Output-feedback Synthesis Orbit Geometry: Quotient Manifolds and LQG Direct Policy Optimization

Spencer Kraisler and Mehran Mesbahi

Abstract— **In this paper, we consider direct policy optimization for the linear-quadratic Gaussian (LQG) setting. Over the past few years, it has been recognized that the landscape of stabilizing output-feedback controllers of relevance to LQG has an intricate geometry, particularly as it pertains to the existence of spurious stationary points. In order to address such challenges, in this paper, we first adopt a Riemannian metric for the space of stabilizing full-order minimal output-feedback controllers. We then proceed to prove that the orbit space of such controllers modulo coordinate transformation admits a Riemannian quotient manifold structure. This geometric structure is then used to develop a Riemannian gradient descent for the direct LQG policy optimization. We prove a local convergence guarantee with linear rate and show the proposed approach exhibits significantly faster and more robust numerical performance as compared with ordinary gradient descent for LQG. Subsequently, we provide reasons for this observed behavior; in particular, we argue that optimizing over the orbit space of controllers is the right theoretical and computational setup for direct LQG policy optimization.**

I. Introduction

Direct policy optimization (PO) synthesizes controllers by formalizing constrained optimization problems over controller parameters rather than solving for value functions or Lyapunov certificates using matrix inequalities. In recent years, PO has been shown to be an effective first order procedure for a number of feedback synthesis problems, while providing a natural bridge between control synthesis and data-driven methods and learning.

In the PO setting, design problems such as LQR, LQG, H_∞ subject to H_2 or H_∞ constraints [1], [2], are first formalized in terms of the corresponding feedback parameters; subsequently, some variants of first order methods are adopted to update these parameters with the goal of at least a local convergence guarantees. Such a "direct" synthesis procedure has let to the need for a deeper analysis of the control objectives in relation to the space of controllers [3], [4]. A primary example of such a perspective is the observation that when LQR is written directly in terms of the control parameters, it is gradient-dominant, allowing gradient descent to have a linear rate of convergence for its solution.

In the context of PO–as it turns out–LQG has a more intricate landscape as compared with the LQR [5], hindering a straight forward adoption of first order methods for LQR for its solution. In fact, to our knowledge, there are no local convergence guarantees for LQG. There are a number of reasons for this. First, the domain of output-feedback controllers similar to the case of state feedback are non-convex. However, as opposed to the LQR case, the LQG cost admits strict and non-strict saddle points. Furthermore, the coordinate-invariance of LQG implies that each stationary point lies within an orbit of stationary points. Moreover, some systems admit degenerate LQG controllers, hence greatly impacting the convergence rate for optimization algorithms. In fact, the LQG controller can be non-minimal and have small stability margins [6]. Last, the search space of full-order controllers is large with $n^2+nm+np$ dimensions, where n, m, and p are the dimensions of the state, control, and output, respectively.

In this paper, we present a geometric approach for resolving many of these issues; the key ingredient is framing PO for LQG over the *Riemannian quotient manifold* of full-order minimal controllers modulo coordinate transformation. Equipping a search space with a Riemannian metric for optimization is a popular technique in controls and machine learning [7]–[9]. We prove this setup is well-defined and show how to perform Riemannian gradient descent (RGD). We show this technique is far faster than GD and performs well at avoiding saddle points, and we present a proof of local convergence.

Although PO for control synthesis is a relatively recent research direction, our work benefits from tools historically developed in geometric system theory pertaining to orbit spaces of linear systems. In fact, examining such orbits was initiated by Kalman and Hazewinkel in the early 1970s [10], [11] for system identification. In this paper, we show that these tools are rather powerful for PO and data driven control since the set of output-feedback controllers is *a family* of linear systems. In the meantime, optimization over the geometry induced by orbits of linear (dynamic) controllers comes hand in hand with a number of technical issues that are addressed in this work. For a brief modern survey of geometric system theory, we recommend [12], as well as the older references [13], [14] that examine a broader set topics in this discipline.

The outline of the paper is as follows. In §II, we introduce notation and mathematical background. In §III, we present the main algorithm, its implementation details, and discuss why it has superior performance over gradient descent (GD). In §IV, we delve into the theoretical setup for the proposed algorithm followed by its convergence analysis in §V. In §VI, we compare the performance of our algorithm with ordinary GD adopted for LQG PO for various examples. Lastly, §VII, future directions for research are discussed.

II. Preliminaries

Consider the continuous-time linear system,

$$\dot{x}(t) = Ax(t) + Bu(t) + w(t),\ y(t) = Cx(t) + v(t), \quad (1)$$

The authors are with the William E. Boeing Department of Aeronautics and Astronautics, University of Washington. Emails: {kraisler+mesbahi}@uw.edu

as our plant model, where the process $w(.)$ and measurement $v(.)$ noise terms are zero-mean Gaussian with covariance matrices $W \in \mathcal{S}_n^+$ (positive semidefinite) and $V \in \mathcal{S}_p^{++}$ (positive definite), respectively. We also assume (A, B) and $(A, W^{1/2})$ are controllable and (A, C) is observable. An output-feedback controller of order $1 \leq q \leq n$ for this plant is now parameterized as,

$$\dot{\xi}(t) = A_K\xi(t) + B_K y(t),\ u(t) = C_K\xi(t), \tag{2}$$

where $A_K \in \mathbb{R}^{q\times q}$, $B_K \in \mathbb{R}^{q\times p}$ and $C_K \in \mathbb{R}^{m\times q}$. Let $\tilde{\mathcal{C}}_q$ be the set of all such qth-order output-feedback controllers, represented as

$$K = \begin{bmatrix} 0_{m\times p} & C_K \\ B_K & A_K \end{bmatrix} \in \mathbb{R}^{(m+q)\times(p+q)} \tag{3}$$

Now, let $J_q : \tilde{\mathcal{C}}_q \to \mathbb{R}$ be the LQG cost for qth-order controllers, with $Q \in \mathcal{S}_n^+$ and $R \in \mathcal{S}_m^{++}$ as the state and control cost matrices. We assume that $(A, Q^{1/2})$ is observable.

In order to examine output-feedback synthesis, we first consider the combined plant/controller closed loop as,

$$\begin{bmatrix} \dot{x} \\ \dot{\xi} \end{bmatrix} = \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} I_n & 0_{n\times p} \\ 0_{n\times n} & B_K \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix} \tag{4a}$$

$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} C & 0_{p\times q} \\ 0_{m\times n} & C_K \end{bmatrix} \begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} 0_{p\times n} & I_p \\ 0_{n\times n} & 0_{n\times p} \end{bmatrix} \begin{bmatrix} w \\ v \end{bmatrix}. \tag{4b}$$

The realized closed-loop system and observation matrices are now, respectively,

$$A_{\mathrm{cl}}(K) \in \mathbb{R}^{(n+q)\times(n+q)}, \qquad B_{\mathrm{cl}}(K) \in \mathbb{R}^{(n+q)\times(n+p)},$$
$$C_{\mathrm{cl}}(K) \in \mathbb{R}^{(m+p)\times(n+q)}, \qquad D_{\mathrm{cl}}(K) \in \mathbb{R}^{(m+p)\times(n+p)}.$$

Hence, (2) is stabilizing when $A_{\mathrm{cl}}(K) \in \mathcal{H}_{n+q}$, where $\mathcal{H}_k$ is the set of $k \times k$ Hurwitz stable matrices.

Let $\tilde{\mathcal{C}}_q^{\min}$ be the set of minimal (i.e., controllable and observable) qth-order output-feedback controllers. Our first observation is as follows.

Lemma 2.1: $\tilde{\mathcal{C}}_q^{\min} \underset{\mathrm{op}}{\subset} \tilde{\mathcal{C}}_q$ and $\tilde{\mathcal{C}}_q^{\min}$ is a generic subset of $\tilde{\mathcal{C}}_q$.

We will see an elegant interaction between $\tilde{\mathcal{C}}_q^{\min}$ and $\tilde{\mathcal{C}}_q$ in the context of PO subsequently. A key construct for our geometric approach is the Lyapunov operator $\mathbb{L}(A, Q)$, mapping $A \in \mathcal{H}_k$ and $Q \in \mathcal{S}_k^+$ to the unique solution of $AP + PA^\intercal = -Q$. In fact, defining the maps

$$Q_{\mathrm{cl}}(K) := \begin{bmatrix} Q & 0_{n\times q} \\ 0_{q\times n} & C_K^\intercal R C_K \end{bmatrix}, \tag{5a}$$

$$W_{\mathrm{cl}}(K) := \begin{bmatrix} W & 0_{n\times q} \\ 0_{q\times n} & B_K V B_K^\intercal \end{bmatrix}, \tag{5b}$$

$$X(K) := \mathbb{L}\left(A_{\mathrm{cl}}(K), W_{\mathrm{cl}}(K)\right), \tag{5c}$$

on $\tilde{\mathcal{C}}_q$, facilitate recognizing that $J_q(K) = \mathrm{tr}\left(Q_{\mathrm{cl}}(K)X(K)\right)$ is the LQG cost. We note that the Euclidean gradient and Hessian of J_q have been computed in [5].

A. Riemannian Geometry

We recommend [15], [16] for optimization-oriented references to Riemannian manifolds. More comprehensive resources on Riemannian geometry can be found in [17]–[19].

Let $\mathcal{M} \subset \mathbb{R}^N$ be a smooth manifold. A smooth curve is a smooth function $c : \mathbb{R} \to \mathcal{M}$. The tangent space at x, denoted as $T_x\mathcal{M}$, is the set of the tangent vectors $\dot{c}(0)$ of all smooth curves $c(\cdot)$ with $c(0) = x$. For example, as an Euclidean open set, the tangent spaces of $\tilde{\mathcal{C}}_q$ identifies as

$$T_K\tilde{\mathcal{C}}_q \equiv \mathcal{V}_q := \left\{ \begin{bmatrix} 0_{m\times p} & G \\ F & E \end{bmatrix} \in \mathbb{R}^{(m+q)\times(p+q)} \right\}.$$

For matrix manifolds, tangent vectors are matrices. We use boldface letters, e.g., $\mathbf{V}$, to denote such. The disjoint union of tangent spaces is called the tangent bundle, written as $T\mathcal{M}$.

Let $F : \mathcal{M} \to \mathcal{N}$ be a smooth function between two smooth manifolds $\mathcal{M}$ and $\mathcal{N}$. The differential $dF_x : T_x\mathcal{M} \to T_{F(x)}\mathcal{N}$ of F at x along $v \in T_x\mathcal{M}$ is the linear mapping defined as

$$dF_x(v) := \frac{d}{dt}\Big|_{t=0} (F \circ c)(t),$$

where $c(\cdot)$ is any smooth curve satisfying $c(0) = x$ and $\dot{c}(0) = v$. For example, the differential of the Lyapunov operator $\mathbb{L}$ is,

$$d\mathbb{L}_{(A,Q)}(\mathbf{V}, \mathbf{W}) = \mathbb{L}(A, \mathbf{V}\mathbb{L}(A, Q) + \mathbb{L}(A, Q)\mathbf{V}^\intercal + \mathbf{W}).$$

When the differential $dF_x(.)$ is independent of x, we will drop $|_x$ and simply write $dF(.)$.

A Riemannian metric is a smooth state-dependent inner product $\langle ., . \rangle_x : T_x\mathcal{M} \times T_x\mathcal{M} \to \mathbb{R}$. Given open $U \subset \mathcal{M}$, a local frame is a set of linearly independent vector fields $(\mathbf{E}_i : U \to T\mathcal{M})_{i=1}^{\dim \mathcal{M}}$. The coordinates $G(x) \in \mathcal{S}_n^{++}$ of the metric at $x \in U$ are

$$G_{ij}(x) = \langle \mathbf{E}_i|_x, \mathbf{E}_j|_x \rangle_x. \tag{6}$$

With $G^{ij} := (G^{-1})_{ij}$, the gradient of $f : \mathcal{M} \to \mathbb{R}$ at x is,

$$\nabla f(x) = \sum_{i=1}^{n}\sum_{j=1}^{n} G^{ij}(x) df_x(\mathbf{E}_i|_x)\mathbf{E}_j|_x. \tag{7}$$

We will denote the Euclidean gradient as $\mathrm{grad} f$ and the Riemannian gradient as ∇f; similarly, Hess f and $\nabla^2 f$ for the corresponding Hessians.

A retraction is a smooth mapping $\mathcal{R} : \mathcal{S} \subset T\mathcal{M} \to \mathcal{M}$ where $\mathcal{S}$ is open, $(x, 0_x) \in \mathcal{S}$ for all x, and the curve $c(t) := \mathcal{R}_x(tv) \equiv \mathcal{R}(x, tv)$ satisfies $c(0) = x$ and $\dot{c}(0) = v$ for each $(x, v) \in \mathcal{S}$. Retractions are the central constructs in Riemannian optimization. When $\mathcal{M} \subset \mathcal{V}_n$, one can use the Euclidean metric and retraction:

$$\langle \mathbf{V}, \mathbf{W} \rangle_K := \mathrm{tr}(\mathbf{V}^\intercal\mathbf{W}), \quad \mathcal{R}_K(\mathbf{V}) := K + \mathbf{V}. \tag{8}$$

When $\mathcal{S} = T\mathcal{M}$, we enjoy an abundance of local convergence guarantees for first-order optimizers. Care must be taken for convergence proofs when $\mathcal{S} \subsetneq T\mathcal{M}$ [20, Remark 2.2].

With these basic ingredients of Riemannian optimization in place, the Riemannian Gradient Descent (RGD) of f under $(\langle ., . \rangle, \mathcal{R})$ is defined as,

$$x_{t+1} := \mathcal{R}_{x_t}(-s_t\nabla f(x_t)), \tag{9}$$

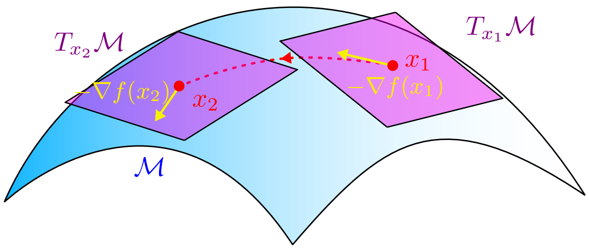

Fig. 1. Visualization of RGD. Here, $x_2 = \mathcal{R}_{x_1}(-s_1\nabla f(x_1))$

where $s_t \geq 0$ is a chosen step-size. Figure 1 is a visual depiction of the RGD procedure.

III. Our Algorithm

In this section, we introduce our proposed first-order optimizer for LQG. Algorithm 1 is RGD over the domain of $\tilde{\mathcal{C}}_n^{\min}$. The optimizer uses the Euclidean retraction and the Riemannian metric defined in §III-A. We optimize over $\tilde{\mathcal{C}}_n^{\min}$ instead of $\tilde{\mathcal{C}}_n$ for two reasons. First, the orbit space of $\tilde{\mathcal{C}}_n^{\min}$ modulo coordinate transformation admits a *quotient manifold structure*. This is not the case for $\tilde{\mathcal{C}}_n$. Second, the metric is coordinate-invariant, and so RGD over the $(n^2+nm+np)$-dimensional Riemannian manifold $\tilde{\mathcal{C}}_n^{\min}$ coincides with RGD over the much smaller $(nm+np)$-dimensional *Riemannian quotient manifold*. This will be explained in detail in §IV.

Algorithm 1 Riemannian Gradient Descent

Require: $K_0 \in \tilde{\mathcal{C}}_n^{\min}$, $\epsilon > 0$, $T \in \mathbb{N}$, $s_t \geq 0$
  $K \leftarrow K_0$, $t \leftarrow 0$
  while $t \leq T$ **and** $\|\nabla J_n(K)\|_K \geq \epsilon$ **do**
    $K \leftarrow K - s_t \nabla J_n(K)$
    $t \leftarrow t+1$
  return K

A few remarks are in order. In the case where $K^+ = \mathcal{R}_K(-s\nabla J_n(K))$ is non-stabilizing, one has to choose a small enough s_t in Algorithm 1; analogously, when K^+ is non-minimal, one can perturb the step direction (Lemma 2.1). Next, $\nabla J_n(K)$ is computed via (6) and (7) in each iteration[1]. In this context, the global frame $(\mathbf{E}_i)_{i=1}^N$, where $N = n^2+nm+np$, can simply be a fixed basis of $\mathcal{V}_n$. Lastly, the differential of the LQG cost is,

$$dJ_n|_K(\mathbf{V}) = \mathrm{tr}\left(dQ_{\mathrm{cl}}|_K(\mathbf{V})X(K) + Q_{\mathrm{cl}}(K)dX_K(\mathbf{V})\right),$$

where

$$dQ_{\mathrm{cl}}|_K(\mathbf{V}) = \begin{bmatrix} 0_{n\times n} & 0_{n\times n} \\ 0_{n\times n} & G^\intercal R C_K + C_K^\intercal R G \end{bmatrix}$$

$$dW_{\mathrm{cl}}|_K(\mathbf{V}) = \begin{bmatrix} 0_{n\times n} & 0_{n\times n} \\ 0_{n\times n} & F V B_K^\intercal + B_K V F^\intercal \end{bmatrix}$$

$$dX_K(\mathbf{V}) = d\mathbb{L}_{(A_{\mathrm{cl}}(K),W_{\mathrm{cl}}(K))}(dA_{\mathrm{cl}}(\mathbf{V}), dW_{\mathrm{cl}}|_K(\mathbf{V})).$$

A. Krishnaprasad-Martin Metric

Let $K \in \tilde{\mathcal{C}}_n^{\min}$ and $\mathbf{V} \in \mathcal{V}_n$ be a tangent vector. Define the mappings

$$\hat{\mathbf{E}}(\mathbf{V}) := dA_{\mathrm{cl}}(\mathbf{V}) = \begin{bmatrix} 0_{n\times n} & BG \\ FC & E \end{bmatrix}$$

[1]One can use Cholesky decomposition to reduce per-iteration time.

$$\hat{\mathbf{F}}(\mathbf{V}) := dB_{\mathrm{cl}}(\mathbf{V}) = \begin{bmatrix} 0_{n\times n} & 0_{n\times p} \\ 0_{n\times n} & F \end{bmatrix}$$

$$\hat{\mathbf{G}}(\mathbf{V}) := dC_{\mathrm{cl}}(\mathbf{V}) = \begin{bmatrix} 0_{p\times n} & 0_{p\times n} \\ 0_{m\times n} & G \end{bmatrix}.$$

Next, let $\mathcal{W}_c(K)$ and $\mathcal{W}_o(K)$ denote the controllability and observability Grammians of $(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K), C_{\mathrm{cl}}(K))$. Consider now the following Riemannian metric,

$$\langle \mathbf{V}_1, \mathbf{V}_2 \rangle_K^{\mathrm{KM}} := w_1 \mathrm{tr}[\mathcal{W}_o(K)\cdot\hat{\mathbf{E}}(\mathbf{V}_1)\cdot\mathcal{W}_c(K)\cdot\hat{\mathbf{E}}(\mathbf{V}_2)^\intercal] \tag{10a}$$
$$+ w_2 \mathrm{tr}[\hat{\mathbf{F}}(\mathbf{V}_1)^\intercal\cdot\mathcal{W}_o(K)\cdot\hat{\mathbf{F}}(\mathbf{V}_2)] \tag{10b}$$
$$+ w_3 \mathrm{tr}[\hat{\mathbf{G}}(\mathbf{V}_1)\cdot\mathcal{W}_c(K)\cdot\hat{\mathbf{G}}(\mathbf{V}_2)^\intercal], \tag{10c}$$

where $w_1 > 0$, $w_2, w_3 \geq 0$ are constants.

This metric was derived from a similar setup in [21], [22]. In literature, the original metric is called the Krishnaprasad-Martin (KM) metric [12]. Although we have slightly augmented the original metric, we will keep the original terminology.

B. Limitation of gradient descent on LQG landscape

For $S \in \mathrm{GL}(q)$ (invertible $q\times q$ matrices), define the diffeomorphism $\mathcal{T}_S : \tilde{\mathcal{C}}_n \to \tilde{\mathcal{C}}_n$:

$$\mathcal{T}_S(K) = \begin{bmatrix} 0_{m\times p} & C_K S^{-1} \\ SB_K & SA_K S^{-1} \end{bmatrix}. \tag{11}$$

We call (11) a coordinate transformation. Abusing notation, we have $d\mathcal{T}_S(\mathbf{V}) = \mathcal{T}_S(\mathbf{V})$ since (11) is linear.

A function $\mathcal{F}$ on $\tilde{\mathcal{C}}_q \times \mathcal{V}_q$ is called coordinate-*invariant* when,

$$\mathcal{F}(K, \mathbf{V}) = \mathcal{F}(\mathcal{T}_S(K), d\mathcal{T}_S(\mathbf{V})). \tag{12}$$

A function $\mathcal{F} : \tilde{\mathcal{C}}_q \times \mathcal{V}_n \to \mathcal{V}_n$ is coordinate-*equivariant* if

$$\mathcal{F}(\mathcal{T}_S(K), d\mathcal{T}_S(\mathbf{V})) = \mathcal{T}_S(\mathcal{F}(K, \mathbf{V})).$$

We can now discuss the limitations of GD over J_n. Take note that the GD procedure lacks coordinate-equivariance. This is due to the fact that the Euclidean metric, despite its simplicity, is not coordinate-invariant. As such, GD has to search through $\dim(\mathrm{GL}(n)) = n^2$ redundant dimensions. Furthermore, if we initialize $K_0 \in \tilde{\mathcal{C}}_n$ with particularly "bad" coordinates, one ends up with a large value of $\|K_0\|_F$; in that case, $\|\mathrm{grad}J_n(K)\|_F$ will also be large. Such ill-conditioned coordinates in turn result in numerical instabilities. These two issues are resolved however when the metric is coordinate-invariant and retraction is coordinate-equivariant.

C. Coordinate-invariance of the KM Metric

In this section, we will prove the above mentioned properties of the KM metric.

Lemma 3.1: For the system (A, B, C) in (1) and $K \in \tilde{\mathcal{C}}_q^{\min}$, we have $(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K), C_{\mathrm{cl}}(K))$ is minimal.

Proof: The proof involves using the Popov-Belevitch-Hautus (PBH) test to verify the controllability and observability of $(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K), C_{\mathrm{cl}}(K))$ and omitted for brevity. ■

Theorem 3.2: The mapping defined in (10) is a Riemannian metric and coordinate-invariant.

Proof: It is clear that for any K, the mapping $\langle .,. \rangle_K^{\mathrm{KM}}$ is smooth, bi-linear, and symmetric. It thus suffices to show that this map is positive-definite. Since $A_{\mathrm{cl}}(K) \in \mathcal{H}_{2n}$, it follows that,

$$\mathcal{W}_c(K) = \mathbb{L}(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K)B_{\mathrm{cl}}(K)^\intercal), \tag{13a}$$
$$\mathcal{W}_o(K) = \mathbb{L}(A_{\mathrm{cl}}(K)^\intercal, C_{\mathrm{cl}}(K)^\intercal C_{\mathrm{cl}}(K)). \tag{13b}$$

Since the closed-loop system $(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K), C_{\mathrm{cl}}(K))$ is minimal, we have $\mathcal{W}_c(K), \mathcal{W}_o(K) \in \mathcal{S}_{2n}^{++}$. Hence, (10) is positive-definite.

Now, we will show that the KM metric is coordinate-invariant. Let $S \in \mathrm{GL}(n)$ and $L := \mathcal{T}_S(K)$. Then

$$(A_{\mathrm{cl}}(L), B_{\mathrm{cl}}(L), C_{\mathrm{cl}}(L)) = \mathcal{T}_{\hat{S}}(A_{\mathrm{cl}}(K), B_{\mathrm{cl}}(K), C_{\mathrm{cl}}(K)),$$

where $\hat{S} := \begin{bmatrix} I_n & 0_{n\times n} \\ 0_{n\times n} & S \end{bmatrix}$. It follows that,

$$\begin{aligned} A_{\mathrm{cl}}(L) &= \hat{S} A_{\mathrm{cl}}(K) \hat{S}^{-1} \\ B_{\mathrm{cl}}(L)B_{\mathrm{cl}}(L)^\intercal &= \hat{S} B_{\mathrm{cl}}(K) B_{\mathrm{cl}}(K) \hat{S}^\intercal \\ C_{\mathrm{cl}}(L)^\intercal C_{\mathrm{cl}}(L) &= \hat{S}^{-\intercal} C_{\mathrm{cl}}(K)^\intercal C_{\mathrm{cl}}(K) \hat{S}^{-1} \end{aligned}$$

Thereby,

$$\mathcal{W}_c(L) = \hat{S}\mathcal{W}_c(K)\hat{S}^\intercal, \quad \mathcal{W}_o(L) = \hat{S}^{-\intercal}\mathcal{W}_o(K)\hat{S}^{-1}, \tag{14}$$

where we have used coordinate transformation property of the Lyapunov operator. We also have,

$$\hat{\mathbf{E}}(d\mathcal{T}_S(\mathbf{V}_i)) = \hat{S}\hat{\mathbf{E}}(\mathbf{V}_i)\hat{S}^{-1} \tag{15a}$$
$$\hat{\mathbf{F}}(d\mathcal{T}_S(\mathbf{V}_i)) = \hat{S}\hat{\mathbf{F}}(\mathbf{V}_i) \tag{15b}$$
$$\hat{\mathbf{G}}(d\mathcal{T}_S(\mathbf{V}_i)) = \hat{\mathbf{G}}(\mathbf{V}_i)\hat{S}^{-1}; \tag{15c}$$

now, plug (14) and (15) into (10) to conclude the proof. ■

IV. Orbit Space of Output-feedback Controllers

In this section, we go over key features for Riemannian quotient manifolds of controllers. We suggest referring to [15], [16] for the salient features of such a construction; we highlight some of the key points of quotient manifolds next.

Let $\tilde{\mathcal{M}}$ be a smooth manifold with group action G. For example, the family of coordinate transformations $\{\mathcal{T}_S(.) : S \in \mathrm{GL}(q)\} \equiv \mathrm{GL}(q)$ is a group action over $\tilde{\mathcal{C}}_q$. The orbit space of $\tilde{\mathcal{M}}$ modulo G is the collection of all orbits under the quotient topology:

$$\mathcal{M} \equiv \tilde{\mathcal{M}}/G := \{[x] : x \in \tilde{\mathcal{M}}\}. \tag{16}$$

Here, $[x] = \{y \in \tilde{\mathcal{M}} : \exists g \in G, g(x) = y\}$ is the orbit of x. See Figure 2 for a visual depiction. We say $U \subset \mathcal{M}$ is G-stable if $x \in U$ implies $[x] \subset U$. If $\mathcal{M}$ is a quotient manifold, then $\dim(\mathcal{M}) = \dim(\tilde{\mathcal{M}}) - \dim(G)$; this is particularly desirable for optimization due to the reduced dimension of the quotient manifold.

Let $\mathcal{V}_x := \ker d\pi_x$ be the tangent space of $[x]$ at x. Here, $d\pi_x$ is the differential of the quotient map $\pi(x) := [x]$. Next, let[2] $\mathcal{H}_x := \mathcal{V}_x^\perp$. Since $d\pi_x|_{\mathcal{H}_x} : \mathcal{H}_x \to T_{[x]}\mathcal{M}$ is a bijection, we identify $\xi \in T_{[x]}\mathcal{M}$ with $\mathrm{lift}_x(\xi) := (d\pi_x|_{\mathcal{H}_x})^{-1}(\xi) \in \mathcal{H}_x$.

Convergence analysis of our proposed RGD procedure involves showing $\mathcal{C}_n^{\min} := \tilde{\mathcal{C}}_n^{\min}/\mathrm{GL}(n)$ is a smooth quotient

[2]The orthogonal complement is taken with respect to the metric, not the dot product. With $\tilde{\mathcal{C}}_n^{\min}$, it does not coincide with [5, Prop. 4.2].

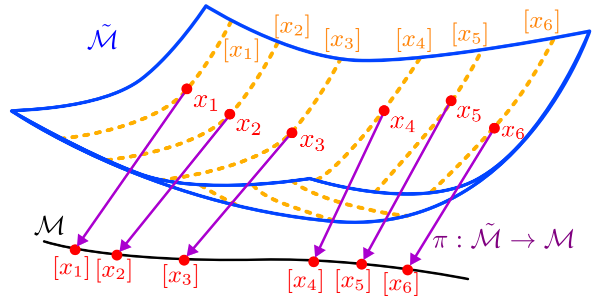

Fig. 2. Illustration of a manifold and its orbit space

manifold and inherits a Riemannian metric and retraction from the KM metric and Euclidean retraction. We state these results and omit their proofs for brevity.

The quotient manifold structure of $\mathcal{C}_n^{\min}$ follows from the remarkable theorem that the quotient space of minimal system realizations $\mathcal{L}_{n,p,m}^{\min}$ forms a smooth quotient manifold [13]. Since $\mathcal{C}_n^{\min}$ is an open subset $\mathcal{L}_{n,p,m}^{\min}$, then our claim is a consequence of the following lemma:

Lemma 4.1: If $\mathcal{M} := \tilde{\mathcal{M}}/G$ is a quotient manifold and $\tilde{\mathcal{N}} \underset{\text{op}}{\subset} \tilde{\mathcal{M}}$ is G-stable, then $\mathcal{N} := \tilde{\mathcal{N}}/G$ is a quotient manifold.

The inherited Riemannian structure on $\mathcal{C}_n^{\min}$ is a result of the invariance properties of the KM metric and retraction. In particular,

$$\langle \xi, \eta \rangle_{[K]}^{\mathrm{KM}} := \langle \mathrm{lift}_K(\xi), \mathrm{lift}_K(\eta) \rangle_K^{\mathrm{KM}} \tag{17}$$
$$\mathcal{R}_{[K]}(\xi) := [\mathcal{R}_K(\mathrm{lift}_K(\xi))], \tag{18}$$

where $K \in \tilde{\mathcal{C}}_n^{\min}$ and $\xi, \eta \in T_{[K]}\mathcal{C}_n^{\min}$. This is the result of the following lemma in the general setting:

Lemma 4.2: Suppose $\tilde{\mathcal{M}}$ is equipped with a G-invariant Riemannian metric. Pick $x \in \tilde{\mathcal{M}}$, $g \in G$, and $\xi \in T_{[x]}\mathcal{M}$. Then $dg_x(\mathrm{lift}_x(\xi)) = \mathrm{lift}_{g(x)}(\xi)$.

Intuitively, this also implies that performing RGD over the higher-dimensional $\tilde{\mathcal{C}}_n^{\min}$ coincides with its performance over the lower-dimensional $\mathcal{C}_n^{\min}$ [15, Sect. 9.9].

V. Convergence Analysis

In this section, we conduct a convergence analysis for our algorithm. Let J_n and $\tilde{J}_n$ denote the LQG cost over $\mathcal{C}_n^{\min}$ and $\tilde{\mathcal{C}}_n^{\min}$, respectively. We make the following assumption on the non-degeneracy of the LQG controller. We found empirically that this property is generic among minimal LQG controllers.

Assumption 5.1: The cost $\tilde{J}_n$ admits a (minimal) global minimum K^* on $\tilde{\mathcal{C}}_n^{\min}$ with $\mathrm{null}(\mathrm{Hess}\ \tilde{J}_n(K^*)) = \mathrm{null}(d\pi_{K^*})$.

We now present the proof of guaranteed local convergence of Algorithm 1 to optimality with a linear rate.

Theorem 5.2: There exists a neighborhood $\mathcal{U} \subset \mathcal{C}_n^{\min}$ of $[K^*]$ and $L > 0$ such that given $[K_0] \in \mathcal{U}$, the resulting sequence $([K_t])_{t\geq 0}$ via $K^+ = \tilde{F}(K) := \mathcal{R}_K(-\frac{1}{L}\nabla \tilde{J}_n(K))$ stays in $\mathcal{U}$ and converges to $[K^*]$ at least linearly.

The key idea for this analysis is to build up the two conditions to execute [15, Thm. 4.19]. The first condition requires that the LQG controller is a non-degenerate global minimum (Assumption 5.1). Non-degeneracy is a corollary of Lemma 5.3, which establishes a relationship between the Euclidean and Riemannian Hessians for LQG.

The second condition requires constructing a domain $\mathcal{L}_0 \subset \mathcal{C}_n^{\min}$ on which our RGD procedure $F(.)$ is well-defined and invariant. To ensure $F(.)$ is well-defined, since our retraction (8) is only defined on a strict subset of $T\tilde{\mathcal{C}}_n^{\min}$, we rely on a stability certificate (Lemma 5.5) for our analysis. With this construction, we then proceed to choose our domain sufficiently small so that $F(.)$ is well-defined. To show that $F(.)$ is $\mathcal{L}_0$-invariant, we present convexity-like (Lemma 5.4) and Lipschitz-like (Lemma 5.6) inequalities useful for analysis of first-order methods adopted for smooth manifolds.

Lemma 5.3: Let $\text{grad}J_n(K^*) = 0$ and (s_-, s_0, s_+) be the signature of Hess $\tilde{J}_n(K^*)$. Then $\nabla J_n([K^*]) = 0$ and the signature of $\nabla^2 J_n([K^*])$ is $(s_-, s_0 - n^2, s_+)$.

Proof: At stationary points, the signature of the Hessian is invariant of the Riemannian metric [15, Prop. 8.71.]. So, $\nabla^2 \tilde{J}_n(K^*)$ and Hess $\tilde{J}_n(K^*)$ share the same signature. By [15, Ex. 9.46.], the eigenvalues of $\nabla^2 \tilde{J}_n(K^*)$ are exactly the eigenvalues of $\nabla^2 J_n([K^*])$ with n^2 additional zeros. ∎

Next, we demonstrate that $\mathcal{R}$-balls[3] about $[K^*]$ satisfy a strong convexity-like inequality:

Lemma 5.4: Let $(\mathcal{M}, \langle .,. \rangle, \mathcal{R})$ be a Riemannian manifold. Let $f : \mathcal{M} \to \mathbb{R}$ be smooth with $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*) > 0$. Then there exists $\rho > 0$ for which $\nabla^2 f > 0$ on $\mathcal{D} := \overline{B_{x^*}(\rho)}$, and $M > 0$ such that $\mathcal{D}$ contains the unique connected component $\mathcal{L}_0$ of the sublevel set $\mathcal{L}(f, M) := \{y \in \mathcal{M} : f(y) \leq M\}$ containing x^*.

Proof: Choose small enough $\rho > 0$ so that $\mathcal{D} \subset \text{dom}(\mathcal{R}_{x^*})$ and $\nabla^2 f > 0$ on $\mathcal{D}$. By [15, Prop. 5.44],

$$f(\mathcal{R}_{x^*}(v)) = f(x^*) + \frac{1}{2}\langle \nabla^2 f(x^*)v, v\rangle_{x^*} + O(\|v\|_{x^*}^3).$$

Since $\frac{1}{2}\langle \nabla^2 f(x^*)v, v\rangle_{x^*} = O(\|v\|_{x^*}^2)$ is positive-definite, then $f(\mathcal{R}_{x^*}(v)) > f(x^*)$ for all $\|v\|_{x^*} \leq \rho$ for a small enough $\rho > 0$. It follows $f(x^*) < f(y)$ for all $y \in \mathcal{D} - \{x^*\}$.

Pick $\epsilon > 0$ small enough so that $M := \min f(\partial\mathcal{D}) - \epsilon > f(x^*)$. Pick $y \in \mathcal{L}_0$. Let $c : [0,1] \to \mathcal{L}_0$ be any curve from x^* to y in $\mathcal{L}_0$. If $c(t) \in \mathcal{D}$ for only $0 \leq t \leq t_{\max} < 1$, then $f(c(t_{\max})) \geq M + \epsilon > M$, a contradiction. So, $y \in \mathcal{D}$. ∎

Let $\rho > 0$ be small enough so that $\tilde{\mathcal{D}} := \overline{B_{K^*}(\rho)} \subset \tilde{\mathcal{C}}_n^{\min}$. Since $\mathcal{D} := \overline{B_{[K^*]}(\rho)} \subsetneq \pi(\tilde{\mathcal{D}})$, then ensure $\rho > 0$ is additionally small enough so $\mathcal{D}$ satisfies the constraints in Lemma 5.4. This lemma grants us $M > 0$ and $\mathcal{L}_0 \subset \mathcal{D} \subset \mathcal{C}_n$.

Next, we must construct a stability certificate [9] for $\tilde{\mathcal{C}}_n^{\min}$:

Lemma 5.5: Define the stability certificate[4]

$$s(K, \mathbf{V}) := \frac{1}{2\|A_{\text{cl}}(\mathbf{V})\|_2 \overline{\lambda}(\mathbb{L}(A_{\text{cl}}(K), I_{2n}))} > 0.$$

Then $\mathcal{R}_K(t\mathbf{V}) \in \tilde{\mathcal{C}}_n$ for $t \in [0, s(K, \mathbf{V}))$.

Proof: Set $P := \mathbb{L}(A_{\text{cl}}(K), I_{2n})$. Then

$$t\overline{\lambda}(A_{\text{cl}}(\mathbf{V})P + PA_{\text{cl}}^\intercal(\mathbf{V})) \leq 2t\|A_{\text{cl}}(\mathbf{V})\|_2 \overline{\lambda}(P) < 1,$$

and so $t(A_{\text{cl}}(\mathbf{V})P + PA_{\text{cl}}(\mathbf{V})^\intercal) \prec I_{2n}$. Since $A_{\text{cl}}(.)$ is linear and $A_{\text{cl}}(K)P + PA_{\text{cl}}^\intercal = -I_{2n}$, we have $A_{\text{cl}}(K^+)P + PA_{\text{cl}}(K^+)^\intercal \prec 0$, where $K^+ = K + t\mathbf{V}$. ∎

Now, we will guarantee a Lipschitz-like inequality. For $K \in \tilde{\mathcal{C}}_n^{\min}$, define $r(K) := \frac{1}{2}\|\nabla \tilde{J}_n(K)\|_K^{-1} \min_{\|\mathbf{V}\|_K = 1} s(K, \mathbf{V})$.

[3]A $\mathcal{R}$-ball is $B_x(\rho) := \{\mathcal{R}_x(\xi) : \|\xi\|_x < \rho\}$.

[4]Take note this does not ensure $K + t\mathbf{V}$ is minimal.

Lemma 5.6: Let $\mathcal{K} \subset \tilde{\mathcal{C}}_n^{\min}$ be compact. Define $\mathcal{T} := \{(K, \mathbf{V}) \in T\tilde{\mathcal{C}}_n^{\min} : K \in \mathcal{K}, \|\mathbf{V}\|_K \leq r(K)\}$. Next, define $\mathcal{T}^* := \{(\mathcal{T}_S(K), d\mathcal{T}_S(\mathbf{V})) : (K, \mathbf{V}) \in \mathcal{T}, S \in \text{GL}(n)\}$. Then there exists $L > 0$ where for all $(K, \mathbf{V}) \in \mathcal{T}^*$,

$$\tilde{J}_n(\mathcal{R}_K(\mathbf{V})) \leq \tilde{J}_n(K) + \langle \nabla \tilde{J}_n(K), \mathbf{V}\rangle_K + \frac{L}{2}\|\mathbf{V}\|_K^2. \quad (19)$$

Proof: Remark that $r(.)$ is continuous and $\mathcal{T} \subset T\tilde{\mathcal{C}}_n^{\min} \subset T\tilde{\mathcal{C}}_n$ is compact. Since $\tilde{J}_n$ is also analytic over $\tilde{\mathcal{C}}_n$, then $\tilde{J}_n \circ \mathcal{R} : \mathcal{T} \to \mathbb{R}$ is well-defined and analytic over compact $\mathcal{T}$. It follows its derivatives are bounded uniformly, and hence satisfies (19) [15, Lemma 10.57].

Fix $(K, \mathbf{V}) \in \mathcal{T}$ and $S \in \text{GL}(n)$. Since $K + \mathbf{V} \in \tilde{\mathcal{C}}_n$, so is $\mathcal{T}_S(K) + d\mathcal{T}_S(\mathbf{V})$. Due to the invariance properties, (19) holds for $(\mathcal{T}_S(K), d\mathcal{T}_S(\mathbf{V}))$. ∎

Let $L > 0$ be the Lipschitz constant from Lemma 5.6 with $\mathcal{K} := \tilde{\mathcal{D}}$. Ensure L sufficiently large so that $\frac{1}{L} \leq \min r(\tilde{\mathcal{D}})$. The Lipschitz-like inequality holds for all $(K, \mathbf{V}) \in \mathcal{T}^*$. Take note that $\tilde{\mathcal{D}}^* := \pi_1(\mathcal{T}^*) = \bigcup_{S \in \text{GL}(n)} \mathcal{T}_S(\tilde{\mathcal{D}}) = \pi^{-1}(\pi(\tilde{\mathcal{D}}))$.

Set $\tilde{\mathcal{L}}_0 := \pi^{-1}(\mathcal{L}_0)$. Then $\tilde{\mathcal{L}}_0 \subset \tilde{\mathcal{D}}^*$

Lemma 5.7: We have $\tilde{F}(K) := \mathcal{R}_K(-\frac{1}{L}\nabla \tilde{J}_n(K)) \in \tilde{\mathcal{L}}_0$ for $K \in \tilde{\mathcal{L}}_0$.

Proof: Let $K^+ := \tilde{F}(K)$. Define the curve $c(t) = \mathcal{R}_K(-\frac{t}{L}\nabla \tilde{J}_n(K))$. Plugging this into (19), it follows $\tilde{J}_n(K^+) \leq \tilde{J}_n(K)$. Also, $\pi \circ c$ is continuous, contained in $\mathcal{L}_0$, and connects $[K]$ to $[K^+]$. Since $\mathcal{L}_0$ is closed, we must have $[K^+] \in \mathcal{L}_0$. Therefore $K^+ \in \tilde{\mathcal{L}}_0$. ∎

This all implies that $F : \mathcal{L}_0 \to \mathcal{L}_0$, $F([K]) := [\tilde{F}(K)]$ is well-defined and smooth. Our local convergence guarantee now follows from [15, Thm. 4.19].

VI. Numerical Experiments and Results

We will now compare RGD with ordinary gradient descent (Figure 3). Our step size procedure is Algorithm 2.

Algorithm 2 Backtracking Line-Search

Require: $K \in \tilde{\mathcal{C}}_n^{\min}$, $\gamma \in (0,1)$, $\beta \in (0,1)$, $\bar{s} > 0$
  $s \leftarrow \bar{s}$
  $K^+ \leftarrow K - s\nabla J_n(K))$
  while $K^+ \notin \tilde{\mathcal{C}}_n^{\min}$ **or** $J_n(K) - J_n(K^+) < \gamma s\|\nabla J_n(K)\|_K^2$ **do**
    $s \leftarrow \beta s$
    $K^+ \leftarrow K - s\nabla J_n(K))$
  return s

We ran our numerical experiments against four representative systems. 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 $J_n(K) - J_n^* < 10^{-10}$. 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)** $w_1 = w_2 = w_3 = 1$ and **(2)** with $w_1 = 1$, $w_2 = w_3 = 0$.

The first system is Doyle's famous counterexample. The second system is a plant whose LQG controller is non-minimal. These examples can be found in [5]. The third system admits saddle points with vanishing Hessian. The fourth system has dimensions $(n, m, p) = (4, 3, 3)$ and entries either set to zero or sampled from the standard Gaussian distribution with probability 0.8.

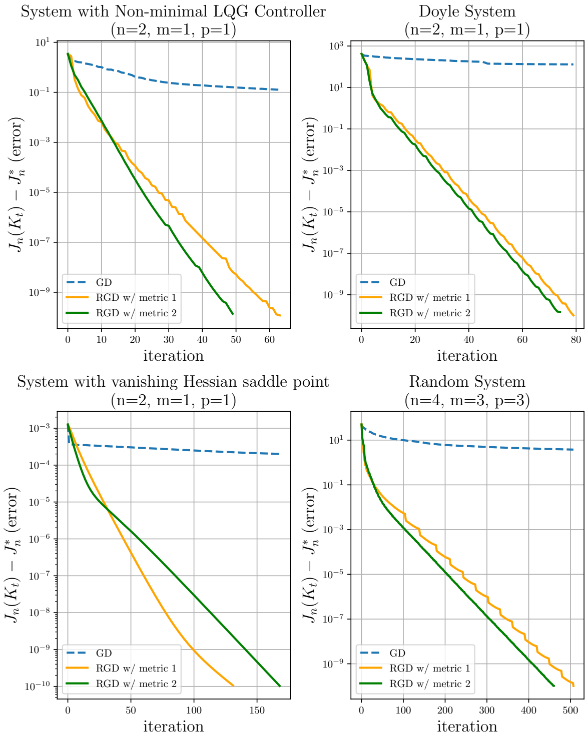

Fig. 3.

As we observe, in all four cases, Algorithm 1 significantly outperforms GD. In fact, for the vanishing Hessian system, GD gets stuck in the non-strict saddle point. Furthermore, note that the first and second metrics used for RGD result in algorithms with comparable performances, with one performing better over the other for some examples. Lastly, we point out the eventual linear rate of convergence of the corresponding RGDs for all these examples.

VII. Future Directions

We intend to study second order PO methods and how they perform for output feedback synthesis problems, particularly their behavior around saddle points and achieving a super-linear convergence rate. We also plan to study LQG PO over reduced-order feedback controllers. In fact, we note that most of the results delineated in this paper do not rely on having full-order controllers. We also plan to expand on the discrete formulation of our setup. Lastly, finite-horizon LQG PO, Kalman filter PO, and connections between the KM metric and the Riemannian metric used in [9], will be examined in our subsequent works.

VIII. Acknowledgements

The authors thank Shahriar Talebi for helpful discussions on policy optimization and Riemannian geometry of stabilizing feedback gains. The research of the authors has been supported by NSF grant ECCS-2149470 and AFOSR grant FA9550-20-1-0053.

References

[1] B. Hu, K. Zhang, N. Li, M. Mesbahi, M. Fazel, and T. Başar, "Toward a theoretical foundation of policy optimization for learning control policies," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, pp. 123–158, 2023.
[2] B. Hu and Y. Zheng, "Connectivity of the feasible and sublevel sets of dynamic output feedback control with robustness constraints," *IEEE Control Systems Letters*, vol. 7, pp. 442–447, 2022.
[3] J. Bu, A. Mesbahi, and M. Mesbahi, "On topological and metrical properties of stabilizing feedback gains: the mimo case," *arXiv preprint arXiv:1904.02737*, 2019.
[4] J. Bu, A. Mesbahi, and M. Mesbahi, "On topological properties of the set of stabilizing feedback gains," *IEEE Transactions on Automatic Control*, vol. 66, no. 2, pp. 730–744, 2020.
[5] Y. Tang, Y. Zheng, and N. Li, "Analysis of the optimization landscape of Linear Quadratic Gaussian (LQG) control," in *Learning for Dynamics and Control*, pp. 599–610, PMLR, 2021.
[6] J. C. Doyle, "Guaranteed margins for LQG regulators," *IEEE Transactions on automatic Control*, vol. 23, no. 4, pp. 756–757, 1978.
[7] S.-I. Amari, "Natural gradient works efficiently in learning," *Neural computation*, vol. 10, no. 2, pp. 251–276, 1998.
[8] B. Mishra and R. Sepulchre, "Riemannian preconditioning," *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 635–660, 2016.
[9] S. Talebi and M. Mesbahi, "Policy optimization over submanifolds for constrained feedback synthesis," *arXiv preprint arXiv:2201.11157*, 2022.
[10] R. Kalman, "Algebraic geometric description of the class of linear systems of constant dimension," in *8th Annual Princeton Conference on Information Sciences and Systems*, vol. 3, Princeton University Princeton, NJ, 1974.
[11] M. Hazewinkel, "Moduli and canonical forms for linear dynamical systems ii: The topological case," *Mathematical Systems Theory*, vol. 10, no. 1, pp. 363–385, 1976.
[12] B. Afsari and R. Vidal, "Bundle reduction and the alignment distance on spaces of state-space lti systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3804–3819, 2017.
[13] M. 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*, pp. 125–193, Springer, 1980.
[14] A. Tannenbaum, *Invariance and System Theory: Algebraic and Geometric Aspects*, vol. 845. Springer, 2006.
[15] N. Boumal, *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023.
[16] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
[17] J. M. Lee, *Smooth Manifolds*. Springer, 2012.
[18] J. M. Lee, *Introduction to Riemannian Manifolds*, vol. 2. Springer, 2018.
[19] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry, Volume 2*, vol. 61. John Wiley & Sons, 1996.
[20] N. Boumal, P.-A. Absil, and C. Cartis, "Global rates of convergence for nonconvex optimization on manifolds," *IMA Journal of Numerical Analysis*, vol. 39, no. 1, pp. 1–33, 2019.
[21] P. S. Krishnaprasad and C. F. Martin, "On families of systems and deformations," *International Journal of Control*, vol. 38, no. 5, pp. 1055–1079, 1983.
[22] P. S. Krishnaprasad, *Geometry of Minimal Systems and the Identification Problem*. Harvard University Press, 1977.