

Learning Task-Space Synergies using Riemannian Geometry

Martijn J.A. Zeestraten¹, Ioannis Havoutis^{2,3}, Sylvain Calinon^{2,1}, Darwin. G. Caldwell¹

Abstract—In the context of robotic control, synergies can form elementary units of behavior. By specifying task-dependent coordination behaviors at a low control level, one can achieve task-specific disturbance rejection. In this work we present an approach to learn the parameters of such low-level controllers by demonstration. We identify a synergy by extracting covariance information from demonstration data. The extracted synergy is used to derive a time-invariant state feedback controller through optimal control. To cope with the non-Euclidean nature of robot poses, we utilize Riemannian geometry, where both estimation of the covariance and the associated controller take into account the geometry of the pose manifold. We demonstrate the efficacy of the approach experimentally in a bimanual manipulation task.

I. INTRODUCTION

Synergies are functional groupings of elements that are constrained to work as a single unit [1]. Their existence potentially explains how the central nervous system is able to perform tasks accurately despite the redundant and noisy sensorimotor system [1]–[4]. They also appear at task level: grasping and manipulation require coordination between objects in the environment and the degrees of freedom of the hands.

Similarly, in the context of robotic control, synergies can be used as elementary units of behavior. By specifying task-dependent coordination behaviors at a low control level, one can achieve task-specific disturbance rejection. In this work we present an approach to learn the parameters of such synergistic controllers by demonstration.

Our approach extracts task-space synergies from demonstration data, and uses them to generate an optimal state feedback controller. It relies on the Linear Quadratic Regulator (LQR)—a control paradigm that simplifies the design of optimal controllers for linear dynamical systems. This optimal regulator is found by minimizing a cost function parameterized by a positive definite tracking cost matrix Q and a control cost matrix R . Most often, the required tracking precision is specified for each state variable individually as a diagonal tracking cost matrix Q . We instead propose to exploit the off-diagonal elements to specify the coupling among state variables. The ability to specify synergies through Q while guaranteeing stability makes the LQR well suited for our goal.

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union’s 7th Framework Programme FP7/2007-2013/ under REA grant agreement number 608022. It is also supported by the DexROV Project through the EC Horizon 2020 programme (Grant #635491) and by the Swiss National Science Foundation through the I-DRESS project (CHIST-ERA).

¹ Department of Advanced Robotics, Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genova, Italy. name.surname@iit.it

² Idiap Research Institute, Rue Marconi 19, 1920 Martigny, Switzerland. name.surname@idiap.ch

³ Oxford Robotics Institute, Department of Engineering Science, University of Oxford, United Kingdom. ioannis@robots.ox.ac.uk

Task-space synergies require a suitable parameterization of robot pose, which involves both position and orientation. Since no parameterization of orientation exists that is global singularity free and Euclidean, the typical methods available for imitation learning and LQR cannot directly be applied. We build upon the probabilistic framework for imitation learning on Riemannian manifolds introduced in our previous work [5]. This framework allows us to learn distributions over robot poses whose support is contained in a regular geodesic ball [6]. In practice, this restricts the orientation data to lie within a $\pm\pi$ radius of the empirical mean (the Riemannian center of mass). This is achieved by encoding robot poses as elements on the manifold $\mathbb{R}^3 \times \mathcal{S}^3$ —The Cartesian product of the 3-dimensional Euclidean space and the unit-quaternion manifold \mathcal{S}^3 , respectively.

Different approaches for LQR on $SO(3)$, $SE(3)$ or coverings of these groups exists. Saccon *et al.* [7] derive a LQR controller on $SO(3)$ through Pontryagin’s Maximum Principle. Marinho *et al.* in [8] use a dual-quaternion representation to derive a LQR tracking controller. The latter involves converting the dual-quaternion transformation invariant error into an affine time-varying system. Such representation can be compared to a pose manifold, yet the position quaternion does not represent well the Cartesian space, and the method requires the manual specification of the control and state error costs. Similarly, Wang and Yu [9] present a dual quaternion controller for rigid-body motion stabilization and tracking, built on a screw theory formulation.

Our approach to learn synergies from demonstration involves the estimation of stiffness and damping matrices from the correlation observed in the demonstration data. Similarly, Rozo *et al.* [10] and Saveriano and Lee [11] estimate the stiffness directly from the covariance information. Smoother stiffness profiles can be obtained from the covariance information through LQR as demonstrated by Medina *et al.* [12], Calinon *et al.* [13] and Zeestraten *et al.* [14]. Kronander and Billard [15] use a combination of tactile and kinesthetic teaching to communicate the desired stiffness of the robot along a trajectory. Unlike these previous works, the presented method considers coordination among position and orientation of multiple end-effectors.

Our contributions are two-fold: i) We demonstrate that the Riemannian Gaussian can be used to encode rich synergies of task-space manipulation that involve positions and orientations; ii) We show how infinite horizon LQR can be used to regulate synergies that are defined on Riemannian manifolds. The method used to perform statistics on manifolds is restricted to simply-connected manifolds [6]. We demonstrate the efficacy of our approach in a bimanual manipulation task, where three different coordination patterns between two end-effector poses are taught. This paper is accompanied by a video of the experimental evaluation, and source codes of

\mathcal{M} :	\mathbb{R}^d	\mathcal{S}^3
$g \in \mathcal{M}$	$\begin{pmatrix} g_{x_1} \\ \vdots \\ g_{x_d} \end{pmatrix}$	$\begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix}$
$\text{Exp}_e(\mathfrak{g})$	$e + \begin{bmatrix} \mathfrak{g}_{x_1} \\ \vdots \\ \mathfrak{g}_{x_d} \end{bmatrix}$	$\begin{cases} \begin{pmatrix} c(\ \mathfrak{g}\) \\ s(\ \mathfrak{g}\) \frac{\mathfrak{g}}{\ \mathfrak{g}\ } \end{pmatrix}, & \ \mathfrak{g}\ \neq 0 \\ \begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix}, & \ \mathfrak{g}\ = 0 \end{cases}$
$\text{Log}_e(\mathbf{g})$	$\begin{bmatrix} g_{x_1} \\ \vdots \\ g_{x_d} \end{bmatrix}$	$\begin{cases} \text{ac}_*(q_0) \frac{\mathbf{q}}{\ \mathbf{q}\ }, & q_0 \neq 1 \\ \begin{bmatrix} 0 \\ \mathbf{0} \\ 0 \end{bmatrix}, & q_0 = 1 \end{cases}$
$\mathcal{A}_g^h(\mathbf{p})$	$\mathbf{p}_g - \mathbf{g} + \mathbf{h}$	$\mathbf{g} * \mathbf{h} * \mathbf{g}^{-1} * \mathbf{p}_g$

TABLE I: Overview of the exponential and logarithmic maps at the origin e , and the action for the Riemannian manifolds considered in this work. $s(\cdot)$, $c(\cdot)$, $\text{ac}_*(\cdot)$ are short notations for the sine, cosine, and a modified version of the arccosine (see Sec. II), respectively. The elements of \mathcal{S}^3 are quaternions, $*$ defines their product, and \mathbf{g}^{-1} the quaternion inverse of \mathbf{g} .

the approach is available through

<http://www.idiap.ch/software/pbdlib/>.

The remainder of this paper is organized as follows: In Section II we introduce some preliminaries of Riemannian geometry and statistics. Section III introduces how we propose to learn and reproduce task-space synergies. The approach is evaluated experimentally in Section IV, and we conclude with a discussion in Section V.

II. PRELIMINARIES

Our objective is to use infinite horizon LQR for system states that are defined on Riemannian manifolds [16]. Unlike the Euclidean space, the Riemannian Manifold is not a vector space where sum and scalar multiplication are defined. Therefore, we cannot directly apply Euclidean methods to data defined on a Riemannian manifold \mathcal{M} . However, these methods can be applied in the tangent spaces of the manifold. These are Euclidean and provide a way to indirectly perform computation on the manifold. We indicate the manifold by \mathcal{M} , and its elements in bold typeface, i.e. $\mathbf{p} \in \mathcal{M}$. Tangent spaces are defined at each point $\mathbf{p} \in \mathcal{M}$ and indicated by $\mathcal{T}_{\mathbf{p}}\mathcal{M}$. Elements on the tangent space are displayed in fraktur typeface, i.e. $\mathfrak{g} \in \mathcal{T}_{\mathbf{p}}\mathcal{M}$. Finally, matrices are displayed as capital boldface (greek) letters.

For each point on a Riemannian manifold we can define an *exponential map* $\text{Exp}_g(\cdot) : \mathcal{T}_g\mathcal{M} \rightarrow \mathcal{M}$, a distance preserving map between the tangent space and the manifold. $\text{Exp}_g(\mathfrak{p})$ maps \mathfrak{p} to \mathbf{p} in such a way that \mathbf{p} lies on the *geodesic*—the generalization of straight lines—through \mathbf{g} with direction \mathfrak{p} , and the distance between \mathbf{g} and \mathbf{p} is $\|\mathfrak{p}\| = \langle \mathfrak{p}, \mathfrak{p} \rangle_g$, see Fig. 1a. The inverse mapping is called the *logarithmic map*, and is defined within the injectivity radius of the exponential map [6].

In general, one exponential and logarithmic map is required for each tangent space. For homogeneous manifolds, however, their function can be moved from the origin e to

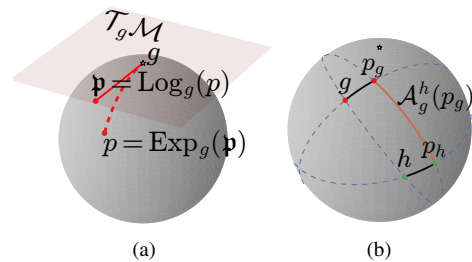


Fig. 1: Manifold mappings and action function illustrated on the 2-sphere. (a) The exponential and the logarithmic map provide local one-to-one mappings between the manifold and its tangent space at point \mathbf{g} . (b) An action $\mathcal{A}_g^h(\mathbf{p}_g)$ maps \mathbf{p}_g (a point defined relative to \mathbf{g}) to \mathbf{p}_h by moving it along a geodesic (dotted lines) until it reaches a point such that the distance between \mathbf{p}_h and \mathbf{h} equals the distance between \mathbf{p}_g and \mathbf{g} (both distances visualized by black lines).

other points on the manifold with

$$\text{Exp}_g(\mathbf{p}_g) = \mathcal{A}_e^g(\text{Exp}_e(\mathbf{p}_g)), \quad (1)$$

$$\text{Log}_g(\mathbf{p}) = \text{Log}_e(\mathcal{A}_g^e(\mathbf{p})), \quad (2)$$

where $\mathcal{A}_g^h(\mathbf{p}_g)$ is called the action function (see Fig. 1b).

The manifolds used in our experiments are the 3-dimensional Euclidean space for position, and the 3-sphere to represent orientation in terms of unit-quaternions. Table I gives an overview of the required mappings.

Orientation controllers used for unit-quaternions can suffer from the ‘unwinding’ problem [17]. The manifold of unit-quaternions, \mathcal{S}^3 , provides a double covering of the space of rotation matrices $SO(3)$, and both \mathbf{q} and $-\mathbf{q}$ represent the same orientation. If the control objective is to steer to \mathbf{q} , while the system is at $-\mathbf{q}$ (which also represents the desired orientation), the controller will leave the desired state to ‘unwind’ to \mathbf{q} . Such behavior can be avoided by explicitly taking into account the sign change, at the cost of making the control law discontinuous. To ensure that the distance between two antipodal rotations is zero we use

$$\text{arccos}_*(\rho) \begin{cases} \arccos(\rho) - \pi & , -1 \leq \rho < 0 \\ \arccos(\rho) & , 0 \leq \rho \leq 1 \end{cases} \quad (3)$$

in the definition of our logarithmic map for \mathcal{S}^3 .

The maximum entropy distribution on a Riemannian manifold, given a center and a covariance, is a distribution of the exponential family [6]. Although the exact solution is computationally impractical, it is often approximated by a Gaussian-like distribution

$$\mathcal{N}_{\mathcal{M}}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^d |\boldsymbol{\Sigma}|}} e^{-\frac{1}{2} \text{Log}_{\boldsymbol{\mu}}(\mathbf{x})^\top \boldsymbol{\Sigma}^{-1} \text{Log}_{\boldsymbol{\mu}}(\mathbf{x})}, \quad (4)$$

where $\boldsymbol{\mu} \in \mathcal{M}$ is the Riemannian center of mass, and $\boldsymbol{\Sigma}$ the covariance defined in the tangent space $\mathcal{T}_{\boldsymbol{\mu}}\mathcal{M}$ (see e.g. [18]–[20]). In [5], we showed how operations commonly used in probabilistic imitation learning can be generalized to Riemannian manifolds. Note that uniqueness of the mean is only guaranteed if support of the distribution lies within a regular geodesic ball of radius $r < \frac{\pi}{2\sqrt{\kappa}}$ (with κ the curvature of the manifold; $\kappa_{\mathbb{R}^n} = 0$, $\kappa_{\mathcal{S}^3} = 1$). Practically, one can verify




Fig. 2: Visualization of step responses obtained by Riemannian LQR on the system state manifold, $\mathcal{M}_s = \mathcal{S}^2 \times \mathbb{R}^2$, for two different covariance matrices (red and blue ellipsoids). The initial state of the system is indicated by \mathbf{p} and the desired state by $\boldsymbol{\mu}$. a) Response path in \mathcal{S}^2 and $\mathcal{T}_\mu \mathcal{S}^2$. The response on the manifold is visualized by the solid lines, and the response on the tangent space by the dotted line of corresponding color. b) Response in \mathbb{R}^2 which is diffeomorphic to $\mathcal{T}_0 \mathbb{R}^2$.

this by ensuring the distance between the computed mean and the data is contained in r , i.e. $r = \max_n \text{dist}(\boldsymbol{\mu}, \mathbf{x}_n) < \frac{\pi}{2\sqrt{\kappa}}$ [6]. For unit-quaternions the maximum allotted distance is $\pi/2$ which corresponds to a rotation of π . Although Euler angles face a similar limitations, they cannot be used in our application because they are not defined in a single vector space¹.

The Cartesian product of two Riemannian manifolds is again a Riemannian manifold. For example, a robot pose can be represented by the Cartesian product of a 3 dimensional Euclidean space and a hypersphere, i.e. $\mathbf{p} \in \mathbb{R}^3 \times \mathcal{S}^3$. The corresponding $\text{Exp}()$, $\text{Log}()$ of the Cartesian product are obtained by concatenating the individual functions, e.g.

$$\text{Log}_{\begin{bmatrix} e_a \\ e_b \end{bmatrix}} \left(\begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} \right) = \begin{bmatrix} \text{Log}_{e_a}(\mathbf{a}) \\ \text{Log}_{e_b}(\mathbf{b}) \end{bmatrix}.$$

III. LEARNING SYNERGETIC CONTROL BY DEMONSTRATION

We start with a training set consisting of N data points, $\mathbf{x} \in \mathcal{M}$. This set potentially contains synergetic coupling among the manifold dimensions. Our aim is to find a controller that preserves these synergies. To identify them, we estimate the center $\boldsymbol{\mu} \in \mathcal{M}$ and covariance $\boldsymbol{\Sigma} \in \mathcal{T}_\mu \mathcal{M}$ of a Riemannian Gaussian using the Maximum Likelihood Estimate (MLE) [5]. The covariance matrix encodes the local synergies around the estimated center. Similarly to our previous work [13], [14], we use LQR to replicate the encoded behavior. LQR is a controller for linear systems of the form $\dot{\boldsymbol{\xi}} = \mathbf{A}\boldsymbol{\xi} + \mathbf{B}\mathbf{u}$ that optimizes a cost function that is quadratic in both state $\boldsymbol{\xi}$ and control input \mathbf{u} ,

$$c = \frac{1}{2} \int (\boldsymbol{\xi}^\top \mathbf{Q} \boldsymbol{\xi} + \mathbf{u}^\top \mathbf{R} \mathbf{u}) dt. \quad (5)$$

The solution to this optimal control problem is a state-feedback controller of the form $\mathbf{u} = \mathbf{L}\boldsymbol{\xi}$. Its gain matrix, \mathbf{L} ,

¹Euler angles are 3-tuples (i.e., not 3-dimensional coordinate vectors) and therefore cannot be used to define a covariance matrix—doing so would yield an incorrect encoding of the observed distribution.

is obtained by solving an algebraic Riccati equation (see e.g. [21]).

The required linear system cannot be defined on the manifold, since it is not a vector space. However, we can exploit the linear tangent spaces to achieve a similar result. The state error between the desired state \mathbf{p}_d and current state \mathbf{p} can be computed using the logarithmic map $\boldsymbol{\epsilon} = \text{Log}_{\mathbf{p}_d}(\mathbf{p})$ that projects the minimum length path between \mathbf{p}_d and \mathbf{p} into the Euclidean space $\mathcal{T}_{\mathbf{p}_d} \mathcal{M}$. We define the linear time-invariant system,

$$\underbrace{\begin{bmatrix} \dot{\boldsymbol{\epsilon}} \\ \dot{\boldsymbol{\epsilon}} \end{bmatrix}}_{\mathbf{s}} = \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{0} & -\mathbf{M}^{-1}\mathbf{C} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \boldsymbol{\epsilon} \\ \dot{\boldsymbol{\epsilon}} \end{bmatrix}}_{\mathbf{s}} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1} \end{bmatrix}}_{\mathbf{B}} \mathbf{u}. \quad (6)$$

with inertia matrix \mathbf{M} and damping matrix \mathbf{C} . The augmentation with $\dot{\boldsymbol{\epsilon}}$ makes the system state manifold to be $\mathcal{M}_s = \mathcal{M} \times \mathbb{R}^d$. As a result the linear system is defined in $\mathcal{T}_{\bar{\mathbf{p}}_d} \mathcal{M}_s$ with $\bar{\mathbf{p}}_d \in \mathcal{M}_s$ (the original desired state augmented with a desired velocity), and the mapping from the manifold to this tangent space

$$\mathbf{s} = \text{Log}_{\bar{\mathbf{p}}_d}(\bar{\mathbf{p}}). \quad (7)$$

The covariance $\boldsymbol{\Sigma}$ of a Riemannian Gaussian $\mathcal{N}_{\mathcal{M}}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ describes the variance and correlation of the state variables in a tangent space defined at $\boldsymbol{\mu} \in \mathcal{M}$. By assuming that the desired LQR tracking performance can be related to the variance, we can define it using the covariance. We formulate the cost function (5) in the tangent space $\mathcal{T}_{\bar{\boldsymbol{\mu}}} \mathcal{M}_s$ using (7)

$$c = \frac{1}{2} \int \left(\underbrace{\text{Log}_{\bar{\boldsymbol{\mu}}}(\bar{\mathbf{p}})^\top}_{\mathbf{s}^\top} \mathbf{Q} \underbrace{\text{Log}_{\bar{\boldsymbol{\mu}}}(\bar{\mathbf{p}})}_{\mathbf{s}} + \mathbf{u}^\top \mathbf{R} \mathbf{u} \right) dt, \quad (8)$$

with $\bar{\boldsymbol{\mu}} = \begin{bmatrix} \boldsymbol{\mu} \\ \mathbf{0} \end{bmatrix}$, $\mathbf{Q} = \begin{bmatrix} \boldsymbol{\Sigma}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$,

and \mathbf{R} the control cost matrix. With the dynamical system (6) and cost function (8), the optimal state feedback controller

$$\mathbf{u} = \mathbf{L}\mathbf{s} \stackrel{(7)}{=} \mathbf{L} \text{Log}_{\bar{\boldsymbol{\mu}}}(\bar{\mathbf{p}}), \quad (9)$$

can be computed. Similarly to classical infinite horizon LQR, the gain matrix, $\mathbf{L} = \mathbf{R}^{-1} \mathbf{B}^\top \mathbf{X}$, is obtained by solving the algebraic Riccati equation

$$\mathbf{A}\mathbf{X} + \mathbf{A}\mathbf{X} - \mathbf{X}\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\top\mathbf{X} + \mathbf{Q} = \mathbf{0}. \quad (10)$$

The control law (9) is only valid when $\text{Log}(\bar{\mathbf{p}})$ is defined. For $SO(3)$ this boundary lies on π . The original logarithmic map for the unit-quaternion manifold \mathcal{S}^3 is not defined at the antipode of its origin (which corresponds to 2π rotation). Since unit quaternions provide a double covering of $SO(3)$, this point can be avoided using the modified arccosine (3).

Fig. 2 demonstrates the approach on the manifold $\mathcal{S}^2 \times \mathbb{R}^2$. It shows how the response of the system changes based on the shape of the covariance matrix. LQR is computed in the tangent space of the attractor $\boldsymbol{\mu}$. Its response is visualized in the tangent space, and projected on the manifold.

Note that LQR controllers defined on (a covering of) $SO(3)$ can only achieve local asymptotic stability, since no




Fig. 3: Visualization of the three different behaviors: i) horizontal planar translation (in blue), ii) vertical planar rotation (in red), iii) coupled rotation and translation (in green).

continuous time-invariant feedback controller can achieve global asymptotic stability on compact manifolds such as $SO(3)$ [17]. Although such controllers can achieve almost global asymptotic stability, they also possess unstable equilibria [22]. Similarly, no globally stable continuous time-invariant feedback controller exists on $\mathbb{R}^3 \times \mathcal{S}^3$, or the space of dual quaternions [23]. However, the attraction domain of the unstable equilibria is nowhere dense (see e.g. [22]). Therefore, we consider the absence of global convergence a theoretical limitation that has no practical consequence to the learning and reproduction of synergies from demonstration.

IV. EVALUATION

A. Synergies in Bi-manual Manipulation

The presented method is tested in a bimanual task. Our aim is to demonstrate the variety of synergies that can be learned and reproduced using our approach.

The experimental setup consists of two Barrett WAMs with three fingered hands. Together, the two end-effectors hold a ball. We evaluate three different synergies: i) horizontal planar translation; ii) vertical planar rotation; iii) translation coupled with rotation. The setup with an illustration of the coordination patterns is shown in Fig. 3.

The synergies are taught through kinesthetic teaching [10]. For each synergy we demonstrated the tolerated motion of the hands around a desired ball pose. For example, behavior iii) is demonstrated by repeatedly moving both hands unidirectionally away from the center while rotating along the axis of motion. The demonstration data consist of hand-pose pairs which are defined at the hand palms.

The demonstration data of the bi-manual skill lie on the 12-dimensional ($d = 12$) manifold

$$\mathcal{M}_{bm} = \underbrace{\mathbb{R}^3 \times \mathcal{S}^3}_{\text{Pose 1}} \times \underbrace{\mathbb{R}^3 \times \mathcal{S}^3}_{\text{Pose 2}}. \quad (11)$$

For each behavior we computed the MLE of the mean and covariance of the Riemannian Gaussian $\mathcal{N}_{\mathcal{M}_{bm}}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and added $1 \cdot 10^{-3} \mathbf{I}_d$ as prior to the covariance. This regularization term prevents high gains in the state variables that have very low variance; it bounds the gains found through LQR.

The resulting models are visualized in Fig. 4. The covariance of the position clearly shows the preferred direction of motion: i) motion in the horizontal plane; ii) motion in the vertical plane; iii) motion along one axis. Similarly, the rotational covariance provides information about the allotted rotation: i) no rotation in any direction; ii) rotation in a single plane; iii) rotation around a single axis.

The learned synergies appear in the correlation matrices². For i), the planar coupling between the hand positions results in the strong positive correlation between $x_{L,1}$ and $x_{R,1}$, and $x_{L,2}$ and $x_{R,2}$. For ii), the synergy involves a rotation around the global x_1 axis. This is correctly captured in the correlation between the $\omega_{L,1}$, and $\omega_{R,2}$. Furthermore, there exists a strong negative correlation between the x_3 axes of the left and right hand. This indicates the opposite upwards/downwards motion made during the rotation. Note that the rotation of the hands around the ball created a circular motion around its center. This requires a nonlinear coupling between the x_2 and x_3 of both hands, something that cannot be encoded in a single Gaussian. For iii), strong correlation is observed between $x_{2,L}$ and $x_{2,R}$, indicating the motion along the axis of translation. The strong correlation between $x_{2,L}$, $x_{2,R}$ and $\omega_{2,L}$, $\omega_{2,R}$ establishes the coupling between the translation and rotation.

To reproduce the demonstrated synergies, we employ the Riemannian LQR presented in Section III. The system state manifold for the bimanual skill

$$\mathcal{M}_s = \underbrace{\mathbb{R}^3 \times \mathcal{S}^3}_{\text{Pose 1}} \times \underbrace{\mathbb{R}^3 \times \mathcal{S}^3}_{\text{Pose 2}} \times \underbrace{\mathbb{R}^3 \times \mathbb{R}^3}_{\text{Pose 1 velocity}} \times \underbrace{\mathbb{R}^3 \times \mathbb{R}^3}_{\text{Pose 2 velocity}}, \quad (12)$$

consists of the skill-manifold augmented with the pose velocities. We define a linear system in the tangent space $\mathcal{T}_{\boldsymbol{\mu}} \mathcal{M}_s$ (6), where \mathbf{M} and \mathbf{C} are the end-effector inertia and damping, which we approximate by

$$\mathbf{M} = \text{diag}(1.17, 1.17, 1.17, 0.009, 0.008, 0.005,$$

$$1.17, 1.17, 1.17, 0.009, 0.008, 0.005),$$

$$\mathbf{C} = \text{diag}(20, 20, 20, 1, 1, 0.1, 20, 20, 20, 1, 1, 0.1),$$

respectively. We run the controller with a frequency of 500 Hz.

The control cost matrix \mathbf{Q} is constructed from the inverse covariance matrix of the Gaussian describing the desired synergy

$$\mathbf{Q} = \begin{bmatrix} \boldsymbol{\Sigma}^{-1} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \end{bmatrix}, \quad (13)$$

hereby setting a zero cost on the desired velocity. Furthermore, we manually defined a constant control cost matrix

$$\mathbf{R} = \text{diag}\left(\frac{1}{80}, \frac{1}{80}, \frac{1}{80}, 12.5, 12.5, 25, \frac{1}{80}, \frac{1}{80}, \frac{1}{80}, 12.5, 12.5, 25\right),$$


which was the same for all three synergies.

By solving the LQR problem we obtain the gain matrix $\mathbf{L} \in \mathbb{R}^{12 \times 24}$, which we use to compute the control command $\mathbf{u} = [\mathbf{u}_L^T, \mathbf{u}_R^T]^T \in \mathbb{R}^{12}$. The desired joint torques are computed using


$$\boldsymbol{\tau} = \boldsymbol{\tau}_g(\mathbf{q}) + \mathbf{J}(\mathbf{q})^T \mathbf{u}, \quad (14)$$

with $\boldsymbol{\tau}_g$ the torques required for gravity compensation, \mathbf{J} the manipulator Jacobian, and \mathbf{q} the joint angles.


²The covariance matrix combines correlation $-1 \leq \rho_{ij} \leq 1$ among random variables X_i and X_j with deviation σ_i of random variables X_i , i.e. it has elements $\Sigma_{ij} = \rho_{ij} \sigma_i \sigma_j$. We prefer to visualize the correlation matrix instead of the covariance matrix, since it only contains the correlation coefficients and highlights the coordination among variables.



(a) Figures depicting the center and covariance of the position and orientation. The colored ellipsoids show the covariance of the left and right end-effectors in red and blue, respectively, with their boundary representing one standard deviation. The orthogonal red, green and blue lines indicate the end-effectors mean orientation. The colored ellipsoids at the end of each line depict the rotational covariance using $\sqrt{0.3}$ -standard deviation.



(b) Correlation matrices. The entries related to the position and orientation of the left and right end-effectors are labeled $x_{R/L}$ and $\omega_{R/L}$, respectively.



(c) Reproduction examples. The orthogonal red, blue and green lines originate from the end-effector position and visualize the end-effector orientation. The lines represent the end-effector x_1 , x_2 and x_3 axes respectively.

Fig. 4: Visualization of coordination encoding described in Sec. IV-A. The three behaviors (i)–(iii) are ordered left-to-right.

Typical reproductions of the encoded synergies are visualized in Fig. 4c and in the video accompanying the paper. Fig. 5 shows the step response of the real system and the linear system (6) for behavior iii). The figure shows the stable response of the real system. Furthermore, the transient of the real system and the linear system are similar. The steady-state errors of the real system are due to the approximated inertia matrix.

V. DISCUSSION AND CONCLUSION

We presented an approach to learning task-space synergy controllers from demonstration. We achieved this by exploiting a Riemannian geometry method to combine different

manifolds through the Cartesian product, which makes the approach easily adaptable to a variety of different manifolds. For example, all experimental evaluations of LQR presented in this work are performed using one single piece of code. Changing from the toy-example to the bimanual manipulation example solely required specifying a different system state manifold.

The presented approach relies on the ability to measure minimum distances on the manifold. This is not naturally achieved on $SE(3)$ because there exists neither a bi-invariant metric nor a natural left or right invariant metric [24]. In this work we choose to use a left-invariant metric because




Fig. 5: Step responses of the real and simulated linear system, visualized in the tangent space of the target state. The elements ϵ_1 – ϵ_3 , and ϵ_4 – ϵ_6 correspond to the position and orientation, respectively.

it allows us to encode the synergy models independent of the inertial frame. In addition, we need to specify a relative weighting of the positional and rotational components. Since a similar trade-off is faced when balancing torques and forces through \mathbf{R} , we choose to equally weight position and orientation contributions in the metric. As a result, we can directly use the metric resulting from our pose manifold $\mathbb{R}^3 \times \mathcal{S}^3$. The proper weighting is postponed until the inevitable tuning of the control cost matrix.

The control cost matrix remains an open parameter that requires manual tuning. In our experimental evaluation, we set the values of \mathbf{R} in such a way that the control effort is well balanced for the translation and rotational degrees of freedom. The fact that we could use a single \mathbf{R} for the three different synergies shows that its selection is more system than task-dependent. In practice, one could specify a fixed system-dependent ratio \mathbf{R}_f for the control variables and allow the user to control the overall control cost of the system by a single parameter β , i.e. $\mathbf{R} = \beta \mathbf{R}_f$. Here, \mathbf{R} should be bounded to prevent unstable controllers due to actuator limitations of the real systems.

Throughout the paper, we focused on the derivation of a controller for a single synergy. Real manipulation tasks will require a variety of synergies along a desired trajectory. By encoding manipulation tasks in a Gaussian Mixture Model (GMM), representing the joint distribution of time and pose $\mathcal{P}(t, \mathbf{p})$, one can for example obtain a time dependent synergy by computing $\mathcal{P}(\mathbf{p}|t)$ using Gaussian Mixture Regression (GMR) [5]. In situations where regression is computationally demanding, one could run the regression at a relative low rate, while the presented controller ensures a synergetic response to disturbance.

REFERENCES

- [1] J. A. S. Kelso, *Synergies: Atoms of Brain and Behavior*. Boston, MA: Springer US, 2009, pp. 83–91.
- [2] E. Todorov and M. I. Jordan, “Optimal feedback control as a theory of motor coordination,” *Nature Neuroscience*, vol. 5, no. 11, pp. 1226–1235, 2002.
- [3] M. L. Latash, J. P. Scholz, and G. Schöner, “Toward a new theory of motor synergies,” *Motor control*, vol. 11, no. 3, pp. 276–308, 2007.
- [4] D. Sternad, M. E. Huber, and N. Kuznetsov, “Acquisition of novel and complex motor skills: stable solutions where intrinsic noise matters less,” in *Progress in Motor Control*. Springer, 2014, pp. 101–124.
- [5] M. J. A. Zeestraten, I. Havoutis, J. Silvério, S. Calinon, and D. G. Caldwell, “An approach for imitation learning on Riemannian manifolds,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 2, no. 3, pp. 1240–1247, June 2017.
- [6] X. Pennec, “Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements,” *Journal of Mathematical Imaging and Vision*, vol. 25, no. 1, pp. 127–154, 2006.
- [7] A. Saccon, J. Hauser, and A. P. Aguiar, “Optimal control on lie groups: The projection operator approach,” *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2230–2245, Sept 2013.
- [8] M. M. Marinho, L. F. C. Figueredo, and B. V. Adorno, “A dual quaternion linear-quadratic optimal controller for trajectory tracking,” in *Proc. IEEE/RSSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, 2015, pp. 4047–4052.
- [9] X. Wang and C. Yu, “Unit dual quaternion-based feedback linearization tracking problem for attitude and position dynamics,” *Systems & Control Letters*, vol. 62, no. 3, pp. 225–233, 2013.
- [10] L. Rozo, S. Calinon, D. G. Caldwell, P. Jimenez, and C. Torras, “Learning collaborative impedance-based robot behaviors,” in *AAAI Conf. on Artificial Intelligence*, Bellevue, WA, USA, 2013, pp. 1422–1428.
- [11] M. Saveriano and D. Lee, “Learning motion and impedance behaviors from human demonstrations,” in *Intl. Conf. on Ubiquitous Robots and Ambient Intelligence (URAI)*, 2014, pp. 368–373.
- [12] J. R. Medina, D. Lee, and S. Hirche, “Risk-sensitive optimal feedback control for haptic assistance,” in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2012, pp. 1025–1031.
- [13] S. Calinon, D. Bruno, and D. G. Caldwell, “A task-parameterized probabilistic model with minimal intervention control,” in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, Hong Kong, China, 2014, pp. 3339–3344.
- [14] M. J. A. Zeestraten, S. Calinon, and D. G. Caldwell, “Variable duration movement encoding with minimal intervention control,” in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016, pp. 497–503.
- [15] K. Kronander and A. Billard, “Learning compliant manipulation through kinesthetic and tactile human-robot interaction,” *IEEE Transactions on Haptics*, vol. 7, no. 3, pp. 367–380, 2014.
- [16] J. M. Lee, *Riemannian Manifolds - An Introduction to Curvature*, 1st ed. New York: Springer-Verlag, 1997.
- [17] S. P. Bhat and D. S. Bernstein, “A topological obstruction to continuous global stabilization of rotational motion and the unwinding phenomenon,” *Systems & Control Letters*, vol. 39, no. 1, pp. 63–70, 2000.
- [18] G. Dubbelman, “Intrinsic statistical techniques for robust pose estimation,” Ph.D. dissertation, University of Amsterdam, September 2011.
- [19] E. Simo-Serra, C. Torras, and F. Moreno-Noguer, “3D human pose tracking priors using geodesic mixture models,” *International Journal of Computer Vision*, vol. 122, no. 2, pp. 388–408, 2017.
- [20] S. Kim, R. Haschke, and H. Ritter, “Gaussian mixture model for 3-dof orientations,” *Robotics and Autonomous Systems*, vol. 87, pp. 28–37, October 2017.
- [21] A. E. Bryson, *Dynamic optimization*. Addison Wesley Longman, 1999.
- [22] N. A. Chaturvedi, A. K. Sanyal, and N. H. McClamroch, “Rigid-body attitude control,” *IEEE Control Systems Magazine*, vol. 31, no. 3, pp. 30–51, June 2011.
- [23] H. T. M. Kussaba, L. F. C. Figueredo, J. Y. Ishihara, and B. V. Adorno, “Hybrid kinematic control for rigid body pose stabilization using dual quaternions,” *Journal of the Franklin Institute*, vol. 354, no. 7, pp. 2769–2787, 2017.
- [24] F. C. Park, “Distance metrics on the rigid-body motions with applications to mechanism design,” *J. of Mech. Des.*, vol. 117, pp. 48–54, 1995.