small, especially considering i-SCvx has significantly more attractive results in the other categories. Figure 4.2 shows the optimal trajectories under both algorithms, and hence both costs. As shown, the trajectories are nearly identical.

Table 4.1:

	$\theta_{\max} = 10^\circ$		$\theta_{\max} = 30^\circ$	
	SCvx	i-SCvx	SCvx	i-SCvx
Avg. Iter.	40.21	24.89	45.8	26.8
Std. Iter.	9.23	2.14	18.28	1.88
Time (s)	6.26	4.40	7.10	4.70
Avg. Geo. Cost	4.98	4.73	6.50	5.67
Avg. Eucl. Cost	7.21	6.91	9.65	8.17

Results compare SCvx and i-SCvx for spacecraft attitude control under different angle constraints with $N = 30$ time steps and discretization rate $\Delta t = 0.1$ sec. Values are averages over multiple runs.

Table 4.2:

		$\theta_{\max} = 10^\circ$		$\theta_{\max} = 30^\circ$	
		SCvx	i-SCvx	SCvx	i-SCvx
Avg. Iter.		67.9	24.75	65.72	25.65
Std. Iter.		34.86	2.22	17.02	2.45
Time (s)		22.18	9.09	21.43	9.37
Avg. Geo.		9.04	8.96	10.59	10.50
Cost					
Avg. Eucl.		11.94	13.34	13.98	15.42
Cost					

$N = 60$ time steps and discretization rate $\Delta t = 0.05$ sec.




Figure 4.2: The iterates of SCvx and i-SCvx for the constrained attitude guidance problem with parameters $N = 30$, $\Delta t = 0.05$, and $\theta_{\max} = 20^\circ$.

4.5 Concluding Remarks

This chapter developed *intrinsic successive convexification* (i-SCvx), a parameterization-invariant extension of SCvx for trajectory optimization problems posed on Riemannian manifolds. The key idea is to localize the optimal control problem using perturbations that live in the appropriate tangent spaces, rather than in an ambient Euclidean embedding. This yields linearized dynamics and constraints expressed as linear operators between tangent spaces, and a localized objective obtained from second-order Taylor expansions of geodesically convex costs. As a result, the convex subproblem is posed in the minimal set of directions that are physically meaningful, and the algorithmic iterates are invariant under changes of coordinates and equivalent state parameterizations.

We illustrated the framework on constrained attitude guidance using unit quaternions. In this setting, i-SCvx admits a natural left-invariant modeling of perturbations and produces convex subproblems in three-dimensional tangent coordinates, avoiding the redundant variables and implicit manifold constraints that appear in Euclidean SCvx. Empirically, i-SCvx converged in fewer iterations with lower variance and reduced wall-clock time, while producing trajectories that were nearly indistinguishable from those obtained by Euclidean SCvx.

The intrinsic formulation also comes with limitations. First, SCvx can exploit *ambient* convexity by imposing convex constraints directly in the subproblem; on a manifold, geodesic convexity does not generally interact with inequality constraints in the same convenient way, so constraints often must be linearized to preserve convexity of the localized problem. Second, the present i-SCvx implementation does not yet incorporate several engineering features that are standard in the SCvx/SCP literature, including higher-order discretization (e.g., first-order holds), free-final-time formulations, fixed terminal constraints, and non-smooth objectives and constraints.

These shortcomings suggest natural directions for future work. On the numerical

side, one can generalize first-order hold discretizations and non-smooth constructions (via sub-gradients and exact penalties) to manifold settings by working systematically with retractions, differentials, and tangent-coordinate parameterizations. On the modeling side, extending i-SCvx to additional SCP variants and incorporating free-final-time and terminal constraint handling would substantially broaden its applicability. Finally, beyond attitude guidance, a compelling next step is to evaluate i-SCvx on higher-dimensional benchmark problems such as 6-DOF powered descent, where parameterization effects and redundant ambient constraints can significantly impact performance.

Chapter 5

ACTUAL AVERAGE CONSENSUS OVER LIE GROUPS

Whenever the whole is different from the sum of the parts—whenever there’s cooperation or competition going on—the governing equations must be nonlinear.

Steven Strogatz [87]

Consensus algorithms are a ubiquitous class of protocols with applications in distributed estimation, optimization, and control of multi-agent systems. At a basic level, consensus algorithms steer a collection of interacting agents towards a common state. In *average* consensus, this common state is the average of the agents’ initial conditions. Average consensus is attractive both statistically and computationally, and it underlies a wide range of applications including distributed resource allocation, formation control, and distributed estimation [87–89].

Most of the classical consensus literature is Euclidean. However, many problems in robotics, vision, and attitude synchronization evolve on curved state spaces, and this has motivated a substantial effort to generalize consensus protocols to Riemannian manifolds [83, 90–95]. A representative application is distributed 3D localization of camera sensors [96]. In this geometric setting, an appropriate analogue of the Euclidean average is the *Riemannian center of mass* (RCM), also known as the *Frechet mean* [97].

The manifold-valued consensus algorithms above typically guarantee agreement, but not *agreement at the Frechet mean*, which we call Riemannian average consensus. This stronger objective has been studied explicitly at least as early as [98, §3.2] and was later extended to broader geometric settings in [95, 99, 100]. A key feature of these approaches is the use of a consensus subroutine within each iteration to enforce

agreement before performing an averaging-type update. While effective, this nested structure can be undesirable in practice and may limit time-efficiency. This motivates a distributed Riemannian average consensus algorithm that is

1. intrinsic, so that its behavior does not depend on how the manifold is parameterized;
2. fully distributed, without relying on a consensus subroutine; and
3. fast, ideally achieving at least a linear convergence rate.

In this chapter, we reformulate Frechet mean computation as an optimization problem on a product manifold subject to a consensus constraint, connecting Riemannian average consensus to distributed optimization. We develop an intrinsic distributed solver on Lie groups equipped with bi-invariant Riemannian metrics. We characterize the stationary points of the resulting closed-loop dynamics and show that any stationary point corresponds to consensus at the Frechet mean of the data. While this result establishes correctness of stationary points, it does not provide a convergence guarantee to those stationary points. In the Euclidean case, we establish a global convergence guarantee and exponential convergence rate. Finally, we compare the proposed method against standard distributed constrained optimization baselines.

5.1 Problem Statement

Let $\mathcal{G}$ be a Lie group equipped with a bi-invariant Riemannian metric, and let $z_1, \dots, z_N \in \mathcal{G}$ be points admitting a unique Frechet mean $\bar{z}$. Consider N agents with states $x_i(t) \in \mathcal{G}$ evolving under a fixed, connected, undirected communication graph $\mathbf{G} = ([N], E)$. By this, we mean the dynamics of each agent must depend only on information from itself and its neighbors.

The objective of this work is to design distributed dynamics of the form above such that

$$\lim_{t \rightarrow \infty} x_i(t) = \bar{z}, \quad \forall i \in [N],$$

using only local communication and intrinsic geometric operations.

5.2 The Average: Riemannian Manifold Case

For detailed treatments of the Frechet mean on Riemannian manifolds, see [83, 97, 101–104].

Let $(\mathcal{M}, g)$ be a geodesically complete Riemannian manifold. For each $p \in \mathcal{M}$, let $\tilde{\mathcal{I}}_p \subset T_p\mathcal{M}$ denote the largest open set on which the Riemannian exponential map $\exp_p$ is a diffeomorphism. Its image

$$\mathcal{I}_p := \exp_p(\tilde{\mathcal{I}}_p) \subset \mathcal{M}$$

is the largest domain on which the Riemannian logarithm $\log_p : \mathcal{I}_p \rightarrow T_p\mathcal{M}$ is well-defined as the local inverse of $\exp_p$.

For $r > 0$, the (open) geodesic ball centered at p of radius r is

$$B(p, r) := \{q \in \mathcal{M} : d_g(p, q) < r\}.$$

The *injectivity radius* at p is

$$\text{inj}_p(\mathcal{M}) := \sup\{r > 0 : B(p, r) \subset \mathcal{I}_p\},$$

and we define the *global injectivity radius* by

$$\text{inj}(\mathcal{M}) := \inf_{p \in \mathcal{M}} \text{inj}_p(\mathcal{M}).$$

Let Δ be an upper bound on the sectional curvature of $\mathcal{M}$. We define the *convexity radius* by

$$r^*(\mathcal{M}) := \frac{1}{2} \min \left\{ \text{inj}(\mathcal{M}), \frac{\pi}{\sqrt{\Delta}} \right\}.$$

The terminology reflects the fact that every geodesic ball $B(p, r)$ with $r \leq r^*(\mathcal{M})$ is (strongly) geodesically convex.

Throughout this work we also use the product manifold

$$\mathcal{M}^N := \underbrace{\mathcal{M} \times \cdots \times \mathcal{M}}_{N \text{ factors}}$$

equipped with the product metric

$$g_{\mathbf{p}}(\mathbf{v}, \mathbf{w}) := \sum_{i=1}^N g_{p_i}(v_i, w_i), \quad \mathbf{p} = (p_1, \dots, p_N) \in \mathcal{M}^N, \quad \mathbf{v} = (v_1, \dots, v_N) \in T_{\mathbf{p}}\mathcal{M}^N.$$

The Frechet mean generalizes the empirical average in Euclidean space. Given $\mathbf{z} = (z_1, \dots, z_N) \in \mathcal{M}^N$, define the *Frechet function*

$$f(p) := \frac{1}{2} \sum_{i=1}^N d_g(p, z_i)^2, \quad p \in \mathcal{M}. \quad (5.1)$$

Any global minimizer of $f(\cdot)$ is called a *Frechet mean* of $\mathbf{z}$, i.e.,

$$\bar{z} \in \arg \min_{p \in \mathcal{M}} f(p). \quad (5.2)$$

In general, a Frechet mean need not exist nor be unique. We will denote $\text{RCM}(\mathbf{z})$ as the set of Frechet means of $\mathbf{z} \equiv \{z_k\}_{k=1}^N$.

Karcher Equation Whenever p lies in a region where all $\log_p(z_i)$ are defined, the gradient of the Frechet function satisfies

$$\nabla_p \left(\frac{1}{2} d_g(p, q)^2 \right) = -\log_p(q),$$

and hence

$$\nabla f_p = - \sum_{i=1}^N \log_p(z_i). \quad (5.3)$$

Therefore, any (local or global) minimizer $\bar{z}$ satisfies the *Karcher equation* $\sum_{i=1}^N \log_{\bar{z}}(z_i) = 0$.

Existence and Uniqueness Define the *convexity domain* (a subset of $\mathcal{M}^N$)

$$\mathcal{C} := \left\{ \mathbf{z} \in \mathcal{M}^N : \exists p \in \mathcal{M} \text{ such that } \{z_1, \dots, z_N\} \subset B(p, r^*(\mathcal{M})) \right\}. \quad (5.4)$$

If $\mathbf{z} \in \mathcal{C}$, then the Frechet function (5.1) admits a unique global minimizer (hence a unique Frechet mean) [97, Thm. 2.1]. With this in mind, RCM becomes a well-defined function:

$$\bar{z} = \text{RCM}(\mathbf{z}), \quad \text{RCM} : \mathcal{C} \subset \mathcal{M}^N \rightarrow \mathcal{M}.$$

Computing the Frechet Mean by Gradient Descent A Frechet mean can be computed via Riemannian gradient descent applied to (5.2). Under the strong geodesic convexity enjoyed on geodesic balls of radius at most $r^*(\mathcal{M})$, a fixed step size yields linear convergence; see [105] for step-size choices and rates.

Algorithm 8 Frechet mean subroutine (Riemannian gradient descent)

Require: $z_1, \dots, z_N \in \mathcal{M}$, step size $\epsilon > 0$, tolerance $\tau > 0$

- 1: Initialize $p \leftarrow z_1$
 - 2: **while** $\|\nabla f_p\|_p > \tau$ **do**
 - 3: Compute $\nabla f_p = -\sum_{i=1}^N \log_p(z_i)$
 - 4: Update $p \leftarrow \exp_p(-\epsilon \nabla f_p)$
 - 5: **end while**
 - 6: **return** p
-

Characterization in a Convex Ball Let $\mathbf{z} \in \mathcal{C}$, and let $B \subset \mathcal{M}$ be a geodesic ball of radius $r^*(\mathcal{M})$ containing $\{z_1, \dots, z_N\}$. Then the Frechet mean $\bar{z}$ is the unique point in B satisfying the Karcher equation (5.3).

5.3 Consensus Reformulation and Distributed Optimization

We now reformulate the Frechet mean problem (5.2) as a consensus-constrained optimization problem over a network of agents. This reformulation makes explicit the connection between Frechet mean computation and distributed consensus.

Consider the lifted objective

$$F(\mathbf{x}) := \frac{1}{2} \sum_{i=1}^N d_g(x_i, z_i)^2, \quad \mathbf{x} = (x_1, \dots, x_N) \in \mathcal{G}^N,$$

defined on the product manifold $\mathcal{G}^N$. Let

$$\mathcal{D} := \{(x_1, \dots, x_N) \in \mathcal{G}^N : x_1 = \dots = x_N\}$$

denote the *diagonal set*. Restricting F to $\mathcal{D}$ recovers the original Frechet function:

$$F(x, \dots, x) = f(x).$$

The Frechet mean problem can therefore be written equivalently as the constrained optimization problem

$$\min_{\mathbf{x} \in \mathcal{G}^N} F(\mathbf{x}) \tag{5.5a}$$

$$\text{s.t. } \mathbf{x} \in \mathcal{D}. \tag{5.5b}$$

If the data $\mathbf{z}$ lie in the convexity domain $\mathcal{C}$, then the Frechet mean $\bar{z}$ is unique, and (5.5) admits the unique solution

$$\mathbf{x}^* = (\bar{z}, \dots, \bar{z}) \in \mathcal{D}.$$

This formulation reveals that computing the Frechet mean is equivalent to minimizing a separable cost subject to a geometric consensus constraint. In the next sections, we exploit this structure to design distributed algorithms that enforce agreement while descending the Frechet objective using only local information.

5.4 Distributed Optimization: Euclidean Case

We derive the proposed distributed algorithm by first reviewing consensus optimization in Euclidean space, explaining why naive distributed gradient descent fails, and introducing gradient tracking as a corrective mechanism. The overall method is called distributed gradient descent with gradient tracking (DGD-GT). See [106–109]. This construction serves as a template for the intrinsic formulation developed later on Lie groups equipped with bi-invariant Riemannian metrics.

Consensus Optimization in Euclidean Space Consider the centralized optimization problem

$$\min_{x \in \mathbb{R}^n} f(x) := \sum_{i=1}^N f_i(x),$$

where each $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and strongly convex. The aggregate cost f is therefore strongly convex and admits a unique global minimizer x^* .

Suppose the cost is distributed across a network of N agents connected by a fixed, undirected, and connected communication graph $\mathbf{G} = ([N], E)$. Each agent i has access only to its local cost f_i and the states of its neighbors. The distributed objective is to compute x^* using only local communication.

This problem can be equivalently written as the consensus-constrained optimization problem

$$\min_{\mathbf{x}=(x_1, \dots, x_N)} F(\mathbf{x}) := \sum_{i=1}^N f_i(x_i), \quad (5.6a)$$

$$\text{s.t. } x_1 = \dots = x_N, \quad (5.6b)$$

where the objective is separable but the constraint couples the agents. When $\mathbf{x} = (x, \dots, x)$, we recover $F(\mathbf{x}) = f(x)$ and $\nabla f(x) = \sum_{i=1}^N \nabla f_i(x)$. Although (5.6) is convex, the consensus constraint precludes a naive decomposition into independent local subproblems.

Failure of Distributed Gradient Descent A standard distributed approach is continuous-time distributed gradient descent (DGD),

$$\dot{x}_i = \sum_{j \in \mathcal{N}_i} (x_j - x_i) - \nabla f_i|_{x_i}, \quad (5.7)$$

which combines a consensus term with local gradient descent. Let L be the graph Laplacian of $\mathbf{G}$ and let $\mathbf{L} = L \otimes I_n$. Identifying $(\mathbb{R}^n)^N$ with $\mathbb{R}^{Nn}$, (5.7) can be written compactly as

$$\dot{\mathbf{x}} = -\mathbf{L}\mathbf{x} - \nabla F_{\mathbf{x}}.$$

At a stationary point, the dynamics satisfy

$$\mathbf{L}\mathbf{x} = -\nabla F_{\mathbf{x}}.$$

In general, this condition neither enforces consensus nor ensures $\sum_{i=1}^N \nabla f_i|_{x_i} = 0$. Consequently, the stationary points of DGD do not coincide with the solutions of (5.6), even when each f_i is strongly convex.

Ideal Consensus Dynamics If each agent had access to the global average gradient $\frac{1}{N} \sum_{j=1}^N \nabla f_j|_{x_j}$, one could instead consider the ideal dynamics

$$\dot{x}_i = \sum_{j \in \mathcal{N}_i} (x_j - x_i) - \frac{1}{N} \sum_{j=1}^N \nabla f_j|_{x_j}. \quad (5.8)$$

Theorem 5.4.1. *Assume the communication graph is undirected and connected. Any stationary point of the dynamics (5.8) satisfies $x_1 = \dots = x_N \equiv x^*$, where x^* is the unique minimizer of f .*

The dynamics (5.8) therefore solve the consensus optimization problem exactly. However, they are not implementable in a distributed setting, as the global average gradient is unavailable to individual agents.

Gradient Tracking Gradient tracking addresses this limitation by allowing each agent to asymptotically reconstruct the global average gradient using only local communication. Introduce auxiliary variables $w_i \in \mathbb{R}^n$ with dynamics

$$\begin{aligned}\dot{w}_i &= \sum_{j \in \mathcal{N}_i} (w_j - w_i + \nabla f_i|_{x_i} - \nabla f_j|_{x_j}), \\ y_i &= -w_i + \nabla f_i|_{x_i}.\end{aligned}$$

In stacked form,

$$\dot{\mathbf{w}} = -\mathbf{L}\mathbf{w} + \mathbf{L}\nabla F_{\mathbf{x}}, \quad (5.9)$$

$$\mathbf{y} = -\mathbf{w} + \nabla F_{\mathbf{x}}, \quad (5.10)$$

with $\mathbf{w}(0) = 0$. The variables y_i asymptotically track the network-wide average gradient $\frac{1}{N} \sum_{j=1}^N \nabla f_j|_{x_j}$, even when the inputs vary with time.

Remark 5.4.1. [88, 89, 110] *The dynamics (5.9) are an instance of dynamic average consensus. For a connected graph with Laplacian L , consider the linear system*

$$\begin{aligned}\dot{x} &= -Lx + Lu, \\ y &= -x + u,\end{aligned}$$

with $x(0) = 0$ and time-varying input $u(t) \in \mathbb{R}^N$. This system preserves averages and asymptotically synchronizes the outputs $y_i(t)$ to the average of the inputs whenever $u(t)$ converges.

Distributed Gradient Descent with Gradient Tracking Replacing the global average gradient in (5.8) with the tracked variables y_i yields the fully distributed algorithm

$$\begin{aligned}\dot{x}_i &= \sum_{j \in \mathcal{N}_i} (x_j - x_i) - y_i, \\ \dot{w}_i &= \sum_{j \in \mathcal{N}_i} (w_j - w_i + \nabla f_i|_{x_i} - \nabla f_j|_{x_j}), \\ y_i &= -w_i + \nabla f_i|_{x_i}.\end{aligned}$$

In stacked form,

$$\begin{aligned}\dot{\mathbf{x}} &= -\mathbf{L}\mathbf{x} - \mathbf{y}, \\ \dot{\mathbf{w}} &= -\mathbf{L}\mathbf{w} + \mathbf{L}\nabla F_{\mathbf{x}}, \\ \mathbf{y} &= -\mathbf{w} + \nabla F_{\mathbf{x}}.\end{aligned}$$

Any stationary point of this system corresponds to consensus at the global minimizer of f . Under standard smoothness assumptions and graph connectivity, the trajectories converge to this stationary point.

5.5 Distributed Optimization: Lie Group Case

Here we will discuss how DGD-GT can be generalized to Lie groups with bi-invariant metrics. Papers introducing a DGD-GT-like algorithm on non-Euclidean spaces include [100, 111–113]. See §A for a necessary background on Lie theory for this work.

The reason we work with Lie groups $\mathcal{G}$ equipped with bi-invariant Riemannian metrics g is because of the following. The Lie and manifold exponential and logarithm operators coincide. By this, we mean

$$\begin{aligned}\exp_x(\mathrm{d}L_p\xi) &= p \cdot \exp(\xi) \\ \log_p(q) &= \mathrm{d}L_p \log(p^{-1} \cdot q)\end{aligned}$$

This has an important consequence. If we define the consensus error for N points $\mathbf{x} = (x_1, \dots, x_N) \subset \mathcal{G}$ under a connected graph $\mathbf{G} = ([N], E)$ as

$$\varphi(x_1, \dots, x_N) := \frac{1}{4} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} d(x_i, x_j)^2,$$

then its gradient along the i th argument becomes

$$\nabla_i \varphi_{\mathbf{x}} = - \sum_{j \in \mathcal{N}_i} \log_{x_i}(x_j) \implies \mathrm{d}L_{x_i^{-1}} \nabla_i \varphi_{\mathbf{x}} = - \sum_{j \in \mathcal{N}_i} \log(x_i^{-1} \cdot x_j).$$

This implies

$$\sum_{i=1}^N \mathrm{d}L_{x_i^{-1}} \nabla_i \varphi_{\mathbf{x}} = - \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} \log(x_i^{-1} \cdot x_j) = 0,$$

where we are using the anti-symmetric property of the Lie logarithm¹:

$$\log(p^{-1} \cdot q) = -\log(q^{-1} \cdot p)$$

In conclusion, we can write

$$\mathrm{d}L_{\mathbf{x}^{-1}} \nabla \varphi_{\mathbf{x}} \in \ker(\mathbf{L}) \subset \mathfrak{g}^N,$$

¹As a side note, the anti-symmetry property holds for all Lie groups. But the relationship with the Riemannian gradient only holds in the bi-invariant case.

where L is the graph Laplacian of $\mathbf{G}$ and $\mathbf{L} : \mathfrak{g}^N \rightarrow \mathfrak{g}^N$ is the graph Laplacian operator augmented onto the vector space $\mathfrak{g}$.

It is this fact that allows us to characterize the stationary points of DGD-GT generalized to bi-invariant Lie groups, as we will show below.

Lifting To Lie Groups We now lift the construction introduced in §5.4 to a bi-invariant Lie group $\mathcal{G}$. Let

$$f_i(x) := \frac{1}{2}d_g(x, z_i)^2, \quad f(x) := \sum_{i=1}^N f_i(x),$$

and define the distributed cost $F(\mathbf{x}) := \sum_{i=1}^N f_i(x_i)$ on $\mathcal{G}^N$. Remark

$$\nabla f_i|_x = -\log_x(z_i), \quad \nabla_i F_{\mathbf{x}} = \nabla f_i|_{x_i}.$$

Consensus is measured by the potential

$$\varphi(\mathbf{x}) := \frac{1}{2} \sum_{\{i,j\} \in E} d_g(x_i, x_j)^2. \quad (5.11)$$

Remark

$$\nabla_i \varphi_{\mathbf{x}} = - \sum_{j \in \mathcal{N}_i} \log_{x_i}(x_j).$$

The natural Riemannian analogue of DGD is

$$\dot{\mathbf{x}} = -\nabla \varphi_{\mathbf{x}} - \nabla F_{\mathbf{x}}, \quad (5.12)$$

which again fails to solve the consensus-constrained problem.

Gradient Tracking On The Lie Algebra To correct this, gradient tracking is implemented on the Lie algebra $\mathfrak{g}$. Define the body-frame gradients

$$\gamma_i := dL_{x_i^{-1}} \nabla_i F_{\mathbf{x}} = -\log(x_i^{-1} \cdot z_i).$$

The gradient tracking dynamics are

$$\dot{\alpha}_i = \sum_{j \in \mathcal{N}_i} (\alpha_j - \alpha_i + \gamma_i - \gamma_j), \quad (5.13a)$$

$$\omega_i = -\alpha_i + \gamma_i, \quad (5.13b)$$

so that each ω_i tracks the average body-frame gradient.

Replacing $-\nabla F_{\mathbf{x}}$ in (5.12) with the tracked signal $\boldsymbol{\omega}$ yields the distributed solver presented in Algorithm 5.14, thereby closing the loop between consensus and optimization. The resulting dynamics are

$$\dot{x}_i = \mathrm{d}L_{x_i} \left(\sum_{j \in \mathcal{N}_i} \log(x_i^{-1} \cdot x_j) - \omega_i \right), \quad (5.14)$$

$$\dot{\alpha}_i = \sum_{j \in \mathcal{N}_i} (\alpha_j - \alpha_i + \gamma_i - \gamma_j) \quad (5.15)$$

where

$$\gamma_i = -\log(x_i^{-1} \cdot z_i),$$

$$\omega_i = -\alpha_i + \gamma_i,$$

and $\alpha_i(0) = 0$.

Remark 5.5.1. *The signals z_i need not be constant. If the data are time-varying, the same gradient-tracking mechanism allows the agents to track the time-varying Frechet mean, as confirmed empirically.*

Remark 5.5.2. *The initialization $\mathbf{x}(0) = \mathbf{z}$ can be relaxed to any $\mathbf{x}(0) \in \mathcal{C}$. Closer initialization to the Frechet mean improves convergence speed, but since the mean is unknown a priori, $\mathbf{x}_i(0) = z_i$ is a natural choice.*

5.6 Main Results

In this section, we first provide the limit point analysis of Algorithm 5.14 in the general case and then guarantee convergence for the Euclidean case.

Proposition 5.6.1. *Let $\mathcal{G}$ be a Lie group equipped with a bi-invariant Riemannian metric, and let $r^* > 0$ denote its convexity radius. Let $\mathcal{B} \subset \mathcal{G}$ be a geodesic ball of radius $r < r^*$. Pick $z_1, \dots, z_N \in \mathcal{B}$. Suppose the agents are initialized with $x_i(0) \in \mathcal{B}$ and $\alpha_i(0) = 0$ for all i , and assume that the trajectory generated by (5.14) converges to a stationary point $(\mathbf{x}^*, \boldsymbol{\alpha}^*)$ with $\mathbf{x}^* \in \mathcal{B}^N$.*

Then the agents reach consensus, i.e.,

$$x_1^* = \dots = x_N^* =: x^*,$$

and the consensus point coincides with the unique Frechet mean $\bar{z}$ of $\{z_1, \dots, z_N\}$.

Moreover,

$$\alpha_i^* = \log(z_i^{-1} \cdot \bar{z}), \quad i = 1, \dots, N.$$

Proof. For $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{G}^N$, define the stacked logarithm

$$\log(\mathbf{z}^{-1} \cdot \mathbf{x}) := (\log(z_1^{-1} \cdot x_1), \dots, \log(z_N^{-1} \cdot x_N)) \in \mathfrak{g}^N,$$

and write $dL_{\mathbf{x}} : \mathfrak{g}^N \rightarrow T_{\mathbf{x}}\mathcal{G}^N$ for the product left-translation differential, i.e., $dL_{\mathbf{x}}(\eta_1, \dots, \eta_N) = (dL_{x_1}\eta_1, \dots, dL_{x_N}\eta_N)$. With this notation, Algorithm 5.14 can be written compactly

as

$$\begin{pmatrix} \dot{\mathbf{x}} \\ \dot{\boldsymbol{\alpha}} \end{pmatrix} = \begin{pmatrix} -\nabla\varphi_{\mathbf{x}} + dL_{\mathbf{x}}[\boldsymbol{\alpha} - \log(\mathbf{z}^{-1} \cdot \mathbf{x})] \\ \mathbf{L}[-\boldsymbol{\alpha} + \log(\mathbf{z}^{-1} \cdot \mathbf{x})] \end{pmatrix}.$$

Since $(\mathbf{x}^*, \boldsymbol{\alpha}^*)$ is a stationary point, we have

$$\mathbf{L}[-\boldsymbol{\alpha}^* + \log(\mathbf{z}^{-1} \cdot \mathbf{x}^*)] = 0. \quad (5.16)$$

Because the communication graph is connected, $\ker(\mathbf{L}) = \text{span}\{\mathbf{1}\}$, and hence there exists $\xi \in \mathfrak{g}$ such that

$$-\boldsymbol{\alpha}^* + \log(\mathbf{z}^{-1} \cdot \mathbf{x}^*) = (\xi, \dots, \xi) =: \xi \mathbf{1}. \quad (5.17)$$

The stationary condition for $\dot{\mathbf{x}}$ gives

$$\nabla \varphi_{\mathbf{x}^*} = dL_{\mathbf{x}^*}[\boldsymbol{\alpha}^* - \log(\mathbf{z}^{-1} \cdot \mathbf{x}^*)].$$

Applying $dL_{(\mathbf{x}^*)^{-1}}$ to both sides and using (5.17) yields

$$dL_{(\mathbf{x}^*)^{-1}} \nabla \varphi_{\mathbf{x}^*} = -\xi \mathbf{1}. \quad (5.18)$$

Taking inner products gives

$$N \|\xi\|^2 = \langle \xi \mathbf{1}, \xi \mathbf{1} \rangle = -\langle dL_{(\mathbf{x}^*)^{-1}} \nabla \varphi_{\mathbf{x}^*}, \xi \mathbf{1} \rangle.$$

Next, we compute the left-trivialized consensus gradient. By bi-invariance,

$$\nabla_i \varphi_{\mathbf{x}} = - \sum_{j \in \mathcal{N}_i} \log_{x_i}(x_j) = -dL_{x_i} \sum_{j \in \mathcal{N}_i} \log(x_i^{-1} \cdot x_j),$$

hence

$$dL_{x_i^{-1}} \nabla_i \varphi_{\mathbf{x}} = - \sum_{j \in \mathcal{N}_i} \log(x_i^{-1} \cdot x_j).$$

Therefore,

$$\langle dL_{(\mathbf{x}^*)^{-1}} \nabla \varphi_{\mathbf{x}^*}, \xi \mathbf{1} \rangle = - \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} \langle \log((x_i^*)^{-1} \cdot x_j^*), \xi \rangle.$$

Using the antisymmetry identity $\log(a^{-1} \cdot b) = -\log(b^{-1} \cdot a)$ (whenever defined) and pairing terms edge-by-edge shows the double sum cancels:

$$\sum_{i=1}^N \sum_{j \in \mathcal{N}_i} \log((x_i^*)^{-1} \cdot x_j^*) = \sum_{\{i,j\} \in E} \left(\log((x_i^*)^{-1} \cdot x_j^*) + \log((x_j^*)^{-1} \cdot x_i^*) \right) = 0.$$

Consequently,

$$\langle dL_{(\mathbf{x}^*)^{-1}} \nabla \varphi_{\mathbf{x}^*}, \xi \mathbf{1} \rangle = 0,$$

and hence $\xi = 0$. It follows from (5.18) that

$$\nabla \varphi_{\mathbf{x}^*} = 0.$$

By [83, Theorem 5], the critical points of φ restricted to $\mathcal{C}$ coincide with the diagonal set $\mathcal{D}$. Since $\mathbf{x}^* \in \mathcal{B}^N \subset \mathcal{C}$ and $\nabla \varphi_{\mathbf{x}^*} = 0$, we conclude that $\mathbf{x}^* \in \mathcal{D}$, i.e.,

$$x_1^* = \cdots = x_N^* =: x^*.$$

Next, observe that $\frac{d}{dt} \sum_{i=1}^N \alpha_i(t) = 0$. Since $\alpha_i(0) = 0$, it follows that $\sum_{i=1}^N \alpha_i^* = 0$. Combining this with (5.17) and $\xi = 0$ yields

$$\sum_{i=1}^N \log(z_i^{-1} \cdot x^*) = 0.$$

Because $x^*, z_1, \dots, z_N \in \mathcal{B}$ and $\mathcal{B}$ lies within the convexity radius, this condition uniquely characterizes the Frechet mean, and therefore $x^* = \bar{z}$. The expression for α_i^* follows immediately. $\square$

Convergence Proof: Euclidean Case Next, we provide the convergence guarantees of our algorithm for $\mathcal{G} = \mathbb{R}^n$; the corresponding analysis for arbitrary Lie groups is the subject of our future work. Herein, we write $\mathbf{x} = \text{vec}(x_1, \dots, x_N) \in \mathbb{R}^{Nn}$ to denote vertical concatenation and use similar notation for $\mathbf{z}$, $\mathbf{x}$, and $\boldsymbol{\alpha}$. Then the dynamics of Algorithm 5.14 reduces to the following matrix form:

$$\dot{\mathbf{x}} = -(L \otimes I_n)\mathbf{x} + \boldsymbol{\alpha} - (\mathbf{x} - \mathbf{z}) \quad (5.19a)$$

$$\dot{\boldsymbol{\alpha}} = -(L \otimes I_n)\boldsymbol{\alpha} + (L \otimes I_n)(\mathbf{x} - \mathbf{z}), \quad (5.19b)$$

where $\otimes$ denote the Kronecker product. The RCM of initial states $\mathbf{z}$ reduces to the Euclidean average

$$\bar{z} = \text{RCM}(\mathbf{z}) = \frac{1}{N} \sum_{i=1}^N z_i \in \mathbb{R}^n,$$

and so $\mathbf{x}^* = \text{vec}(\bar{z}, \dots, \bar{z}) = \mathbf{1}_N \otimes \bar{z}$. The next result establishes guaranteed average consensus starting from arbitrary initial points.

Theorem 5.6.1 (Convergence in $\mathbb{R}^n$). *Suppose $\mathbf{x}(t), \boldsymbol{\alpha}(t)$ is the trajectory generated by (5.19) over a connected graph $\mathbf{G}$ with $\boldsymbol{\alpha}(0) = 0$. Then, starting from any arbitrary*

starting point $\mathbf{x}(0)$, we have $\lim_{t \rightarrow \infty} x_i(t) = \bar{z}$ and $\lim_{t \rightarrow \infty} \alpha_i(t) = \bar{z} - z_i$ for all i with a linear convergence rate.

Proof. First, by noting that $\mathbf{v} = -\boldsymbol{\alpha} + \mathbf{x} - \mathbf{z}$ we reformulate the system in (5.19) as

$$\begin{bmatrix} \dot{\mathbf{v}} \\ \dot{\mathbf{x}} \end{bmatrix} = (A \otimes I_n) \begin{bmatrix} \mathbf{v} \\ \mathbf{x} \end{bmatrix}, \quad (5.20)$$

where

$$A := \begin{bmatrix} -\mathbf{L} - I_N & -\mathbf{L} \\ -I_N & -\mathbf{L} \end{bmatrix}. \quad (5.21)$$

Before proceeding, we need the following results on characterizing the spectrum of A with its proof deferred to the end of this section.

Lemma 5.6.1. *Suppose $\mathbf{G}$ is connected. Then zero is a simple eigenvalue of the matrix A in (5.21). Furthermore, all non-zero eigenvalues of this matrix have negative real parts.*

Lemma 5.6.1 implies that the linear system $\dot{\xi} = (A \otimes I_n)\xi$ is marginally stable: it has an n -dimensional eigenspace at $\lambda = 0$ (corresponding to the consensus directions), while all remaining modes decay exponentially. Let p, q be, respectively, the right and left eigenvectors of A associated with $\lambda = 0$ such that $p^\top q = 1$. So, $p \in \mathcal{N}(A)$. Also, since 0 is a simple eigenvalue of $A^\top$, we get

$$\mathcal{N}(A^\top) = \text{span} \left(\begin{bmatrix} 1_N \\ -1_N \end{bmatrix} \right).$$

Thus, we set $p = \begin{bmatrix} 0_N \\ 1_N \end{bmatrix}$, $q = \frac{1}{N} \begin{bmatrix} -1_N \\ 1_N \end{bmatrix}$. By properties of Kronecker product, we obtain that $A \otimes I_n$ is also marginally stable with a zero eigenvalue of algebraic and geometric multiplicity n . This eigenvalue has corresponding right eigenvectors $p \otimes \mathbf{e}_i$ and left eigenvectors $q \otimes \mathbf{e}_i$. Here, $\mathbf{e}_i \in \mathbb{R}^n$ is the i -th standard basis for $i = 1, 2, \dots, n$.

Next, by computing $e^{(A \otimes I_n)t}$ using the Jordan decomposition and taking the limit as $t \rightarrow \infty$ we obtain that (see [114, Proposition 3.11] for a similar computation)

$$\lim_{t \rightarrow \infty} e^{(A \otimes I_n)t} = \sum_{i=1}^n (p \otimes \mathbf{e}_i)(q \otimes \mathbf{e}_i)^\top = (pq^\top) \otimes I_n.$$

Therefore, 5.21 converges as follows

$$\begin{bmatrix} \mathbf{v}^* \\ \mathbf{x}^* \end{bmatrix} := \lim_{t \rightarrow \infty} \begin{bmatrix} \mathbf{v}(t) \\ \mathbf{x}(t) \end{bmatrix} = ((pq^\top) \otimes I_n) \begin{bmatrix} \mathbf{v}(0) \\ \mathbf{x}(0) \end{bmatrix},$$

where $\boldsymbol{\alpha}(0) = 0_{Nn}$ and $\mathbf{x}(0)$ is the arbitrary starting point. Thus, $\mathbf{v}(0) = -\boldsymbol{\alpha}(0) + \mathbf{x}(0) - \mathbf{z} = \mathbf{x}(0) - \mathbf{z}$, and

$$\begin{aligned} \begin{bmatrix} \mathbf{v}^* \\ \mathbf{x}^* \end{bmatrix} &= \frac{1}{N} \left(\begin{bmatrix} 0_{N \times N} & 0_{N \times N} \\ -1_N 1_N^\top & 1_N 1_N^\top \end{bmatrix} \otimes I_n \right) \begin{bmatrix} \mathbf{x}(0) - \mathbf{z} \\ \mathbf{x}(0) \end{bmatrix} \\ &= \begin{bmatrix} 0_{Nn} \\ \frac{1}{N}(1_N 1_N^\top \otimes I_n) \mathbf{z} \end{bmatrix} = \begin{bmatrix} 0_{Nn} \\ 1_N \otimes \bar{\mathbf{z}} \end{bmatrix}. \end{aligned}$$

Therefore $\mathbf{x}^* = 1_N \otimes \bar{\mathbf{z}}$ which completes the proof. $\square$

Proof of 5.6.1. Let $\begin{bmatrix} \mathbf{v} \\ \mathbf{x} \end{bmatrix} \in \mathcal{N}(A)$ be non-zero. Then

$$-\mathbf{L}\mathbf{v} - \mathbf{v} - \mathbf{L}\mathbf{x} = 0, \tag{5.22a}$$

$$-\mathbf{v} - \mathbf{L}\mathbf{x} = 0. \tag{5.22b}$$

By combining the two, we get $\mathbf{L}\mathbf{v} = 0$ implying that $\mathbf{v} \in \text{span}(1_N)$. It follows from (5.22b) that $\mathbf{L}\mathbf{x} = -\mathbf{v} = \alpha 1_N$. Then $1^\top \mathbf{L}\mathbf{x} = 0 = \alpha 1_N^\top 1_N$, and thus $\alpha = 0$. But this implies that $\mathbf{v} = 0$. Thus $\mathbf{L}\mathbf{x} = 0$, which implies that $\mathbf{x} \in \text{span}(1_N)$. Therefore, we can conclude that

$$\mathcal{N}(A) = \left\{ \begin{bmatrix} 0_N \\ \mathbf{x} \end{bmatrix} : \mathbf{x} \in \text{span}(1_N) \right\}.$$

Since the nullspace has dimension 1, it follows that 0 is a simple eigenvalue of A .

Next, let $\begin{bmatrix} \mathbf{v} \\ \mathbf{x} \end{bmatrix}$ be an eigenvector of A associated with a non-zero eigenvalue λ . For the sake of contradiction, suppose $\Re[\lambda] \geq 0$. Then

$$\begin{cases} \lambda \mathbf{v} &= -\mathbf{L}\mathbf{v} - \mathbf{v} - \mathbf{L}\mathbf{x} \\ \lambda \mathbf{x} &= -\mathbf{v} - \mathbf{L}\mathbf{x} \end{cases}$$

Note, if $\mathbf{x} = 0$, then $\mathbf{v} = 0$. So, in order for $\begin{bmatrix} \mathbf{v} \\ \mathbf{x} \end{bmatrix}$ to be an eigenvector, we need $\mathbf{x} \neq 0$.

By solving for $\mathbf{x}$, we obtain

$$-\lambda \mathbf{x} = (\lambda I + \mathbf{L})^2 \mathbf{x}. \quad (5.23)$$

So, $(-\lambda, \mathbf{x})$ is an eigenpair of the matrix $(\lambda I + \mathbf{L})^2$. The eigenvalues of $(\lambda I + \mathbf{L})^2$ are $\{(\lambda + \rho)^2 : \rho \in \sigma(\mathbf{L})\}$. So, there exists some $\rho \in \sigma(\mathbf{L})$ such that $-\lambda = (\lambda + \rho)^2$. Let $\lambda = a + bi$ and note the hypothesis of contradiction is $a \geq 0$ and $a^2 + b^2 \neq 0$. By separating the real and imaginary parts, we obtain

$$\begin{aligned} -a &= (a + \rho)^2 - b^2, \\ -b &= 2(a + \rho)b. \end{aligned}$$

Suppose $b = 0$, then the first equation and the hypothesis $a \geq 0$ implies that $a = 0$. But, this implies $\lambda = 0$ which is a contradiction. Next, suppose $b \neq 0$, then the second equation implies $-1 = 2a + 2\rho$. Thus $a = \frac{-1-2\rho}{2} < 0$, which is a contradiction with $a \geq 0$ as $\rho \geq 0$. Therefore, if $\lambda \neq 0$ then we must have $a < 0$, and so all non-zero eigenvalues of A have negative real part. This completes the proof. $\square$

5.7 Simulations

In this section and for illustration purposes, we consider the special orthogonal Lie group $\mathcal{G} = \text{SO}(3)$, with its parameterization of 3×3 rotation matrices. The associated Lie algebra, denoted $\mathfrak{so}(3)$, is the space of 3×3 skew-symmetric matrices. We equip $\text{SO}(3)$ with metric induced by the Frobenius inner product. The corresponding geodesic distance reduces to

$$d(R, S)^2 = \|\log(R^\top S)\|^2 = -\frac{1}{2}\text{tr}[\log(R^\top S)^2] .$$

In the following, two scenarios are in order where we randomly initialize agents on $\text{SO}(3)$.

5.7.1 Scenario 1 (*Comparison of Consensus Algorithms*)

In this section, we run simulations comparing our solver to three other consensus algorithms. We randomly initialized $N = 10$ agents $\mathbf{R} \in \mathcal{C}$ and run each algorithm with the same initial conditions in order to compare them. We choose the following two metrics for comparison: the *consensus error* (5.11), and the *Frechet mean Error* $E_{\text{RCM}}(\mathbf{R}) = \sum_{i=1}^N d(R_i, \bar{Z})^2$ which aggregates the error of each point $R_1, \dots, R_N$ from the RCM $\bar{Z}$. Using these two metrics, we compare our solver against three other algorithms described in the following and illustrate the results in 5.1. For technical details of their implementation see the code [115].

Algorithm 5.14

To implement the continuous dynamics of our algorithm, we consider their forward-Euler discretization 9 on each tangent space with step size $\epsilon = 0.1$ and use the Lie group exponential mapping as our choice of retraction. We run $K = 200$ iterations.

Algorithm 9 Discretization of Algorithm 5.14

Input: $\mathbf{Z} \in \mathcal{C}$, $K \in \mathbb{N}$, $\epsilon > 0$

Initialize: $R_i \leftarrow Z_i$, $\alpha_i \leftarrow 0 \in \mathfrak{so}(3)$ for $1 \leq i \leq N$

for $k = 1, 2, \dots, K$ **do**

for $i = 1, 2, \dots, N$ **do**

$$\gamma_i \leftarrow -\log(R_i^{-1}Z_i)$$

$$\omega_i \leftarrow -\alpha_i + \gamma_i$$

$$U_i \leftarrow \sum_{j \in \mathcal{N}_i} \log(R_i^{-1}R_j) - \omega_i$$

$$R_i \leftarrow R_i \exp(\epsilon U_i)$$

$$\alpha_i \leftarrow \alpha_i + \epsilon \sum_{j \in \mathcal{N}_i} (\alpha_j - \alpha_i + \gamma_i - \gamma_j)$$

end for

end for

Riemannian Consensus Algorithm

As stated before, [83] was one of the first works to introduce a consensus algorithm for arbitrary Riemannian manifolds.

Given N agents and a communication graph $\mathbf{G}$, the dynamics of the i -th agent is given by

$$R_i(k+1) = R_i(k) \exp \left[\epsilon \sum_{j \in \mathcal{N}_i} \log(R_i(k)^\top R_j(k)) \right]. \quad (5.24)$$

Penalty method

The penalty approach is a commonly used solver for constrained optimization problems. We implemented [79, Algorithm 14.3.1], while introducing a distributed implementation for it.

The Penalty Method proceeds as follows:

1. **(initialization)** Pick $\mathbf{Z} \in \mathcal{C}$. Set $\mathbf{R}(0) = \mathbf{Z}$. Pick a step size $\epsilon > 0$ and a decreasing sequence $(\mu(s))$ that (slowly) converges to 0.

2. **(inner minimization)** Approximately minimize

$$\Phi_0(\mathbf{R}, \mu(s)) = f(\mathbf{R}) + \frac{1}{2\mu(s)} \sum_{i=1}^N \varphi_i(\mathbf{R})^2$$

over $\mathcal{G}^N$ via gradient descent, where

$$\varphi_i(\mathbf{R}) := \frac{1}{2} \sum_{j \in \mathcal{N}_i} d(R_i, R_j)^2.$$

As such, each agent i updates as follows:

$$R_i(k+1) \leftarrow R_i(k) \exp(-\epsilon R_i(k)^\top \nabla_{R_i} \Phi_0(\mathbf{R}(k), \mu(s))).$$

3. **(Update the penalty parameter)** $\mu(s+1) \leftarrow \mu(s)$.
4. **(Choose the next starting point)** Repeat steps 1,2, starting at the current $\mathbf{R}(k)$.

Lagrangian method

A Lagrangian method is an approach to solving constrained optimization problems that involves finding saddle points of the Lagrangian function. We implemented the solver described in [116, 4.4.1] and [117].

Recall the Lagrangian function of (5.5) is

$$L(\mathbf{R}, \lambda) = f(\mathbf{R}) + \sum_{i=1}^N \lambda_i \varphi_i(\mathbf{R}).$$

The Lagrangian method then proceeds with the following updates:

$$\begin{aligned} R_i(k+1) &= R_i(k) \exp(-\epsilon \nabla_{R_i} L(\mathbf{R}(k), \lambda(k))) \\ \lambda_i(k+1) &= \lambda_i(k) + \epsilon \nabla_{\lambda_i} L(\mathbf{R}(k), \lambda(k)). \end{aligned}$$

Since $\mathcal{G} = \text{SO}(3)$ is a Lie group equipped with a bi-invariant metric, we can explicitly write

$$\nabla_{R_i} L(\mathbf{R}, \lambda) = R_i \log(Z_i^{-1} R_i) - \lambda_i \sum_{j \in \mathcal{N}_i} R_i \log(R_i^{-1} R_j),$$

$$\nabla_{\lambda_i} L(\mathbf{R}, \lambda) = \varphi_i(\mathbf{R}).$$

All algorithms are implemented in discrete time and initialized from the same initial conditions. While different step sizes and update schedules were used across methods, the qualitative rate behavior (i.e., linear versus sublinear rate) was consistent across reasonable parameter choices.

Figure 5.1 compares the evolution of the consensus error and the RCM error for the proposed algorithm and three baseline methods. The proposed algorithm exhibits a clear linear convergence rate in both metrics, as indicated by the straight-line decay on the logarithmic scale.

The Riemannian consensus algorithm of [83] achieves rapid reduction of the consensus error and also displays an approximately linear rate in this metric. However, this method does not converge to the Riemannian center of mass: while consensus is reached, the RCM error stagnates at a nonzero value, indicating convergence to a configuration that is not the Frechet mean.

In contrast, both the penalty-based method and the Lagrangian method reduce the consensus and RCM errors more slowly. Their convergence appears sublinear in both metrics, consistent with the use of diminishing penalty parameters or dual updates and the absence of explicit gradient-tracking mechanisms.

5.7.2 Scenario 2 (Linear rate of convergence)

In this scenario, we generate 100 different problem instances with $N = 10$ agents randomly initialized on $\mathcal{C} \subset \text{SO}(3)^{10}$. We run the algorithm on each instance for 200 iterations for a visualization of the convergence rate. We illustrate the result




Figure 5.1: The Consensus Error and the RCM Error at each iteration, comparing Algorithm 9 (in blue) with three other algorithms.

in Figure 5.2 showing the statistics of the RCM Error at each time step, confirming a linear convergence rate. We note that although convergence analysis of the algorithm has only been provided for the Euclidean case in this work, these numerical experiments—using different communication graphs and initial states—demonstrate the effectiveness of the proposed approach for distributed RCM consensus on Lie groups.




Figure 5.2: Statistical visualization of the Consensus error and RCM error of our proposed 9 over 100 randomly sampled problem instances.

5.8 *Concluding Remarks*

In this chapter, we reformulated the computation of the Frechet mean as a consensus-constrained distributed optimization problem and proposed an intrinsic distributed algorithm for its solution on Lie groups equipped with bi-invariant Riemannian metrics. By exploiting the geometric structure of such Lie groups, the proposed method couples a consensus mechanism on the group with a dynamic consensus mechanism on the associated Lie algebra, yielding a gradient-tracking formulation that avoids nested consensus subroutines.

We characterized the stationary points of the resulting closed-loop dynamics within a convexity domain and showed that they correspond to consensus at the Frechet mean. In the Euclidean setting, we further established global convergence guarantees and exponential convergence rates. For Lie groups, the analysis is restricted to correctness of stationary points, while convergence remains an open problem.

Numerical simulations on $SO(3)$ corroborate the theoretical findings and illustrate favorable convergence behavior compared to standard distributed optimization baselines. Extending the convergence analysis to broader classes of Lie groups and to general Riemannian manifolds, as well as developing tools to control curvature-induced transport effects in gradient tracking, constitute important directions for future work.

BIBLIOGRAPHY

- [1] Constantin Udriste. *Convex functions and optimization methods on Riemannian manifolds*, volume 297. Springer Science & Business Media, 2013.
- [2] Yang Zheng, Chih-fan Pai, and Yujie Tang. Benign nonconvex landscapes in optimal and robust control, part I: Global optimality. arXiv preprint, 2023.
- [3] Yang Zheng, Chih-Fan Pai, and Yujie Tang. Benign nonconvex landscapes in optimal and robust control, part ii: Extended convex lifting. arXiv preprint, 2024.
- [4] Nicolas Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023.
- [5] John M Lee. *Smooth Manifolds*. Springer, 2012.
- [6] John M Lee. *Introduction to Riemannian Manifolds*, volume 2. Springer, 2018.
- [7] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [8] Michael Spivak. *A Comprehensive Introduction to Differential Geometry, Volume 2*. Publish or Perish, Inc., Houston, TX, 2 edition, 1979.
- [9] Hermann Weyl. *Philosophy of mathematics and natural science*. Princeton University Press, 2021.
- [10] Michael Spivak. *A Comprehensive Introduction to Differential Geometry, Volume 1*. Publish or Perish, Inc., Houston, TX, 2 edition, 1979.
- [11] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry*, volume 1. New York, London, 1963.
- [12] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of Differential Geometry, Volume 2*, volume 61. John Wiley & Sons, 1996.
- [13] Hermann Weyl and Gerald R MacLane. *The concept of a Riemann surface*. Courier Corporation, 2009.

- [14] John Willard Milnor. *Morse theory*. Princeton university press, 1963.
- [15] María Barbero-Liñán and David Martín de Diego. Retraction maps: a seed of geometric integrators. *Foundations of Computational Mathematics*, 23(4):1335–1380, 2023.
- [16] Uwe Helmke. The topology of a moduli space for linear dynamical systems. *Commentarii Mathematici Helvetici*, 60:630–655, 1985.
- [17] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*. SIAM, 1994.
- [18] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [19] Steven Bradtke. Reinforcement learning applied to linear quadratic regulation. *Advances in neural information processing systems*, 5, 1992.
- [20] Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2:253–279, 2019.
- [21] Bin Hu, Kaiqing Zhang, Na Li, Mehran Mesbahi, Maryam Fazel, and Tamer Başar. Toward a theoretical foundation of policy optimization for learning control policies. *Annual Review of Control, Robotics, and Autonomous Systems*, 6:123–158, 2023.
- [22] Yujie Tang, Yang Zheng, and Na Li. Analysis of the optimization landscape of Linear Quadratic Gaussian (LQG) control. In *Learning for Dynamics and Control*, pages 599–610. PMLR, 2021.
- [23] Jingjing Bu, Afshin Mesbahi, Maryam Fazel, and Mehran Mesbahi. LQR through the lens of first-order methods: Discrete-time case. arXiv preprint, 2019.
- [24] Zhaolin Ren, Yang Zheng, Maryam Fazel, and Na Li. On controller reduction in linear quadratic gaussian control with performance bounds. In *Learning for Dynamics and Control Conference*, pages 1008–1019. PMLR, 2023.
- [25] Yang Zheng, Yue Sun, Maryam Fazel, and Na Li. Escaping high-order saddles in policy optimization for Linear Quadratic Gaussian (LQG) control. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5329–5334. IEEE, 2022.

- [26] Christopher I Byrnes and Peter L Falb. Applications of algebraic geometry in system theory. *American Journal of Mathematics*, 101(2):337–363, 1979.
- [27] Christopher I. Byrnes and Norman E. Hurt. On the moduli of linear dynamical systems. In Gian-Carlo Rota, editor, *Studies in Analysis*, volume 4 of *Advances in Mathematics Supplementary Studies*, pages 83–121. Academic Press, New York, NY, 1979. Third paper in the collection, Total pages: 323.
- [28] Christopher I Byrnes. Algebraic and geometric aspects of the analysis of feedback systems. *Geometrical Methods for the Theory of Linear Systems*, pages 85–124, 1980.
- [29] Christopher I Byrnes and Tyrone E Duncan. On certain topological invariants arising in system theory. In *New Directions in Applied Mathematics: Papers Presented April 25/26, 1980, on the Occasion of the Case Centennial Celebration*, pages 29–71. Springer, 1982.
- [30] Christopher I Byrnes and Clyde F Martin. *Geometrical Methods for the Theory of Linear Systems: Proceedings of a NATO Advanced Study Institute and AMS Summer Seminar in Applied Mathematics held at Harvard University, Cambridge, Mass., June 18–29, 1979*, volume 62. Springer Science & Business Media, 2012.
- [31] Christopher I Byrnes and Norman Hurt. On the moduli of linear dynamical systems. *Adv. in Math.: Studies in Analysis*, 4, 1978.
- [32] Roger Brockett. Some geometric questions in the theory of linear systems. *IEEE Transactions on Automatic Control*, 21(4):449–455, 1976.
- [33] R Brockett and C Byrnes. Multivariable nyquist criteria, root loci, and pole placement: a geometric viewpoint. *IEEE Transactions on Automatic Control*, 26(1):271–284, 1981.
- [34] Roger Brockett. The early days of geometric nonlinear control. *Automatica*, 50(9):2203–2224, 2014.
- [35] Michiel Hazewinkel. (fine) moduli (spaces) for linear systems: what are they and what are they good for? In *Geometrical Methods for the Theory of Linear Systems: Proceedings of a NATO Advanced Study Institute and AMS Summer Seminar in Applied Mathematics held at Harvard University, Cambridge, Mass., June 18–29, 1979*, pages 125–193. Springer, 1980.

- [36] Michiel Hazewinkel. Moduli and canonical forms for linear dynamical systems II: The topological case. *Mathematical Systems Theory*, 10(1):363–385, 1976.
- [37] Michiel Hazewinkel. Moduli and canonical forms for linear dynamical systems iii: The algebraic geometric case. In *Proc. of the 76 Ames Research Center (NASA) Conference on Geometric Control Theory*, pages 291–336, 1977.
- [38] Michiel Hazewinkel and Rudolf E Kalman. On invariants, canonical forms and moduli for linear, constant, finite dimensional, dynamical systems. In *Mathematical Systems Theory: Proceedings of the International Symposium Udine, Italy, June 16–27, 1975*, pages 48–60. Springer, 1976.
- [39] Perinkulam Sambamurthy Krishnaprasad. *Geometry of Minimal Systems and the Identification Problem*. Harvard University Press, 1977.
- [40] P. S. Krishnaprasad and C. F. Martin. On families of systems and deformations. *International Journal of Control*, 38(5):1055–1079, 1983.
- [41] Uwe Helmke. Topology of the moduli space for reachable linear dynamical systems: The complex case. *Mathematical Systems Theory*, 19:155–187, 1986.
- [42] U Helmke and PA Fuhrmann. Tangent spaces of rational matrix functions. *Linear Algebra and its Applications*, 271(1-3):1–40, 1998.
- [43] Uwe Helmke and Christopher I Byrnes. The cohomology of the moduli space of controllable linear systems. *Acta Applicandae Mathematica*, 28:161–188, 1992.
- [44] Robert Hermann. *Linear Systems Theory & Introductory Algebraic Geometry*, volume 8. Math-Sci Press, 1974.
- [45] Robert Hermann. *Topics in the geometric theory of linear systems*, volume 22. Math Sci Press, 1984.
- [46] RE Kalman. Algebraic geometric description of the class of linear systems of constant dimension. In *8th Annual Princeton Conference on Information Sciences and Systems*, volume 3. Princeton University Princeton, NJ, 1974.
- [47] Samuel Eilenberg. *Automata, Languages, and Machines*. Academic press, 1974.
- [48] Allen Tannenbaum. *Invariance and System Theory: Algebraic and Geometric Aspects*, volume 845. Springer, 2006.

- [49] Athanasios C Antoulas. *Mathematical system theory: the influence of RE Kalman*. Springer Science & Business Media, 2013.
- [50] Edmond A Jonckheere. *Algebraic and differential topology of robust stability*. Oxford University Press, 1997.
- [51] Kemin Zhou and Pramod P Khargonekar. An algebraic riccati equation approach to $\mathcal{H}_\infty$ optimization. *Systems & Control Letters*, 11(2):85–91, 1988.
- [52] Spencer Kraisler and Mehran Mesbahi. Output-feedback synthesis orbit geometry: Quotient manifolds and lqg direct policy optimization. *IEEE Control Systems Letters*, 2024.
- [53] Shahriar Talebi, Yang Zheng, Spencer Kraisler, Na Li, and Mehran Mesbahi. Policy optimization in control: Geometry and algorithmic implications. *arXiv preprint arXiv:2406.04243*, 2024.
- [54] Shahriar Talebi and Mehran Mesbahi. Riemannian constrained policy optimization via geometric stability certificates. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 1472–1478, 2022.
- [55] Petar Mlinarić, Christopher A Beattie, Zlatko Drmač, and Serkan Gugercin. Irka is a riemannian gradient descent method. *IEEE transactions on automatic control*, 2024.
- [56] Jingjing Bu, Afshin Mesbahi, and Mehran Mesbahi. On topological and metrical properties of stabilizing feedback gains: The MIMO case. *arXiv preprint*, 2019.
- [57] Jingjing Bu, Afshin Mesbahi, and Mehran Mesbahi. Policy gradient-based algorithms for continuous-time linear quadratic control. *arXiv preprint*, 2020.
- [58] Bijan Afsari and Rene Vidal. Bundle reduction and the alignment distance on spaces of state-space LTI systems. *IEEE Transactions on Automatic Control*, 62(8):3804–3819, 2017.
- [59] Tim Netzer and Daniel Plaumann. *Geometry of Linear Matrix Inequalities*. Springer, 2023.
- [60] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real algebraic geometry*, volume 36. Springer Science & Business Media, 2013.

- [61] Didier Henrion and J-B Lasserre. Convergent relaxations of polynomial matrix inequalities and static output feedback. *IEEE Transactions on Automatic Control*, 51(2):192–202, 2006.
- [62] Joao P Hespanha. *Linear Systems Theory*. Princeton university press, 2018.
- [63] John C Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on automatic Control*, 23(4):756–757, 1978.
- [64] Nicolas Boumal, Pierre-Antoine Absil, and Coralia Cartis. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2019.
- [65] Unsik Lee and Mehran Mesbahi. Optimal power descent guidance with 6-DoF line of sight constraints via unit dual quaternions. In *AIAA Guidance, Navigation, and Control Conference*, 2015.
- [66] Unsik Lee and Mehran Mesbahi. Dual quaternions, rigid body mechanics, and powered-descent guidance. In *IEEE Conference on Decision and Control*, pages 3386–3391, 2012.
- [67] Michael Szmuk, Taylor Reynolds, Behcet Acikmese, Mehran Mesbahi, and John M Carson. Successive convexification for 6-DoF powered descent guidance with compound state-triggered constraints. In *AIAA Scitech Forum*, 2019.
- [68] Steven Thomas Smith. Optimization techniques on Riemannian manifolds. *arXiv preprint arXiv:1407.5965*, 2014.
- [69] Changshuo Liu and Nicolas Boumal. Simple algorithms for optimization on riemannian manifolds with constraints. *Applied Mathematics & Optimization*, 82(3):949–981, 2020.
- [70] Jiang Hu, Xin Liu, Zai-Wen Wen, and Ya-Xiang Yuan. A brief introduction to manifold optimization. *Journal of the Operations Research Society of China*, 8:199–248, 2020.
- [71] Ronny Bergmann and Roland Herzog. Intrinsic formulation of KKT conditions and constraint qualifications on smooth manifolds. *SIAM Journal on Optimization*, 29(4):2423–2444, 2019.
- [72] Alessandro Saccon, John Hauser, and A Pedro Aguiar. Optimal control on Lie groups: The projection operator approach. *IEEE Transactions on Automatic Control*, 58(9):2230–2245, 2013.

- [73] John Hauser and David G Meyer. The trajectory manifold of a nonlinear control system. In *IEEE Conference on Decision and Control*, pages 1034–1039, 1998.
- [74] Yuanqi Mao, Michael Szmuk, Xiangru Xu, and Behçet Açikmese. Successive convexification: A superlinearly convergent algorithm for non-convex optimal control problems. *arXiv preprint arXiv:1804.06539*, 2018.
- [75] Danylo Malyuta, Taylor P. Reynolds, Michael Szmuk, Thomas Lew, Riccardo Bonalli, Marco Pavone, and Behçet Açikmese. Convex optimization for trajectory generation: A tutorial on generating dynamically feasible trajectories reliably and efficiently. *IEEE Control Systems Magazine*, 42(5):40–113, 2022.
- [76] Riccardo Bonalli, Andrew Bylard, Abhishek Cauligi, Thomas Lew, and Marco Pavone. Trajectory optimization on manifolds: a theoretically-guaranteed embedded sequential convex programming approach. *arXiv preprint arXiv:1905.07654*, 2019.
- [77] Spencer Kraisler, Mehran Mesbahi, and Behçet Açikmese. Intrinsic successive convexification: Trajectory optimization on smooth manifolds. *IEEE Control Systems Letters*, 2025.
- [78] Taylor P Reynolds and Mehran Mesbahi. The crawling phenomenon in sequential convex programming. In *2020 American Control Conference (ACC)*, pages 3613–3618. IEEE, 2020.
- [79] Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- [80] Fulvio Forni and Rodolphe Sepulchre. Differential analysis of nonlinear systems: Revisiting the pendulum example. In *53rd IEEE Conference on Decision and Control*, pages 3848–3859. IEEE, 2014.
- [81] Fulvio Forni and Rodolphe Sepulchre. Differentially positive systems. *IEEE Transactions on Automatic Control*, 61(2):346–359, 2015.
- [82] Cyrus Mostajeran and Rodolphe Sepulchre. Positivity, monotonicity, and consensus on lie groups. *SIAM Journal on Control and Optimization*, 56(3):2436–2461, 2018.
- [83] Roberto Tron, Bijan Afsari, and René Vidal. Riemannian consensus for manifolds with bounded curvature. *IEEE Transactions on Automatic Control*, 58(4):921–934, 2012.

- [84] Joan Sola. Quaternion kinematics for the error-state Kalman filter. *arXiv preprint arXiv:1711.02508*, 2017.
- [85] Wulf Rossmann. *Lie Groups: an Introduction Through Linear Groups*. Oxford University Press, 2006.
- [86] Ernst Hairer, Marlis Hochbruck, Arieh Iserles, and Christian Lubich. Geometric numerical integration. *Oberwolfach Reports*, 3(1):805–882, 2006.
- [87] Steven H Strogatz. *Sync: How Order Emerges from Chaos in the Universe, Nature, and Daily Life*. Hachette UK, 2012.
- [88] Reza Olfati-Saber, J Alex Fax, and Richard M Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [89] Solmaz S Kia, Bryan Van Scoy, Jorge Cortes, Randy A Freeman, Kevin M Lynch, and Sonia Martinez. Tutorial on dynamic average consensus: The problem, its applications, and the algorithms. *IEEE Control Systems Magazine*, 39(3):40–72, 2019.
- [90] Reza Olfati-Saber. Swarms on sphere: A programmable swarm with synchronous behaviors like oscillator networks. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 5060–5066. IEEE, 2006.
- [91] Alain Sarlette, Silvere Bonnabel, and Rodolphe Sepulchre. Coordinated motion design on Lie groups. *IEEE Transactions on Automatic Control*, 55(5):1047–1058, 2010.
- [92] Rodolphe Sepulchre. Consensus on nonlinear spaces. *IFAC Proceedings Volumes*, 43(14):1029–1039, 2010.
- [93] Sheng Chen, Lindu Zhao, Weigong Zhang, and Peng Shi. Consensus on compact Riemannian manifolds. *Inf. Sci.*, 268:220–230, June 2014.
- [94] Sheng Chen, Peng Shi, Weigong Zhang, and Lindu Zhao. Finite-time consensus on strongly convex balls of Riemannian manifolds with switching directed communication topologies. *J. Math. Anal. Appl.*, 409(2):663–675, January 2014.
- [95] Spencer Kraisler, Shahriar Talebi, and Mehran Mesbahi. Distributed consensus on manifolds using the Riemannian center of mass. *IEEE Conference on Control Technology and Applications*, 2023.

- [96] Roberto Tron and Rene Vidal. Distributed 3-d localization of camera sensor networks from 2-d image measurements. *IEEE Transactions on Automatic Control*, 59(12):3325–3340, 2014.
- [97] Bijan Afsari. Riemannian L^p center of mass: Existence, uniqueness, and convexity. *Proceedings of the American Mathematical Society*, 139(2):655–673, 2011.
- [98] Roberto Tron, René Vidal, and Andreas Terzis. Distributed pose averaging in camera networks via consensus on $SE(3)$. In *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pages 1–10. IEEE, 2008.
- [99] Roberto Tron, Bijan Afsari, and René Vidal. Average consensus on riemannian manifolds with bounded curvature. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7855–7862. IEEE, 2011.
- [100] Spencer Kraisler, Shahriar Talebi, and Mehran Mesbahi. Consensus on lie groups for the riemannian center of mass. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 4461–4466. IEEE, 2023.
- [101] Bijan Afsari. *Means and averaging on Riemannian manifolds*. University of Maryland, College Park, 2009.
- [102] Hermann Karcher. Riemannian center of mass and so called Karcher mean. *arXiv preprint arXiv:1407.2087*, 2014.
- [103] Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on pure and applied mathematics*, 30(5):509–541, 1977.
- [104] Xavier Pennec. Barycentric subspace analysis on manifolds. *Ann. Stat.*, 46(6A):2711–2746, December 2018.
- [105] Bijan Afsari, Roberto Tron, and René Vidal. On the convergence of gradient descent for finding the riemannian center of mass. *SIAM Journal on Control and Optimization*, 51(3):2230–2260, 2013.
- [106] Solmaz S Kia, Jorge Cortés, and Sonia Martínez. Distributed convex optimization via continuous-time coordination algorithms with discrete-time communication. *Automatica*, 55:254–264, 2015.
- [107] Angelia Nedic, Alex Olshevsky, and Wei Shi. Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4):2597–2633, 2017.

- [108] Guannan Qu and Na Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 5(3):1245–1260, 2017.
- [109] Angelia Nedić and Alex Olshevsky. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615, 2014.
- [110] Reza Olfati-Saber and Richard M Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on automatic control*, 49(9):1520–1533, 2004.
- [111] Shixiang Chen, Alfredo Garcia, Mingyi Hong, and Shahin Shahrampour. Decentralized riemannian gradient descent on the stiefel manifold. In *International Conference on Machine Learning*, pages 1594–1605. PMLR, 2021.
- [112] Suhail M Shah. Distributed optimization on Riemannian manifolds for multi-agent networks. *arXiv preprint arXiv:1711.11196*, 2017.
- [113] Jun Chen, Lina Liu, Tianyi Zhu, Yong Liu, Guang Dai, Yunliang Jiang, and Ivor W Tsang. Decentralized optimization on compact submanifolds by quantized riemannian gradient tracking. *IEEE Transactions on Signal Processing*, 2025.
- [114] Mehran Mesbahi and Magnus Egerstedt. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [115] Spencer Kraisler, Shahriar Talebi, and Mehran Mesbahi. RCM-Consensus-Lie-Group, 2023.
- [116] Dimitri P Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*. Academic press, 2014.
- [117] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.
- [118] Andreas Arvanitogeorgos. *An introduction to Lie groups and the geometry of homogeneous spaces*, volume 22. American Mathematical Soc., 2003.
- [119] Brian C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*. Springer, 2013.

- [120] John Milnor. Curvatures of left invariant metrics on Lie groups. *Advances in mathematics*, 21(3):293–329, 1976.
- [121] J.R. Munkres. *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated, 2000.
- [122] David Mumford, John Fogarty, and Frances Kirwan. *Geometric invariant theory*, volume 34. Springer Science & Business Media, 1994.

Appendix A

LIE THEORY

A comprehensive reference on Lie groups is [118, 119], which inspired these notes.

Lie Groups A *Lie group* $\mathcal{G}$ is a smooth manifold equipped with a group structure such that the multiplication and inverse maps

$$\begin{aligned}(p, q) &\mapsto p \cdot q, \\ p &\mapsto p^{-1},\end{aligned}$$

are smooth. The identity element is denoted by e .

Example A.0.1. *The real numbers $\mathbb{R}$, complex numbers $\mathbb{C}$, and the quaternions $\mathbb{H}$ are Lie groups under their usual multiplication.*

Example A.0.2. *The general linear group $\mathrm{GL}(n)$ under matrix multiplication, as well as any closed subgroup of $\mathrm{GL}(n)$, are Lie groups.*

A subgroup $\mathcal{H} \subset \mathcal{G}$ is called a *Lie subgroup* if the inclusion map $\iota : \mathcal{H} \hookrightarrow \mathcal{G}$ is an immersion. Every closed subgroup of a Lie group is an embedded Lie subgroup.

If $\mathcal{G}$ and $\mathcal{H}$ are Lie groups and $F : \mathcal{G} \rightarrow \mathcal{H}$ is a smooth group homomorphism, then F is called a *Lie group homomorphism*. When F is a diffeomorphism, $\mathcal{G}$ and $\mathcal{H}$ are said to be *Lie group isomorphic*.

Translations For $p \in \mathcal{G}$, define the left- and right-translations

$$L_p(q) := p \cdot q, \quad R_p(q) := q \cdot p.$$

Both are diffeomorphisms. Their differentials are written

$$dL_p : T_q\mathcal{G} \rightarrow T_{p \cdot q}\mathcal{G}, \quad dR_p : T_q\mathcal{G} \rightarrow T_{q \cdot p}\mathcal{G},$$

with the base point understood from context.

Parallelizability Every Lie group is parallelizable. If (ξ_i) is a basis of $T_e\mathcal{G}$, define vector fields

$$E_i|_p := dL_p(\xi_i).$$

Then (E_i) is a global frame. Consequently, every tangent space $T_p\mathcal{G}$ is canonically identified with $T_e\mathcal{G}$ via left translation by p^{-1} .

Invariant Vector Fields and the Lie Algebra A vector field $V \in \mathfrak{X}(\mathcal{G})$ is *left-invariant* if

$$V_{p \cdot q} = dL_p(V_q), \quad \forall p, q \in \mathcal{G}.$$

Such a vector field is uniquely determined by its value at the identity:

$$V_p = dL_p(V_e).$$

The set of left-invariant vector fields is denoted $\mathfrak{g}$. A remarkable fact is $\mathfrak{g}$ is closed under the Lie bracket. That is, if $V, W \in \mathfrak{g}$, then $[V, W] \in \mathfrak{g}$. Since $\mathfrak{g}$ is identified with $T_e\mathcal{G}$, then $[\cdot, \cdot]$ can be defined on $T_e\mathcal{G}$ as

$$[\xi, \eta] := [V, W]_e, \quad \xi, \eta \in T_e\mathcal{G}.$$

We call $(\mathfrak{g}, [\cdot, \cdot]) \equiv (T_e\mathcal{G}, [\cdot, \cdot])$ the *Lie algebra* of $\mathcal{G}$.

Matrix Lie Groups Let $GL(n)$ denote the set of invertible $n \times n$ real matrices. As an open subset of $\mathbb{R}^{n \times n}$, $GL(n)$ forms an open embedded submanifold. Equipped with matrix multiplication, $GL(n)$ forms a Lie group. Any closed subgroup of $GL(n)$ is also a closed embedded submanifold, and hence a Lie group. As a result, we call closed subgroups of $GL(n)$ *matrix Lie groups*.

Theorem A.0.1. *If $\mathcal{G} \subset \mathrm{GL}(n)$ is a matrix Lie group and $A, B \in T_I \mathcal{G}$, then*

$$[A, B] = AB - BA.$$

One-Parameter Subgroups and the Lie Exponential A *one-parameter subgroup* is a smooth curve $\phi : \mathbb{R} \rightarrow \mathcal{G}$ satisfying

$$\phi(s+t) = \phi(s) \cdot \phi(t), \quad \phi(0) = e.$$

Theorem A.0.2. *Every one-parameter subgroup is uniquely determined by its derivative at $t = 0$.*

This defines the *Lie exponential map*

$$\exp : \mathfrak{g} \equiv T_e \mathcal{G} \rightarrow \mathcal{G}, \quad \exp(\xi) := \phi(1),$$

where ϕ is the unique one-parameter subgroup with $\dot{\phi}(0) = \xi$.

The differential $d\exp_0$ is the identity, hence $\exp$ admits a local inverse near 0, called the *Lie logarithm*

$$\log : \exp(U) \subset \mathcal{G} \rightarrow U \subset \mathfrak{g}.$$

The exponential satisfies

$$\exp(-\xi) = \exp(\xi)^{-1}, \quad \exp((s+t)\xi) = \exp(s\xi) \exp(t\xi),$$

but generally $\exp(\xi + \eta) \neq \exp(\xi) \exp(\eta)$ unless $\mathcal{G}$ is abelian.

For matrix Lie groups,

$$\exp(A) = \sum_{k=0}^{\infty} \frac{A^k}{k!}, \quad \log(R) = - \sum_{k=1}^{\infty} \frac{(I - R)^k}{k},$$

when the series converge.

Bi-Invariant Riemannian Metrics A Riemannian metric g on $\mathcal{G}$ is *left-invariant* if

$$g_p(v, w) = g_{q \cdot p}(\mathrm{d}L_q v, \mathrm{d}L_q w),$$

and *right-invariant* if

$$g_p(v, w) = g_{p \cdot q}(\mathrm{d}R_q v, \mathrm{d}R_q w).$$

A metric that is both left- and right-invariant is called *bi-invariant*.

Every Lie group admits left- and right-invariant Riemannian metrics [6, Lemma 3.10]. However, a Lie group admits a bi-invariant metric if and only if it is isomorphic to the Cartesian product of a compact Lie group and a vector space [120, Lemma 7.5]. A canonical example is the group of unit quaternions equipped with the Euclidean inner product.

For bi-invariant Lie groups, the Lie exponential and the Riemannian exponential coincide; and likewise for the logarithm operator [6, Prop. 5.9]. In this case, both left- and right-translations are isometries, and the Levi–Civita connection agrees with the Cartan–Schouten connection. Consequently, for any $p, q \in \mathcal{G}$ and $\xi \in \mathfrak{g}$, the following identities hold (where properly defined) [6, Prop. 5.9]:

$$\begin{aligned} \exp_p(\mathrm{d}L_p \xi) &= p \cdot \exp(\xi), \\ \log_p(q) &= \mathrm{d}L_p \log(p^{-1} \cdot q), \end{aligned}$$

where $\exp$ and $\exp_p$ denote the Lie and manifold exponential maps, respectively, and likewise for $\log$ and $\log_p$. Also, for any $p, q \in \mathcal{G}$ such that the logarithms are defined, we have

$$\log(p^{-1} \cdot q) = -\log(q^{-1} \cdot p).$$

[6, Prop. 5.9]

Lie Retraction The Lie exponential induces a canonical retraction

$$R_p(v) := p \cdot \exp(\mathrm{d}L_{p^{-1}} v), \quad R_p^{-1}(q) := \mathrm{d}L_p \log(p^{-1} \cdot q).$$

This retraction is equivariant, smooth, and compatible with bi-invariant metrics.

Appendix B

A NECESSARY BACKGROUND ON TOPOLOGY

These notes follow [121]. Let X be a nonempty set. A *topology* on X is a choice of subsets that we declare to be *open*, intended to formalize when points are “near” one another without assuming any additional structure (such as a metric).

Topological Spaces A topology on X is a collection of subset τ of X satisfying

- $\emptyset, X \in \tau$,
- if $\{U_\alpha\} \subset \tau$, then $\bigcup_\alpha U_\alpha \in \tau$,
- if $U_1, \dots, U_m \in \tau$, then $\bigcap_{i=1}^m U_i \in \tau$.

The pair (X, τ) is a *topological space*. Elements of τ are the *open sets*. A set $C \subset X$ is *closed* if $X \setminus C$ is open. A subset of X may be both open and closed (clopen), or neither.

If $x \in X$, any open set $U \in \tau$ with $x \in U$ is a *neighborhood* of x .

Hausdorff Spaces And Limits A topological space (X, τ) is *Hausdorff* if for every pair of distinct points $p \neq q$ there exist disjoint open sets $U, V \in \tau$ such that $p \in U$ and $q \in V$.

A sequence $(x_k) \subset X$ *converges* to $x^* \in X$ if for every neighborhood U of x^* there exists $N \in \mathbb{N}$ such that $x_k \in U$ for all $k \geq N$. In that case we write $x_k \rightarrow x^*$ and call x^* the *limit* of (x_k) . In a general topological space, limits of a sequence need not be unique. In a Hausdorff space, every convergent sequence has a unique limit.

Bases And Generated Topologies Constructing a topology directly from the axioms can be awkward since τ is often huge. A standard device is to specify a *basis*, which is a “smaller” collection of open sets that uniquely generates a topology. A collection $\mathcal{B}$ of subsets of a set X is called a basis if

- for every $x \in X$ there exists $B \in \mathcal{B}$ with $x \in B$,
- if $x \in B_1 \cap B_2$ with $B_1, B_2 \in \mathcal{B}$, then there exists $B_3 \in \mathcal{B}$ such that $x \in B_3 \subset B_1 \cap B_2$.

The relationship between bases and topologies is analogous to the relationship between smooth atlases and smooth structures in differential geometry. The topology *generated* by $\mathcal{B}$ is the collection of all unions of elements of $\mathcal{B}$. This is the unique topology τ for which $\mathcal{B}$ is a basis. Conversely, if (X, τ) is a topological space, a subcollection $\mathcal{B} \subset \tau$ is a basis for τ if it satisfies the two axioms above.

Example B.0.1. *On $\mathbb{R}^n$, the collection of open balls*

$$B(x, r) := \{y \in \mathbb{R}^n : \|x - y\| < r\}, \quad x \in \mathbb{R}^n, \ r > 0,$$

forms a basis and generates the standard Euclidean topology.

A topological space is *second-countable* if it admits a countable basis. Intuitively, this means that (X, τ) is not “too big”.

Example B.0.2 (The Long Line). *Let*

$$X := \mathbb{R} \times [0, 1).$$

Intuitively, we attach a copy of the half-open interval $[0, 1)$ to each real number and order the resulting set lexicographically.

For points $x = (x_0, x_1)$ and $y = (y_0, y_1)$ in X , define

$$x \preceq y \quad \text{if and only if} \quad x_0 < y_0 \text{ or } (x_0 = y_0 \text{ and } x_1 \leq y_1),$$

and write $x \prec y$ when $x \preceq y$ and $x \neq y$. With this order, $(X, \preceq)$ is a totally ordered set.

Given $x, y \in X$ with $x \prec y$, define the open interval

$$(x, y) := \{a \in X : x \prec a \prec y\}.$$

Let

$$\mathcal{B} := \{(x, y) : x, y \in X, x \prec y\}.$$

Then $\mathcal{B}$ is a basis for a topology on X , called the order topology induced by $\preceq$. The resulting topological space (X, τ) is known as the long line.

The long line is not second-countable: no countable collection of open intervals can form a basis for its topology. Intuitively, the space is “too long” to be captured by countably many local neighborhoods.

Continuous Maps And Homeomorphisms Let $f : X \rightarrow Y$ be a map between topological spaces. The map f is *continuous at* $x \in X$ if for every neighborhood V of $f(x)$ there exists a neighborhood U of x such that $f(U) \subset V$. If f is continuous at every point, it is *continuous*.

A bijection $f : X \rightarrow Y$ is a *homeomorphism* if both f and f^{-1} are continuous. If $f : X \rightarrow Y$ is a homeomorphism onto its image $f(X) \subset Y$, then f is called a *(topological) embedding*.

A *curve* in X is a continuous map $\gamma : I \rightarrow X$ where $I \subset \mathbb{R}$ is an interval.

Connectedness And Path-Connectedness A topological space (X, τ) is *disconnected* if there exist nonempty, disjoint open sets $U, V \in \tau$ such that $X = U \cup V$. If no such decomposition exists, then X is *connected*.

The space X is *path-connected* if for every $x, y \in X$ there exists a continuous curve $\gamma : [0, 1] \rightarrow X$ with $\gamma(0) = x$ and $\gamma(1) = y$. Every path-connected space is connected, but not conversely; a standard example of a connected but not path-connected set is

the *topologist's sine curve*

$$\left\{ \left(x, \sin\left(\frac{1}{x}\right) \right) : x \in \mathbb{R} \setminus \{0\} \right\} \cup (\{0\} \times [-1, 1]).$$

Subspace Topology If $S \subset X$, the *subspace topology* on S is

$$\tau_S := \{U \cap S : U \in \tau\}.$$

Equivalently, $W \subset S$ is open in S if and only if $W = U \cap S$ for some open set $U \subset X$.

Open Covers And Compactness A collection $\{U_\alpha\}_{\alpha \in A} \subset \tau$ is an *open cover* of X if $X = \bigcup_{\alpha \in A} U_\alpha$. The space X is *compact* if every open cover admits a finite subcover, i.e., there exist $\alpha_1, \dots, \alpha_m$ such that $X = \bigcup_{i=1}^m U_{\alpha_i}$.

Appendix C

A NECESSARY BACKGROUND ON ORBIT GEOMETRY

C.1 Smooth Orbit Geometry

Let $\widetilde{\mathcal{M}}$ be a smooth manifold and let $\mathcal{G}$ be a finite-dimensional Lie group acting smoothly on $\widetilde{\mathcal{M}}$. Concretely, we assume each $\alpha \in \mathcal{G}$ is a diffeomorphism $\alpha : \widetilde{\mathcal{M}} \rightarrow \widetilde{\mathcal{M}}$ and the action map

$$\mathcal{G} \times \widetilde{\mathcal{M}} \rightarrow \widetilde{\mathcal{M}}, \quad (\alpha, p) \mapsto \alpha(p)$$

is smooth. We refer to this as a *smooth group action* of $\mathcal{G}$ on $\widetilde{\mathcal{M}}$.

Stabilizers, Orbits, And Invariant Sets For $p \in \widetilde{\mathcal{M}}$, the *stabilizer* (or isotropy subgroup) is

$$\mathcal{G}_p := \{\alpha \in \mathcal{G} : \alpha(p) = p\}.$$

The action is *transitive* if for all $p, q \in \widetilde{\mathcal{M}}$ there exists $\alpha \in \mathcal{G}$ such that $\alpha(p) = q$. For example, the standard action of $SO(3)$ on S^2 is transitive. The action is *free* if $\alpha(p) = p \Rightarrow \alpha = e$ for all p , equivalently if every stabilizer is trivial: $\mathcal{G}_p = \{e\}$.

A subset $U \subset \widetilde{\mathcal{M}}$ is called *$\mathcal{G}$ -invariant* if $\alpha(U) \subset U$ for all $\alpha \in \mathcal{G}$.

For $p \in \widetilde{\mathcal{M}}$, the *orbit* of p is

$$\pi(p) := \{\alpha(p) : \alpha \in \mathcal{G}\}.$$

The map $\pi : \widetilde{\mathcal{M}} \rightarrow \widetilde{\mathcal{M}}/\mathcal{G}$ is called the *orbit map*.

Orbit Space and Orbit Topology The *orbit space* is

$$\mathcal{M} := \widetilde{\mathcal{M}}/\mathcal{G} := \{\pi(p) : p \in \widetilde{\mathcal{M}}\}.$$

We equip $\mathcal{M}$ with the *orbit topology*, in which $U \subset \mathcal{M}$ is open if and only if $\pi^{-1}(U) \subset \widetilde{\mathcal{M}}$ is open. In general, the orbit space need not be Hausdorff; geometric invariant theory studies how such orbit spaces can be decomposed into well-behaved pieces, even when the global orbit is pathological [122].

Smooth Orbit Structure To perform calculus or optimization on $\mathcal{M}$, one seeks a smooth structure on $\mathcal{M}$ compatible with the action. The natural compatibility requirement is that the orbit map $\pi : \widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ be a submersion, since then smooth functions on $\mathcal{M}$ are precisely those whose pullbacks are smooth:

$$f : \mathcal{M} \rightarrow \mathbb{R} \text{ is smooth} \iff f \circ \pi : \widetilde{\mathcal{M}} \rightarrow \mathbb{R} \text{ is smooth.}$$

When such a smooth structure exists, one expects

$$\dim(\mathcal{M}) = \dim(\widetilde{\mathcal{M}}) - \dim(\mathcal{G}).$$

Theorem C.1.1. *Let $\widetilde{\mathcal{M}}$ be a smooth manifold equipped with a smooth action of a Lie group $\mathcal{G}$. The orbit space $\mathcal{M} := \widetilde{\mathcal{M}}/\mathcal{G}$ admits at most one smooth manifold structure for which the orbit map $\pi : \widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ is a submersion. In that case,*

$$\dim(\mathcal{M}) = \dim(\widetilde{\mathcal{M}}) - \dim(\mathcal{G}).$$

When it exists, this is called the *smooth orbit structure*. Existence is not automatic: if the orbit space fails to be Hausdorff, then it cannot even be a topological manifold and therefore cannot carry a smooth structure. Nevertheless, whenever the smooth orbit structure exists, it is unique. For a detailed discussion and its implications for optimization, see [4, §9.3].

Proper Actions A key condition ensuring good orbit behavior is *properness*. The action of $\mathcal{G}$ on $\widetilde{\mathcal{M}}$ is called *proper* if for every convergent sequence $(p_i) \subset \widetilde{\mathcal{M}}$ and every sequence $(\alpha_i) \subset \mathcal{G}$ such that $(\alpha_i(p_i))$ converges in $\widetilde{\mathcal{M}}$, there exists a subsequence of (α_i) that converges in $\mathcal{G}$ [5, Prop. 21.5].

Properness guarantees that orbits behave like embedded submanifolds.

Proposition C.1.1. [5, Prop. 21.7] *Let $\mathcal{G}$ act properly on $\widetilde{\mathcal{M}}$. Then every orbit $\pi(p) \subset \mathcal{M}$ is a closed embedded submanifold of $\mathcal{M}$.*

This viewpoint interprets the orbit manifold $\mathcal{M}$ as the space of embedded submanifolds (orbits) of $\widetilde{\mathcal{M}}$.

Theorem C.1.2. *[5, Thm. 21.10] Let $\mathcal{G}$ act smoothly, freely, and properly on $\widetilde{\mathcal{M}}$. Then the orbit space $\mathcal{M} := \widetilde{\mathcal{M}}/\mathcal{G}$ admits the smooth orbit structure, and the orbit map $\pi : \widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ is a surjective submersion.*

It can also be shown that if $\mathcal{G}$ is a compact Lie group acting smoothly and freely on $\widetilde{\mathcal{M}}$, then it follows $\mathcal{G}$ is a proper action, and so $\mathcal{M}$ inherits the smooth orbit structure.

C.2 Riemannian Orbit Geometry

We do not provide a full development of Riemannian geometry on smooth orbit manifolds here; comprehensive treatments may be found in [4, 7]. While Lie group actions and orbit spaces are classical objects, the systematic study of Riemannian geometry and optimization on smooth orbit manifolds is comparatively recent, emerging primarily in the 1990s and early 2000s in applied mathematics.

Throughout this section, let $\widetilde{\mathcal{M}}$ be a smooth manifold equipped with a *smooth, free, and proper* action of a finite-dimensional Lie group $\mathcal{G}$. Under these assumptions, the orbit space

$$\mathcal{M} := \widetilde{\mathcal{M}}/\mathcal{G}$$

admits the smooth orbit structure, and the orbit map $\pi : \widetilde{\mathcal{M}} \rightarrow \mathcal{M}$ is a surjective submersion. We refer to $\mathcal{M}$ as the *smooth orbit manifold*.

Invariant Riemannian Metrics Let $\widetilde{g}$ be a Riemannian metric on $\widetilde{\mathcal{M}}$. The metric $\widetilde{g}$ is said to be $\mathcal{G}$ -*invariant* if

$$\widetilde{g}_{\alpha(p)}(d\alpha_p(\xi), d\alpha_p(\eta)) = \widetilde{g}_p(\xi, \eta) \quad \forall p \in \widetilde{\mathcal{M}}, \xi, \eta \in T_p\widetilde{\mathcal{M}}, \alpha \in \mathcal{G}.$$

Invariance ensures that the metric assigns identical inner products to tangent vectors related by the group action, and therefore that distances, gradients, and other geometric quantities respect the underlying symmetry.

A $\mathcal{G}$ -invariant Riemannian metric on $\widetilde{\mathcal{M}}$ induces a uniquely defined Riemannian metric on the orbit manifold $\mathcal{M}$ via horizontal lifting. With respect to these metrics, the orbit map

$$\pi : (\widetilde{\mathcal{M}}, \widetilde{g}) \rightarrow (\mathcal{M}, g)$$

is a Riemannian submersion. In this context, we call $(\mathcal{M}, g)$ the Riemannian orbit manifold of $(\widetilde{\mathcal{M}}, \widetilde{g})$.

Vertical and Horizontal Spaces For $p \in \widetilde{\mathcal{M}}$, the orbit $\pi(p) \subset \widetilde{\mathcal{M}}$ is an embedded submanifold. Its tangent space at p is called the *vertical space*:

$$\mathcal{V}_p := T_p \pi(p) = \ker d\pi_p \subset T_p \widetilde{\mathcal{M}}.$$

The vertical space consists of directions tangent to the group orbit and hence corresponds to infinitesimal symmetry transformations.

Given a $\mathcal{G}$ -invariant Riemannian metric $\widetilde{g}$ on $\widetilde{\mathcal{M}}$, the orthogonal complement of $\mathcal{V}_p$ is called the *horizontal space*:

$$\mathcal{H}_p := \{\eta \in T_p \widetilde{\mathcal{M}} : \widetilde{g}_p(\eta, \xi) = 0 \ \forall \xi \in \mathcal{V}_p\} = \mathcal{V}_p^\perp.$$

This decomposition yields the direct sum

$$T_p \widetilde{\mathcal{M}} = \mathcal{V}_p \oplus \mathcal{H}_p,$$

which plays a fundamental role in orbit geometry and optimization.

Identification with the Orbit Tangent Space The differential of the orbit map induces a canonical identification between horizontal vectors in $\widetilde{\mathcal{M}}$ and tangent vectors on the orbit manifold.

Lemma C.2.1. *If $\widetilde{\mathcal{M}}/\mathcal{G}$ admits the smooth orbit structure and $\widetilde{\mathcal{M}}$ is equipped with a $\mathcal{G}$ -invariant metric, then the restriction*

$$d\pi_p|_{\mathcal{H}_p} : \mathcal{H}_p \rightarrow T_{\pi(p)} \mathcal{M}$$

is a linear bijection.

The inverse mapping

$$\text{lift}_p := (d\pi_p|_{\mathcal{H}_p})^{-1} : T_{\pi(p)} \mathcal{M} \rightarrow \mathcal{H}_p$$

is called the *horizontal lift*. It allows tangent vectors on the orbit to be represented concretely as horizontal vectors in $T_p \widetilde{\mathcal{M}}$. The orthogonal projection onto the horizontal space is given by

$$\text{Proj}_{\mathcal{H}_p} = \text{lift}_p \circ d\pi_p : T_p \widetilde{\mathcal{M}} \rightarrow \mathcal{H}_p.$$

Lifting of Functions and Vector Fields Let $\mathcal{N}$ be a smooth manifold and let $f : \mathcal{M} \rightarrow \mathcal{N}$ be smooth. The *lift* of f is the mapping

$$\tilde{f} := f \circ \pi : \widetilde{\mathcal{M}} \rightarrow \mathcal{N}.$$

By construction, $\tilde{f}$ is constant along orbits and therefore $\mathcal{G}$ -invariant. Conversely, every smooth $\mathcal{G}$ -invariant mapping on $\widetilde{\mathcal{M}}$ arises uniquely as the lift of a smooth mapping on $\mathcal{M}$.

Similarly, if $V \in \mathfrak{X}(\mathcal{M})$ is a vector field on the orbit manifold, its *lift* is the vector field $\tilde{V} \in \mathfrak{X}(\widetilde{\mathcal{M}})$ defined by

$$\tilde{V}_p := \text{lift}_p(V_{\pi(p)}).$$

This definition ensures the compatibility condition

$$d\pi_p(\tilde{V}_p) = V_{\pi(p)}, \quad \text{equivalently} \quad d\pi \circ \tilde{V} = V \circ \pi.$$

When the Riemannian metric g is $\mathcal{G}$ -invariant, then the lift of any vector field is smooth; otherwise, we are not guaranteed smoothness and could get a non-smooth vector field (known as a rough vector field).

Equivariant Retractions Let $\tilde{R}$ be a retraction on $\widetilde{\mathcal{M}}$. We say that $\tilde{R}$ is $\mathcal{G}$ -equivariant if

$$\tilde{R}_{\alpha(p)}(d\alpha_p(\eta)) = \alpha(\tilde{R}_p(\eta)) \quad \forall p \in \widetilde{\mathcal{M}}, \eta \in T_p\widetilde{\mathcal{M}}, \alpha \in \mathcal{G}.$$

Equivariance ensures that $\tilde{R}$ respects the group action and maps group orbits to group orbits in a consistent manner.

Let R be a retraction on the orbit manifold $\mathcal{M}$. A retraction $\tilde{R}$ on $\widetilde{\mathcal{M}}$ is said to be a *lift* of R if

$$R_{\pi(p)}(\xi) = (\pi \circ \tilde{R}_p)(\text{lift}_p(\xi)) \quad \forall p \in \widetilde{\mathcal{M}}, \xi \in T_{\pi(p)}\mathcal{M},$$

or equivalently,

$$R = \pi \circ \tilde{R} \circ \text{lift}.$$

Every $\mathcal{G}$ -equivariant retraction $\tilde{R}$ is the lift of a uniquely defined retraction R on the orbit manifold $\mathcal{M}$.

Theorem C.2.1. *Let $\tilde{\mathcal{M}}$ be a smooth manifold equipped with a smooth, free, and proper action of $\mathcal{G}$, a $\mathcal{G}$ -invariant metric $\tilde{g}$, and a $\mathcal{G}$ -equivariant retraction $\tilde{R}$. Let $\tilde{f} \in \mathfrak{F}(\tilde{\mathcal{M}})$ be $\mathcal{G}$ -invariant. Then, for all $p \in \tilde{\mathcal{M}}$ and $\alpha \in \mathcal{G}$,*

$$\begin{aligned}\nabla \tilde{f}_p &\in \mathcal{H}_p, \\ \nabla \tilde{f}_{\alpha(p)} &= d\alpha_p(\nabla \tilde{f}_p), \\ \tilde{R}_{\alpha(p)}(-t\nabla \tilde{f}_{\alpha(p)}) &= \alpha\left(\tilde{R}_p(-t\nabla \tilde{f}_p)\right).\end{aligned}$$

Appendix D

SMOOTH MANIFOLDS: AN EXTRINSIC APPROACH

In this section, we will present the concept of a Riemannian manifold useful for a control theorist. These notes were partially taken from [4].

Definition D.0.1 (Differential). *Let $\mathcal{V}$ and $\mathcal{V}'$ be vector spaces. Let $F : \mathcal{V} \rightarrow \mathcal{V}'$ be smooth at $p \in \mathcal{V}$. The differential of F at p is a linear operator $DF_p : \mathcal{V} \rightarrow \mathcal{V}'$ defined as*

$$DF_p(\eta) := \lim_{\epsilon \rightarrow 0} \frac{F(p + \epsilon\eta) - F(p)}{\epsilon}$$

Definition D.0.2 (Classical Gradient). *Let $(\mathcal{V}, \langle \cdot, \cdot \rangle)$ be an inner product space. Let $f : \mathcal{V} \rightarrow \mathbb{R}$ be smooth. The gradient of f at $p \in \mathcal{V}$ is a vector $\nabla f_p \in \mathcal{V}$ defined as the unique solution to $Df_p(\eta) = \langle \eta, \nabla f_p \rangle$ for all $\eta \in \mathcal{V}$.*

The above definition implies that for a scalar field $f \in \mathfrak{F}(\mathcal{V})$, the gradient ∇f_p is the adjoint of the differential DF_p , which is a co-vector.

D.1 Smooth Manifolds

We now have the necessary tools to define a smooth manifold which is smoothly embedded in $\mathcal{V}$:

Definition D.1.1 (Smooth Manifold). *Let $\mathcal{V}$ be a vector space of dimension d . A nonempty subset $\mathcal{M} \subset \mathcal{V}$ is a smooth manifold of dimension n if either*

1. $n = d$ and $\mathcal{M}$ is an open subset; or,
2. $1 \leq n < d$ and, for each $p \in \mathcal{M}$, there exists a neighborhood U of p in $\mathcal{V}$ and a smooth function $h : U \rightarrow \mathbb{R}^{d-n}$ such that

$$(a) \mathcal{M} \cap U = h^{-1}(0); \text{ and}$$

$$(b) \text{rank}(Dh_p) = d - n.$$

We call h a local defining function. We also call $\mathcal{M}$ a smooth manifold smoothly embedded in $\mathcal{V}$.

Example D.1.1. Take $\mathbb{R}^d$. Let $\mathcal{S} \subset \mathbb{R}$ denote the unit vectors. Set $U := [-2, 2]^d$ and let $h : U \rightarrow \mathbb{R}$ be defined as $h(p) = p^\top p - 1$. Clearly $\mathcal{S} \cap U = \mathcal{S} = h^{-1}(0)$. Also $Dh_p = 2p^\top$ implies $\text{rank}(Dh_p) = 1$ for $p \in \mathcal{S}$. Therefore $\mathcal{S}$ is a smooth manifold of dimension $d - 1$. This implies the unit quaternions and unit dual quaternions are smooth manifolds, being spheres.

From the above example, one can show that the set of rotation matrices $SO(3) \subset \mathbb{R}^{3 \times 3}$ forms a smooth manifold using the local defining function $h(R) := (1 - \det(R)) \cdot (R^\top R - I_3)$.

As a non-example, consider the unit cube $C = \{x \in \mathbb{R}^n : \|x\|_\infty = \max_i |x_i| = 1\}$. No such local defining function exists for the vertices. This is because the cusp renders the shape not smooth.

The most important concept of a smooth manifold is that of a tangent space, which is the linear space tangent to a point on a smooth manifold.

Definition D.1.2 (Tangent Space). *For $p \in \mathcal{M}$, define:*

$$T_p\mathcal{M} := \{c'(0) : c : (-\epsilon, \epsilon) \rightarrow \mathcal{M} \text{ is smooth and } c(0) = p\}.$$

We call $T_p\mathcal{M}$ the tangent space of $\mathcal{M}$ at p . We call elements of $T_p\mathcal{M}$ tangent vectors.

The follow two propositions characterize the geometry of tangent spaces.

Proposition D.1.1. *Let $p \in \mathcal{M}$ and suppose $h : U \rightarrow \mathbb{R}^{d-n}$ is a local defining function, where U is a neighborhood of p in $\mathcal{V}$. Then*

$$T_p\mathcal{M} = \ker Dh_p.$$

Proposition D.1.2. *Tangent spaces are vector spaces whose dimension coincides with the dimension of $\mathcal{M}$.*

The following two definitions gives us a way to characterize the topology of a smooth manifold. By this, we mean what subsets of a smooth manifold are considered open and closed. From here, one can easily figure out how to define notions of closure, interiors, boundaries, and so forth.

Definition D.1.3 (Smooth Manifold Topology). *A subset $U \subset \mathcal{M} \subset \mathcal{V}$ is called open (resp., closed) in $\mathcal{M}$ if there exists an open (resp., closed) subset $\bar{U}$ of $\mathcal{V}$ for which $U = \bar{U} \cap \mathcal{M}$.*

Definition D.1.4 (Neighborhood). *Let $U \subset \mathcal{M}$. A neighborhood of U in $\mathcal{M}$ is an open subset V in $\mathcal{M}$ which contains U . Similarly, a neighborhood of U in $\mathcal{V}$ is an open subset $\bar{V}$ in $\mathcal{V}$ which contains U .*

In the above definition, remark that while U is open in $\mathcal{M}$, it is certainly not open in $\mathcal{V}$. Thus, a neighborhood of p in $\mathcal{M}$ is certainly not one in $\mathcal{V}$.

D.2 Smooth Maps and Derivatives

A common tool we will utilize is that of a smooth extension of a function. Very often analytic constructs on a smooth manifold are simply the projection of an extension of that construct onto the manifold. This is useful for computational purposes, as we will later see.

Definition D.2.1 (Smooth Extension). *Let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. Let $p \in \mathcal{M}$ and $\bar{U}$ be a neighborhood of p in $\mathcal{V}$. Next, let $\bar{F} : \bar{U} \rightarrow \mathcal{V}'$ be a smooth function that coincides with F over $\mathcal{M} \cap \bar{U}$. We call $\bar{F}$ a (local) smooth extension of F about p to $\mathcal{V}$. In the case that $\bar{U}$ contains $\mathcal{M}$, we call $\bar{F}$ a smooth extension of F to $\mathcal{V}$.*

We now have the tools to define notions of smoothness.

Definition D.2.2 (Smooth Functions). *A mapping $F : \mathcal{M} \rightarrow \mathcal{M}'$ is called smooth at a point $p \in \mathcal{M}$ if there exists a local smooth extension $\bar{F}$ about p . When F is smooth for all $p \in \mathcal{M}$, we call F smooth. When F is smooth, invertible, and F^{-1} is also smooth, we call F a diffeomorphism. When F is a diffeomorphism onto its image, we call F a smooth embedding.*

A smooth function $f : \mathcal{M} \rightarrow \mathbb{R}$ is called a scalar field. We will denote the set of all smooth scalar fields $F : \mathcal{M} \rightarrow \mathbb{R}$ as $\mathfrak{F}(\mathcal{M})$.

Example D.2.1. *The inclusion mapping $\iota : \mathcal{M} \rightarrow \mathcal{V}$ is a smooth embedding.*

With smoothness comes the notion of a derivative. In differential geometry, we make use of the differential, a notion of a directional derivative we see in calculus.

Definition D.2.3. *Let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. The differential of F at $p \in \mathcal{M}$ is a linear map $dF_p : T_p\mathcal{M} \rightarrow T_{F(p)}\mathcal{M}'$ defined by*

$$dF_p(\eta) := \left. \frac{d}{dt} F(c(t)) \right|_{t=0},$$

where $\eta \in T_p\mathcal{M}$, and c is a smooth curve on $\mathcal{M}$ with $c(0) = p$ and $\dot{c}(0) = \eta$.

Theorem D.2.1. *Let $\mathcal{M}'$ be a smooth manifold and let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. The differential of F at any point is a linear mapping. Let $\bar{F}$ be a local smooth extension of F at a point $p \in \mathcal{M}$. Then dF_p is the restriction of the differential $D\bar{F}_p$ to $T_p\mathcal{M}$. That is, $dF_p = D\bar{F}_p|_{T_p\mathcal{M}}$.*

D.3 Vector Fields

Definition D.3.1. *The tangent bundle is defined as*

$$T\mathcal{M} := \{(p, \eta) : p \in \mathcal{M}, \eta \in T_p\mathcal{M}\}.$$

With some abuse of notation, we sometimes conflate the notions of η and (p, η) . Remark that tangent spaces at different points may intersect. The tangent spaces of a circle at two antipodal points coincide.

Proposition D.3.1. *The tangent bundle $T\mathcal{M}$ is a smooth manifold embedded in $\mathcal{V}^2$ of dimension $2 \dim(\mathcal{M})$.*

Definition D.3.2. *A vector field is a smooth map $V : \mathcal{M} \rightarrow T\mathcal{M}$ for which $V_p \in T_p\mathcal{M}$. The set of vector fields is denoted $\mathfrak{X}(\mathcal{M})$.*

D.4 Retraction

Definition D.4.1. A (local) retraction is a smooth mapping

$$R : \mathcal{S} \subset T\mathcal{M} \rightarrow \mathcal{M}, (p, \eta) \mapsto R_p(\eta)$$

such that

1. $\mathcal{S}$ is an open subset of $T\mathcal{M}$ and $(p, 0) \in \mathcal{S}$ for $p \in \mathcal{M}$
2. For $(p, \eta) \in \mathcal{S}$, the curve defined as $c(t) := R_p(t\eta)$ satisfies $c(0) = p$ and $\dot{c}(0) = \eta$

When $\mathcal{S} = T\mathcal{M}$, we call R a global retraction.

The second condition is equivalent to requiring $R_p(0) = p$ and

$$dR_p|_0 : T_p\mathcal{M} \rightarrow T_p\mathcal{M}$$

is the identity map: $dR_p|_0(\eta) = \eta$.

Theorem D.4.1. For $p \in \mathcal{M}$, the retraction R_p always admits a local inverse

$$R_p^{-1} : U \subset \mathcal{M} \rightarrow T_p\mathcal{M}$$

where U is some neighborhood of p in $\mathcal{M}$. We call R_p^{-1} the inverse retraction.

D.5 Riemannian Manifolds

Definition D.5.1. Suppose $\mathcal{V}$ is equipped with an inner product $\langle \cdot, \cdot \rangle$. For $p \in \mathcal{M}$, let $g_p : T_p\mathcal{M} \times T_p\mathcal{M} \rightarrow \mathbb{R}$ be defined as

$$g_p(\eta, \xi) := \langle \eta, \xi \rangle.$$

Then g_p is an inner product on $T_p\mathcal{M}$. We call g the Riemannian metric induced from $\mathcal{V}$. We call the pairing $(\mathcal{M}, g)$ the Riemannian manifold isometrically embedded in $\mathcal{V}$.

We use $g_p(\cdot, \cdot)$ notation instead of $\langle \cdot, \cdot \rangle$ notation to emphasize that g_p is the restriction of $\langle \cdot, \cdot \rangle$ to $T_p\mathcal{M}$.

D.6 Riemannian Gradients and Hessians

Definition D.6.1 (Riemannian Gradient). *Let $\mathcal{M}$ be a Riemannian manifold isometrically embedded in $\mathcal{V}$. Let $f \in \mathfrak{F}(\mathcal{M})$ be a scalar field. The Riemannian gradient of f is defined as*

$$\nabla f_p = \text{Proj}_{T_p\mathcal{M}}[\nabla \bar{f}_p],$$

where $\bar{f}$ is a smooth extension of f about $\mathcal{M}$ in $\mathcal{V}$, $\nabla \bar{f}_p \in \mathcal{V}$ is the classical gradient of $\bar{f}$ at p , and

$$\text{Proj}_{T_p\mathcal{M}} : \mathcal{V} \rightarrow T_p\mathcal{M}$$

is the orthogonal projection onto $T_p\mathcal{M}$.

Definition D.6.2. *Let $\mathcal{M}$ be isometrically embedded in $\mathcal{V}$. Consider $f \in \mathfrak{F}(\mathcal{M})$. Let $\bar{G}$ be a smooth extension of the Riemannian gradient vector field ∇f about $\mathcal{M}$ in $\mathcal{V}$. Then the Riemannian Hessian of f at p is a linear map $\nabla^2 f_p : T_p\mathcal{M} \rightarrow T_p\mathcal{M}$ defined as*

$$\nabla^2 f_p(\eta) := \text{Proj}_{T_p\mathcal{M}}[D\bar{G}_p(\eta)]$$

Proposition D.6.1. *The Riemannian Hessian $\nabla^2 f_p$ is self-adjoint with respect to the inner product g_p :*

$$g_p(\nabla^2 f_p(\eta), \xi) = g_p(\eta, \nabla^2 f_p[\xi]).$$

D.7 Local Frames and Components

Definition D.7.1. Let $U \subset \mathcal{M}$ be an open subset. Let (E_i) be n vector fields on $\mathcal{M}$ such that for $p \in U$, we have

$$\text{span}\{E_1|_p, \dots, E_n|_p\} = T_p\mathcal{M}.$$

That is, $(E_i|_p)$ forms a basis for $T_p\mathcal{M}$. We call (E_i) a (local) frame. When $U = \mathcal{M}$, (E_i) is called a global frame.

Suppose also $g_p(E_i|_p, E_j|_p) = \delta_{ij}$. Then we call (E_i) an orthonormal frame.

Definition D.7.2. Let V be a vector field on $\mathcal{M}$. The components of V are the unique set of smooth functions $V^i : \mathcal{M} \rightarrow \mathbb{R}$ satisfying

$$V_p = \sum_{i=1}^n V^i(p) E_i|_p.$$

We will write $[V_p]_i := V^i(p)$. Remark $[V_p] := (V^1(p), \dots, V^n(p)) \in \mathbb{R}^n$.

Proposition D.7.1. Let V be a vector field on $\mathcal{M}$. Suppose (E_i) is an orthonormal frame. Then

$$[V_p]_i = g_p(V_p, E_i|_p)$$

Definition D.7.3. The components of a Riemannian metric g are the unique set of smooth functions $g_{ij} : \mathcal{M} \rightarrow \mathbb{R}$ satisfying

$$g_p(V, W) = \sum_{i,j=1}^n g_{ij}(p) [V_p]_i [W_p]_j$$

for any vector fields V, W on $\mathcal{M}$. We will write $[g_p]_{ij} := g_{ij}(p)$. Remark $[g_p] \in \mathbb{R}^{n \times n}$.

Proposition D.7.2. For $\eta, \xi \in T_p\mathcal{M}$, we have

$$g_p(\eta, \xi) = [\eta]^\top [g_p] [\xi]$$

Proposition D.7.3. *We have*

$$[g_p]_{ij} = g_p(E_i|_p, E_j|_p).$$

Also, $[g_p]$ is always symmetric positive-definite. In the case that (E_i) is orthonormal, we have $[g_p]_{ij} = \delta_{ij}$. In the case that (E_i) is merely orthogonal, then $[g_p]$ is a diagonal matrix with positive diagonals.

Definition D.7.4. *Let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. Let $p \in \mathcal{M}$ and $p' := F(p)$. The components of dF are the unique set of smooth functions $f_{ij} : \mathcal{M} \rightarrow \mathbb{R}$ satisfying*

$$dF_p(\eta) = \sum_{i,j=1}^n f_{ij}(p) [\eta]_j E'_i|_{p'}$$

We will write $[dF_p]_{ij} := f_{ij}$. Remark $[dF_p] \in \mathbb{R}^{n' \times n}$.

Proposition D.7.4. *Let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. Then*

$$[dF_p(\eta)] = [dF_p][\eta] \in \mathbb{R}^{n'}.$$

Proposition D.7.5. *Let $F : \mathcal{M} \rightarrow \mathcal{M}'$ be smooth. Let (E'_i) be an orthonormal frame on $\mathcal{M}'$. Suppose (E_i) and (E'_i) are orthonormal frames on $\mathcal{M}$ and $\mathcal{M}'$. Let $p \in \mathcal{M}$ and $p' := F(p)$. Then*

$$[dF_p]_{ij} = g'_{p'}(dF_p(E_j|_p), E'_i|_{p'}).$$

Proposition D.7.6. *Let $f \in \mathfrak{F}(\mathcal{M})$. Then the components of the Riemannian gradient and differential of f have the following relation:*

$$[df_p] = [g_p][\nabla f_p] \iff [g_p]^{-1}[df_p] = [\nabla f_p]$$

Definition D.7.5. *Let $f \in \mathfrak{F}(\mathcal{M})$. The components of the Riemannian Hessian $\nabla^2 f_p : T_p \mathcal{M} \rightarrow T_p \mathcal{M}$ are the unique set of smooth functions $H_{ij} : \mathcal{M} \rightarrow \mathbb{R}$ satisfying*

$$\nabla^2 f_p(\eta) = H_{ij}(p) [\eta]_j E_i|_p.$$

We will write $[\nabla^2 f_p]_{ij} := H_{ij}(p)$. Take note that $[\nabla^2 f_p] \in \mathbb{R}^{n \times n}$.

Proposition D.7.7. *Let $f \in \mathfrak{F}(\mathcal{M})$ and $p \in \mathcal{M}$. Then*

$$[\nabla^2 f_p(\eta)] = [\nabla^2 f_p][\eta] \in \mathbb{R}^n$$

Proposition D.7.8. *Let $f \in \mathfrak{F}(\mathcal{M})$. Then while ∇f_p is a self-adjoint operator, the matrix $[\nabla^2 f_p] \in \mathbb{R}^{n \times n}$ is not symmetric. However, the matrix $[g_p][\nabla^2 f_p]$ is symmetric. This follows from*

$$\begin{aligned} g_p(\nabla^2 f_p(\eta), \xi) &= g_p(\eta, \nabla^2 f_p[\xi]) \\ \implies [\eta]^\top [\nabla^2 f_p]^\top [g_p][\xi] &= [\eta]^\top [g_p][\nabla^2 f_p][\xi]. \end{aligned}$$

In the case that (E_i) is orthonormal, $[\nabla^2 f_p]$ is symmetric.

Appendix E

FUN STUFF 1: AN ALGEBRAIC DEFINITION OF SMOOTH MANIFOLDS

The definition we gave for smooth manifolds in this work is inherently analytic. It begs to ask whether there is an algebraic way of defining smooth manifolds. It is a common sentiment within the mathematic communities that differential geometry is quite incompatible with algebra. The category of smooth manifolds is a particularly ugly one, which is the reason behind why orbit manifolds, infinite-dimensional manifolds, and submanifolds are so difficult to deal with, compared to say the straightforward categorical behavior of topological spaces or groups.

An algebraic definition of smooth manifolds can give us a satisfying answer to many questions on the path to mathematically capturing the concept of smoothness. First, what exactly is the structure of the set of smooth functions over a smooth manifold? What specifically about that structure is it that gives birth to a calculus framework, which allows optimization and local geometry? How exactly is it generalized from the structure of the set of smooth functions over $\mathbb{R}^n$? Can we take this structure and use it to generalize beyond smooth manifolds? Is there an explicit algebraic way of defining global vs. local properties?

With that said, an algebraic reformulation of smooth manifolds, given in [12], offers a compelling perspective by leveraging sheaf theory—a tool famously employed in algebraic geometry. In this framework, a smooth manifold is a topological manifold $\mathcal{M}$ equipped with a sheaf of functions that is locally sheaf isomorphic to the sheaf of smooth functions on $\mathbb{R}^n$.

Rings A **ring** is a set R equipped with two binary operations: addition $+$: $R \times R \rightarrow R$ and multiplication $\cdot$: $R \times R \rightarrow R$, satisfying the following axioms. For all $a, b, c \in R$,

1. $(a + b) + c = a + (b + c)$,
2. $a + b = b + a$,
3. there exists $0 \in R$ such that $a + 0 = a$,
4. for each $a \in R$, there exists $-a \in R$ such that $a + (-a) = 0$,
5. $(a \cdot b) \cdot c = a \cdot (b \cdot c)$,
6. $a \cdot (b + c) = a \cdot b + a \cdot c$,
7. $(a + b) \cdot c = a \cdot c + b \cdot c$.

The rings considered in this chapter will primarily be rings of scalar-valued functions $f : X \rightarrow \mathbb{R}$, with addition and multiplication defined pointwise.

Pre-Sheaves Let X be a topological space. A *pre-sheaf* $\mathcal{F}$ on X assigns to each open set $U \subset X$ a ring $\mathcal{F}(U)$ of scalar-valued functions $f : U \rightarrow \mathbb{R}$. Furthermore, associated with any pre-sheaf $\mathcal{F}$ is a family of restriction maps. For open sets $V \subset U \subset X$, the restriction mapping is defined by

$$\rho_U^V : \mathcal{F}(U) \rightarrow \mathcal{F}(V), \quad \rho_U^V : f \mapsto f|_V.$$

Remark each ρ_U^V is a ring homomorphism, ρ_U^U is the identity on $\mathcal{F}(U)$, and the compatibility condition $\rho_U^W = \rho_V^W \circ \rho_U^V$ holds whenever $W \subset V \subset U$.

Sheaves We have arrived at the central object of interest: a *sheaf*. A sheaf on X is a pre-sheaf $\mathcal{F}$ satisfying the following two axioms. For any open set $U \subset X$ and any open cover $\{U_i\}$ of U ,

1. **(Locality)** If $f, g \in \mathcal{F}(U)$ satisfy $\rho_U^{U_i}(f) = \rho_U^{U_i}(g)$ for all i , then $f = g$.
2. **(Gluing)** If for each i there exists $f_i \in \mathcal{F}(U_i)$ such that

$$\rho_{U_i}^{U_{ij}}(f_i) = \rho_{U_j}^{U_{ij}}(f_j) \quad \text{for all } i, j,$$

where $U_{ij} := U_i \cap U_j$, then there exists a unique $f \in \mathcal{F}(U)$ such that $\rho_U^{U_i}(f) = f_i$ for all i .

A pair $(X, \mathcal{F})$ consisting of a topological space and a sheaf of rings is called a *ringed space*.

Example On $\mathbb{R}^n$, the assignment $U \mapsto \mathfrak{F}(U)$ defines a sheaf: the sheaf of smooth scalar fields. Likewise, the assignment $U \mapsto C^0(U)$ defines the sheaf of continuous functions.

Non-example Let $\mathcal{F}(U)$ denote the ring of bounded functions $f : U \rightarrow \mathbb{R}$ for $U \subset \mathbb{R}$. This defines a pre-sheaf, but not a sheaf. Indeed, consider functions $f_i : (i-1, i+1) \rightarrow \mathbb{R}$ defined by $f_i(x) = x$. Each f_i is bounded on its domain, but gluing them together yields the globally defined function $f(x) = x$, which is unbounded. The gluing axiom therefore fails. A similar failure occurs for the pre-sheaf of uniformly continuous functions.

Informally, a pre-sheaf is a sheaf precisely when any compatible collection of local functions—finite or infinite—can be glued into a global function that still belongs to the same class. Boundedness fails this test; continuity and smoothness do not.

Smooth Manifolds as a Sheaf of Smooth Scalar Fields Let $\mathcal{S}$ denote the sheaf of smooth scalar fields over $\mathbb{R}^n$, and let $\mathcal{M}$ be a topological manifold of dimension n . Suppose $\mathcal{F}$ is a sheaf of scalar-valued functions on $\mathcal{M}$ such that for every open set $U \subset \mathcal{M}$, there exists a homeomorphism $\varphi : U \rightarrow V \subset \mathbb{R}^n$ for which $\mathcal{F}(U)$ is ring isomorphic to $\mathcal{S}(V)$. Then the pair $(\mathcal{M}, \mathcal{F})$ is called a *smooth manifold*.

Equivalently, a smooth manifold is a topological manifold equipped with a sheaf of scalar-valued functions that is locally isomorphic to the sheaf of smooth functions on $\mathbb{R}^n$. In this formulation, the sheaf $\mathcal{F}$ consists precisely of the functions we declare to be smooth on $\mathcal{M}$. Smoothness is therefore encoded algebraically, rather than through atlases and coordinate transitions.

This perspective makes clear what the “correct” algebraic object underlying smoothness actually is: a sheaf of rings. It is this structure that allows the machinery of calculus to be extended beyond Euclidean space, and it provides a clean separation between local analytic behavior and global geometric structure.

Appendix F

FUN STUFF 2: MOTIVATION FOR THE LEVI-CIVITA CONNECTION'S DEFINITION

These notes are adapted in part from [8, 10]. In this section, we will present four identities of the classical directional derivative on vector spaces which contain no instances of adding or subtracting tangent vectors from different base points. By doing so, we completely motivate the bizarre axiomatic definition of the Levi-Civita connection on Riemannian manifolds.

Let $\mathbb{R}^n$ be equipped with the dot product. $\langle \cdot, \cdot \rangle$. Here, we view $\mathbb{R}^n$ as a Riemannian manifold. Recall vector fields $U \in \mathfrak{X}(\mathbb{R}^n)$ may be identified with *geometric* vector fields

$$U : \mathbb{R}^n \rightarrow \mathbb{R}^n,$$

since each tangent space $T_p\mathbb{R}^n$ canonically identifies with $\mathbb{R}^n$ itself.

A key property of vector spaces is that the Levi-Civita connection

$$\nabla : \mathfrak{X}(\mathbb{R}^n) \times \mathfrak{X}(\mathbb{R}^n) \rightarrow \mathfrak{X}(\mathbb{R}^n), \quad (U, V) \mapsto \nabla_U V,$$

which coincides with the classical directional derivative, can be expressed in terms of the *differential*. That is,

$$\nabla_U V|_p = DV_p(U_p),$$

where the differential (also known as the Jacobian operator) of $V \in \mathfrak{X}(\mathbb{R}^n)$ at $p \in \mathbb{R}^n$ is the linear map

$$DV_p : \mathbb{R}^n \rightarrow \mathbb{R}^n, \quad DV_p(v) := \lim_{t \rightarrow 0} \frac{V_{p+tv} - V_p}{t}.$$

In summary, we may write

$$\nabla_U V|_p = DV_p(U_p) = \lim_{t \rightarrow 0} \frac{V_{p+tU_p} - V_p}{t}. \tag{F.1}$$

The unprincipled identity (F.1) subtracts tangent vectors belonging to two *different* tangent spaces: $T_{p+tU_p}\mathbb{R}^n$ and $T_p\mathbb{R}^n$. Such an operator is unavailable on a general Riemannian manifold.

We will introduce 4 identities of the directional derivative operator that have no such instances. For all $U, V, W \in \mathfrak{X}(\mathbb{R}^n)$ and $f \in \mathfrak{F}(\mathbb{R}^n)$:

1. ∇ is $\mathbb{R}$ -bilinear.
2. ∇ satisfies the product rule in its second argument. That is,

$$\nabla_U(fV)|_p = D(fV)_p(U_p) = Df_p(U_p)V_p + f(p)DV_p(U_p).$$

3. U distributes over the dot product. Defining

$$h(p) := \langle V_p, W_p \rangle,$$

we have

$$(Uh)(p) = Dh_p(U_p) = \langle DV_p(U_p), W_p \rangle + \langle V_p, DW_p(U_p) \rangle.$$

4. We have the identity relating the commutativity of the directional derivative operator:

$$\nabla_U V|_p - \nabla_V U|_p = DV_p(U_p) - DU_p(V_p) =: [U, V]_p.$$

Again, since none of these four identities rely on the addition or subtraction of tangent vectors based at different points of $\mathbb{R}^n$, each remains meaningful on a general Riemannian manifold, where no canonical identification exists between distinct tangent spaces. Remarkably, these four properties uniquely characterize the directional derivative on vector spaces:

Lemma F.0.1. *Any operator*

$$\nabla : \mathfrak{X}(\mathbb{R}^n) \times \mathfrak{X}(\mathbb{R}^n) \rightarrow \mathfrak{X}(\mathbb{R}^n)$$

satisfying the four properties above coincides with the ordinary directional derivative.

This is none other than a manifestation of the fundamental theorem of Riemannian geometry and explains why on a general Riemannian manifold there exists a unique connection characterized by these four properties.

We conclude this dissertation by a restatement of my favorite theorem in mathematics:

Theorem F.0.1 (Fundamental Theorem of Riemannian Geometry). *Let $(\mathcal{M}, g)$ be a Riemannian manifold. There exists a unique mapping*

$$\nabla : \mathfrak{X}(\mathcal{M}) \times \mathfrak{X}(\mathcal{M}) \rightarrow \mathfrak{X}(\mathcal{M}), \quad (U, V) \mapsto \nabla_U V,$$

such that for all $U, V, W \in \mathfrak{X}(\mathcal{M})$ and all $f \in \mathfrak{F}(\mathcal{M})$,

∇ is $\mathbb{R}$ -bilinear,

$$\nabla_U(fV) = (Uf)V + f\nabla_U V,$$

$$Ug(V, W) = g(\nabla_U V, W) + g(V, \nabla_U W),$$

$$\nabla_U V - \nabla_V U = [U, V].$$

The unique connection satisfying these properties is called the Levi–Civita connection of $(\mathcal{M}, g)$.