

Nonparametric Continuous Sensor Registration

William Clark

*Department of Mathematics
Cornell University
Ithaca, NY 14850, USA*

WAC76@CORNELL.EDU

Maani Ghaffari

*Department of Naval Architecture and Marine Engineering
University of Michigan
Ann Arbor, MI 48109 USA*

MAANIGJ@UMICH.EDU

Anthony Bloch

*Department of Mathematics
University of Michigan
Ann Arbor, MI 48109 USA*

ABLOCH@UMICH.EDU

Editor: Sayan Mukherjee

Abstract

This paper develops a new mathematical framework that enables nonparametric joint semantic and geometric representation of continuous functions using data. The joint embedding is modeled by representing the processes in a reproducing kernel Hilbert space. The functions can be defined on arbitrary smooth manifolds where the action of a Lie group aligns them. The continuous functions allow the registration to be independent of a specific signal resolution. The framework is fully analytical with a closed-form derivation of the Riemannian gradient and Hessian. We study a more specialized but widely used case where the Lie group acts on functions isometrically. We solve the problem by maximizing the inner product between two functions defined over data, while the continuous action of the rigid body motion Lie group is captured through the integration of the flow in the corresponding Lie algebra. Low-dimensional cases are derived with numerical examples to show the generality of the proposed framework. The high-dimensional derivation for the special Euclidean group acting on the Euclidean space showcases the point cloud registration and bird's-eye view map registration abilities. An implementation of this framework for RGB-D cameras outperforms the state-of-the-art robust visual odometry and performs well in texture and structure-scarce environments.

Keywords: Kernel methods, nonparametric representation, sensor registration, Lie groups, Riemannian geometry, equivariant models

1. Introduction

We consider the problem of sensor registration, which is defined as finding the transformation between two sensors with overlapping measurements or tracking a moving sensor using its sequential measurements. This problem frequently arises in engineering domains such as point cloud registration (Chen and Medioni, 1991; Censi, 2008; Segal et al., 2009; Stoyanov et al., 2012; Servos and Waslander, 2017; Parkison et al., 2018; Zaganidis et al., 2018;




Figure 1: Conceptual illustration of nonparametric continuous sensor registration. Continuous functions are constructed from measurements shown in red circles. The domain of the functions can be an arbitrary smooth manifold, M , where the action of a Lie group, \mathcal{G} , is used to align them via $h.f_Z$ and $h \in \mathcal{G}$.

Parkison et al., 2019), visual odometry and SLAM (Nistér et al., 2006; Scaramuzza and Fraundorfer, 2011; Fraundorfer and Scaramuzza, 2012; Strasdat, 2012; Kerl et al., 2013a; Engel et al., 2015; Ghaffari et al., 2019; Rosinol et al., 2019), satellite image registration (Kim and Im, 2003; Bentoutou et al., 2005), and photogrammetry (Mikhail et al., 2001). In terms of applications, the above mentioned problems are fundamental to many computer vision (Hartley and Zisserman, 2003; Szeliski, 2010) and robotics (Thrun et al., 2005; Barfoot, 2017) algorithms.

In this paper, we formulate a fundamentally novel nonparametric framework for sensor registration that constructs continuous functions from raw measurements; hence our approach can be considered as a direct method. The functions can be thought of as curves, surfaces, or hyper-surfaces that live in a Reproducing Kernel Hilbert Space (RKHS). For example, given an RGB-D image, a function can be constructed that maps a three-dimensional (3D) point to the RGB color space. Therefore, such functions inherently represent a joint model for semantic/appearance and geometry. Furthermore, the domain of the functions can be an arbitrary smooth manifold where the action of a Lie group is used to align them. This concept is illustrated in Figure 1. The continuous functions allow the registration to be independent of a specific signal resolution and the framework is fully analytical with a closed-form derivation of the Riemannian gradient and Hessian.

In the present work, we study a restriction of the framework where the Lie group acts on functions isometrically. This restriction has a variety of applications and its high-dimensional derivation for the special Euclidean group acting on the Euclidean space showcases the point cloud registration and bird’s-eye view map registration abilities. In particular, this work has the following contributions:

1. a new mathematical framework for sensor registration that enables nonparametric joint semantic/appearance and geometric representation of continuous functions using data;
2. low-dimensional cases are derived with numerical examples to show the generality of the proposed framework;
3. a high-dimensional derivation for the special Euclidean group acting on the Euclidean space showcases the point cloud registration and bird’s-eye view map registration abilities. A specific derivation and implementation of this framework for RGB-D cameras performs well in texture and structure-scarce environments.
4. the open-source implementation of the derived cases in this work is available for download at
<https://github.com/MaaniGhaffari/c-sensor-registration>
<https://github.com/MaaniGhaffari/cvo-rgbd>

An early idea of this work in the context of RGB-D visual odometry was presented at the 2019 Robotics: Science and Systems conference in Freiburg, Germany (Ghaffari et al., 2019). This work generalizes the previous framework and places it into a broader context. In particular, an abstract sensor registration problem is defined in Section 2, where low-dimensional and high-dimensional examples together with their second-order geometry on flat or curved spaces are studied. The generalization to curved spaces also connects this work to an ongoing research direction regarding positive-definite kernel design and hyperparameter learning on curved spaces (Jayasumana et al., 2015; Feragen et al., 2015; Borovitskiy et al., 2020). Furthermore, the derivation of the Riemannian Hessian has important implications in quantifying the pose estimation covariance (Censi, 2007; Bonnabel et al., 2016; Landry et al., 2019; Brossard et al., 2020) and symmetry of the problem by studying vector fields generated by the eigenvalues of Hessian (Zhu et al., 2021), and therefore is included for completeness. However, its applications in those research areas and second-order solvers are left for future studies. Finally, the new experiments on RGB-D visual odometry in Section 9 are done via an open-source and efficient C++ implementation and compared against DVO (Kerl et al., 2013a,b).

The remainder of this paper is organized as follows. A review of the required mathematical machinery and notation is given in Appendix A. Section 2 presents the novel continuous sensor registration framework. Section 3 provides low-dimensional examples of the framework including point cloud registration on circles and torus. We discuss some theoretical limitations of the kernel design on compact Riemannian manifolds in Section 4. The high-dimensional example for the special Euclidean group is given in Section 5. A brief theoretical analysis for the verification of the idea is provided in Section 6. The integration of the flow for the special Euclidean group to obtain the solution is explained in Section 7. Section 8 deals with the Riemannian geometry of the special Euclidean groups. Experimental evaluations of the proposed method for registration and tracking using RGB-D images are presented in Section 9. Section 10 provides discussions regarding several important topics related to the proposed approach. Finally, Section 11 concludes the paper and provides future work ideas.

2. Problem Setup

Let M be a smooth manifold and consider two (finite) collections of points, $X = \{x_i\}$, $Z = \{z_j\} \subset M$. Also, suppose we have a (Lie) group, \mathcal{G} , acting on M . We want to determine which element $h \in \mathcal{G}$ aligns “best” the two point clouds X and $hZ = \{hz_j\}$. To assist with this, we will assume that each point contains information described by a point in an inner product space, $(\mathcal{I}, \langle \cdot, \cdot \rangle_{\mathcal{I}})$. To this end, we will introduce two labeling functions, $\ell_X : X \rightarrow \mathcal{I}$ and $\ell_Z : Z \rightarrow \mathcal{I}$.

In order to measure their alignment, we will be turning the clouds, X and Z , into functions $f_X, f_Z : M \rightarrow \mathcal{I}$ that live in some reproducing kernel Hilbert space, $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$.

Remark 1 *The action, $\mathcal{G} \curvearrowright M$ induces an action $\mathcal{G} \curvearrowright C^\infty(M)$ by*

$$h.f(x) := f(h^{-1}x).$$

Inspired by this observation, we will set $h.f_Z := f_{h^{-1}Z}$.

In the following, we give a summary of the recipe for cloud alignment as well as formal sensor registration problems.

Definition 2 *A sensor registration problem is given by a 5-tuple: $(\mathcal{G}, M, \varphi, \langle \cdot, \cdot \rangle_{\mathfrak{g}}, k)$ where*

(G1) \mathcal{G} is a Lie group,

(G2) M is a smooth manifold,

(G3) $\varphi : \mathcal{G} \rightarrow \text{Diff}(M)$ is a smooth group action,

(G4) $\langle \cdot, \cdot \rangle_{\mathfrak{g}}$ is an inner product on $\mathfrak{g} = \text{Lie}(\mathcal{G})$, and

(G5) $k : M \times M \rightarrow \mathbb{R}$ is a symmetric positive definite function, called the kernel,

while the information is given by a 3-tuple: $(\mathcal{I}, (X, \ell_X), (Z, \ell_Z))$ where

(I1) \mathcal{I} is an inner product space, called the information space,

(I2) $X \subset M$ is a finite collection of points called the target and $\ell_X : X \rightarrow \mathcal{I}$ is its label, and

(I3) $Z \subset M$ is a finite collection of points called the source and $\ell_Z : Z \rightarrow \mathcal{I}$ is its label.

The information (G1)-(G5) is required to build the general form of the gradient while the information (I1)-(I3) encodes the actual clouds which is subsequently plugged into the gradient. The remainder of this work consists of understanding cases with varying (G1)-(G5) information. We first define the general form of the problem as follows.

Problem 1 *Suppose f_X and f_Z are the constructed functions over point clouds X and Z , respectively. The problem of aligning two point clouds can be formulated as minimizing the distance between f_X and $h.f_Z$ in the sense of the RKHS norm. That is, we want to solve*

$$\arg \min_{h \in \mathcal{G}} J(h), \quad J(h) := \|f_X - h.f_Z\|_{\mathcal{H}}^2, \quad (1)$$

where \mathcal{H} is the RKHS induced by the kernel k from (G5).

In this work, we restrict the problem to a special case. If $\mathcal{G} \curvearrowright C^\infty(M)$ is an isometry, we have $\|f_X - h.f_Z\|_{\mathcal{H}}^2 = \|f_X\|_{\mathcal{H}}^2 + \|f_Z\|_{\mathcal{H}}^2 - 2\langle f_X, h.f_Z \rangle_{\mathcal{H}}$. Then we can reformulate Problem 1 into a reduced form as follows.

Problem 2 *The problem of aligning the point clouds can now be rephrased as maximizing the scalar products of f_X and $h.f_Z$, i.e., we want to solve*

$$\arg \max_{h \in \mathcal{G}} F(h), \quad F(h) := \langle f_X, h.f_Z \rangle_{\mathcal{H}}. \quad (2)$$

2.1 Constructing the Functions

We first choose a symmetric function $k : M \times M \rightarrow \mathbb{R}$ to be the kernel of our RKHS, \mathcal{H} . This allows us to turn the point clouds to functions via

$$\begin{aligned} f_X(\cdot) &:= \sum_{x_i \in X} \ell_X(x_i) k(\cdot, x_i), \\ f_Z(\cdot) &:= \sum_{z_j \in Z} \ell_Z(z_j) k(\cdot, z_j). \end{aligned} \quad (3)$$

We can now define the inner product of f_X and f_Z by

$$\langle f_X, f_Z \rangle_{\mathcal{H}} := \sum_{x_i \in X, z_j \in Z} \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}} \cdot k(x_i, z_j). \quad (4)$$

We note that it is possible to use the well-known kernel trick in machine learning (Bishop, 2006; Williams and Rasmussen, 2006; Murphy, 2012) to substitute the inner products in (4) with the appearance/semantic kernel, i.e., $k_c : \mathcal{I} \times \mathcal{I} \rightarrow \mathbb{R}$. The kernel trick can be applied to carry out computations implicitly in the high dimensional space, which leads to computational savings when the dimensionality of the feature space is large compared to the number of data points (Williams and Rasmussen, 2006). After applying the kernel trick to (4), we get

$$\langle f_X, f_Z \rangle_{\mathcal{H}} = \sum_{x_i \in X, z_j \in Z} k_c(\ell_X(x_i), \ell_Z(z_j)) \cdot k(x_i, z_j) := \sum_{x_i \in X, z_j \in Z} c_{ij} \cdot k(x_i, z_j). \quad (5)$$

Remark 3 *We note two advantages of measuring the alignment of X and Z by (4). The first is that we do not need identification of which point of X should be paired with what point of Z . The second is that the number of points in X does not even need to be equal to the number of points in Z !*

Remark 4 *Although the double sum in (4) seems to make the problem computationally intractable when dealing with a large number of data points, the inner product in (4) is sparse. The sparse structure is due to the kernelized formulation and the fact that the kernel bandwidth is spatially limited. As such, most kernel terms are zero. See Fig. 1 in Ghaffari et al. (2019). This notion of sparsity extends to labels, i.e., semantic information, once we kernelize the information inner product as done in (5). Furthermore, to exploit this sparse structure is the motivation for the development of a custom solver. However, the specific solver used in this work is not tied to the proposed framework.*

2.2 Building the Gradient Flow

In order to (at least locally) solve (2), we will construct a gradient flow: $\dot{h} = \nabla F(h)$ (this is similar to the treatment in Bloch et al. (1992) where a gradient flow is used to maximize the inner product between elements of a Lie algebra corresponding to a compact Lie group). Before we proceed, we will first determine the differential, dF . In order to do this, we will need the notion of an infinitesimal generator for a group action (see chapter 4 of Berndt and Klucznik 2001).

Definition 5 *Suppose that a Lie Group \mathcal{G} acts diffeomorphically on a smooth manifold M via φ ; that is*

$$\begin{aligned}\varphi : \mathcal{G} &\rightarrow \text{Diff}(M) \\ g &\mapsto \varphi_g.\end{aligned}$$

For a given $\xi \in \mathfrak{g} = \text{Lie}(\mathcal{G})$, we denote the vector field ξ_M (called the infinitesimal generator) on M given by the rule:

$$df_x(\xi_M) := \left. \frac{d}{dt} \right|_{t=0} f(\varphi_{\exp(t\xi)}(x)), \quad f \in C^\infty(M). \quad (6)$$

This lets us compute the differential, dF .

Theorem 6 *Suppose that $F(h) = \langle f_X, h.f_Z \rangle_{\mathcal{H}}$ as described above. Then*

$$dF_e(\xi) = \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j}(-\xi_M(z_j)), \quad (7)$$

where $\tilde{k}_{x_i} = k(x_i, \cdot)$.

Remark 7 *The notation for the differential of a function used throughout this paper is $df_x(v)$, where $x \in M$ and $v \in T_x M$:*

$$df_x(v) = \left. \frac{d}{dt} \right|_{t=0} f(c(t)), \quad c(0) = x, \quad c'(0) = v.$$

Proof This follows from a straightforward application of the chain rule.

$$\begin{aligned}dF_e(\xi) &= \left. \frac{d}{dt} \right|_{t=0} \langle f_X, \exp(t\xi).f_Z \rangle_{\mathcal{H}} \\ &= \sum c_{ij} \cdot \left. \frac{d}{dt} \right|_{t=0} k(x_i, \exp(-t\xi)z_j) \\ &= \sum c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j} \cdot \left. \frac{d}{dt} \right|_{t=0} \exp(-t\xi)z_j \\ &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j}(-\xi_M(z_j)).\end{aligned}$$

This matches equation (7). ■

Of course, to construct the gradient flow we are interested in computing dF_h instead of just dF_e . We will accomplish this via left-translation. Left-translation is given by the smooth map $\ell_h : \mathcal{G} \rightarrow \mathcal{G}$ where $x \mapsto hx$. Its differential gives rise to an isomorphism of tangent spaces, $(\ell_h)_* : \mathfrak{g} \xrightarrow{\sim} T_h\mathcal{G}$.

Corollary 8 *Under the identification $T_h\mathcal{G} \cong (\ell_h)_*\mathfrak{g}$, we have*

$$dF_h((\ell_h)_*\xi) = \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{h^{-1}z_j}(-\xi_M(h^{-1}z_j)), \quad (8)$$

In order to turn the co-vector $dF_h \in T_h^*\mathcal{G}$ into a vector $\nabla F_h \in T_h\mathcal{G}$, we will use a left-invariant metric. This can be accomplished by defining an inner-product, $\langle \cdot, \cdot \rangle_{\mathfrak{g}}$ on \mathfrak{g} and lifting to a (Riemannian) metric on \mathcal{G} via left-translation, i.e.

$$\langle (\ell_h)_*\eta, (\ell_h)_*\xi \rangle_{T_h\mathcal{G}} := \langle \eta, \xi \rangle_{\mathfrak{g}}.$$

This allows us to define the gradient of F as

$$\langle \nabla F_h, (\ell_h)_*\xi \rangle_{T_h\mathcal{G}} = dF_h((\ell_h)_*\xi). \quad (9)$$

This allows for a way to obtain a (local) solution to (2) by following

$$\dot{h} = \nabla F(h). \quad (10)$$

To demonstrate the generality of this registration algorithm, the cases in Table 1 will be examined in the remainder of this paper.

M	S^1	T^2	S^2	\mathbb{R}^n
\mathcal{G}	S^1	T^2	$\text{SO}(3)$	$\text{SE}(n)$

Table 1: A list of the manifolds / Lie groups to be studied in this paper. Particular interest will be placed on the $(\mathbb{R}^2, \text{SE}(2))$ and $(\mathbb{R}^3, \text{SE}(3))$ examples.

3. Low-dimensional Examples

While the primary focus of this paper is to align RGB-D images in \mathbb{R}^3 , we take the opportunity to show how the above algorithm works on other spaces: specifically the spaces S^1 , T^2 , and S^2 : the circle, the 2-torus, and the sphere.

3.1 The Circle

To demonstrate our algorithm, we will first work out the case where the manifold containing the clouds is the circle, $M = S^1$. S^1 acts naturally on itself via left-translations so we shall take $G = S^1$ as well and φ to be left-translation, i.e.

$$\varphi(\alpha)(\theta) = \alpha + \theta \pmod{2\pi}.$$

The Lie algebra for the circle is the real line and we will therefore take the inner product to be scalar multiplication,

$$\langle a, b \rangle_{\mathbb{R}} = ab.$$

The last piece of geometry we need is the kernel (the information space and the clouds will be determined later). To build k , we will use the Euclidean distance of points on the circle: let $x, y \in [0, 2\pi) \sim S^1$.

$$\begin{aligned} d_2(x, y)^2 &:= (\cos x - \cos y)^2 + (\sin x - \sin y)^2 \\ &= 2(1 - \cos(x - y)). \end{aligned} \quad (11)$$

We will use the Gaussian kernel based on this distance function:

$$k(x, y) = \sigma^2 \exp\left(\frac{-d_2(x, y)^2}{2\ell^2}\right), \quad (12)$$

where $\ell > 0$ is a parameter called the length-scale.

If we again call $c_{ij} = \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}}$, then using (8), the differential is (for $s \in \mathbb{R}$)

$$dF_h((\ell_h)_*s) = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, z_j - h) \sin(z_j - h - x_i) \cdot s.$$

Therefore, the gradient is

$$\nabla F(h) = \frac{1}{\ell^2} \sum_{x_i \in X, z_j \in Z} c_{ij} k(x_i, z_j - h) \sin(z_j - h - x_i). \quad (13)$$

Finally, the Hessian can be found by simply differentiating the gradient as is

$$H_{S^1} = \frac{1}{\ell^2} \sum_{x_i \in X, z_j \in Z} c_{ij} k(x_i, z_j - h) [\sin^2(z_j - h - x_i) - \cos(z_j - h - x_i)]. \quad (14)$$

Example 1 *We perform the registration problem on the circle. $X = \mathbf{linspace}(0, \pi/2, 10)$ and $Z = \mathbf{linspace}(\pi/2, \pi, 12)$. Figure 2 shows two point clouds before and after registration. The labels, ℓ_X and ℓ_Z , for all points are set to 1.*

3.2 The Torus

The torus is the product of two circles: $T^2 = S^1 \times S^1$. Therefore, both the gradient and Hessian will be two copies of the corresponding object from the circle case.

Similarly to the circle case, the group will act via left-translations and we will take $(x^1, x^2) \in [0, 2\pi) \times [0, 2\pi) \sim T^2$ for coordinates. The kernel will be

$$\begin{aligned} k(x, y) &= \sigma^2 \exp\left(\frac{-d_{T^2}(x, y)^2}{2\ell^2}\right), \\ d_{T^2}(x, y)^2 &= d_2(x^1, y^1)^2 + d_2(x^2, y^2)^2. \end{aligned} \quad (15)$$




Figure 2: Left: The clouds X and $A \cdot Z$. Right: The clouds X and $T \cdot A \cdot Z$.

The gradient is then

$$\nabla F(h) = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, z_j - h) \begin{bmatrix} \sin(z_j^1 - h^1 - x_i^1) \\ \sin(z_j^2 - h^2 - x_i^2) \end{bmatrix}. \quad (16)$$

The Hessian is straight-forward to calculate due to the fact that T^2 is abelian. The Hessian is

$$H_{T^2} = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, z_j - h) \cdot h_{ij}, \quad (17)$$

where

$$h_{ij} = \text{diag} \left(\begin{aligned} &\sin^2(z_j^1 - h^1 - x_i^1) - \cos(z_j^1 - h^1 - x_i^1), \\ &\sin^2(z_j^2 - h^2 - x_i^2) - \cos(z_j^2 - h^2 - x_i^2) \end{aligned} \right)$$

This procedure naturally generalizes to T^n .

Example 2 We perform the registration problem on the torus with a 5 pointed star as shown in Figure 3. X contains 50 points while Z contains 75. The displacement is $[1, 1]$ and the transformation found was $[1.0077, 0.9962]$ which is an error of $[-0.0077, 0.0038]$ which has norm of 0.0085.

3.3 The Sphere

The other initial example we wish to develop is the case when $M = S^2$, the sphere. The Lie group we choose to act will be SO_3 , the special orthogonal group. This will act via usual matrix multiplication: let $x \in S^2 \subset \mathbb{R}^3$, then

$$\varphi(g)(x) = gx. \quad (18)$$




Figure 3: Top: The clouds X and $A \cdot Z$. Bottom: The clouds X and $T \cdot A \cdot Z$.

In terms of the point cloud data, assume that we have clouds with points in \mathbb{R}^3 . These points will be projected to S^2 (via normalizing) which is compatible with the group action described above. Given two points $x, y \in \mathbb{R}^3$, their arc-length distance, once projected to S^2 , is

$$d_{S^2}(x, y) = \arccos \left(\frac{\langle x, y \rangle_3}{\|x\| \|y\|} \right). \quad (19)$$

With this distance (which is different from the distance used in the circle example), we define the kernel to be

$$k(x, y) = \exp \left(\frac{-d_{S^2}(x, y)^2}{2\ell^2} \right). \quad (20)$$

In order to determine the gradient, we still need a metric for \mathfrak{so}_3 for which we will use the (negative of the) Killing form. This gives a gradient of $\nabla F(R) = R\hat{\omega}$ where $R \in \text{SO}_3$, $\hat{\cdot} : \mathbb{R}^3 \rightarrow \text{skew}_3$ (three-by-three skew-symmetric matrices), and

$$\omega = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \frac{k(x_i, R^{-1}z_j) d_{S^2}(x_i, R^{-1}z_j)}{\sqrt{\|x_i\|^2 \|z_j\|^2 - \langle x_i, R^{-1}z_j \rangle^2}} (x_i \times R^{-1}z_j). \quad (21)$$

3.3.1 THE HESSIAN

Unlike the circle and torus examples, calculating the Hessian for the sphere is substantially more involved. This follows from the fact that $\text{SO}(3)$ is no longer abelian. In order to determine the Hessian, we need to know the connection on $\text{SO}(3)$ induced by the Killing form on $\mathfrak{so}(3)$. Let us choose the standard basis $\{e_x, e_y, e_z\} \in \mathfrak{so}(3)$ where

$$e_x = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, \quad e_y = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \quad e_z = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Additionally, let $\{E_x, E_y, E_z\}$ be their corresponding left-invariant vector fields. The connection, calculated from (58), is given by

$$\begin{aligned}\nabla_{E_x} E_y &= \frac{1}{2} E_z, & \nabla_{E_x} E_z &= -\frac{1}{2} E_y, \\ \nabla_{E_y} E_z &= \frac{1}{2} E_x, & \nabla_{E_y} E_x &= -\frac{1}{2} E_z, \\ \nabla_{E_z} E_x &= \frac{1}{2} E_y, & \nabla_{E_z} E_y &= -\frac{1}{2} E_x.\end{aligned}$$

Writing (21) as $\hat{\omega} = \omega_x e_x + \omega_y e_y + \omega_z e_z$, the Hessian is given by

$$H(e_\alpha, e_\beta) = \sum_{\gamma \in \{x, y, z\}} \langle (\mathcal{L}_{E_\alpha} \omega_\gamma) E_\gamma + \omega_\gamma \nabla_{E_\alpha} E_\gamma, e_\beta \rangle = \mathcal{L}_{E_\alpha} \omega_\beta + \frac{1}{2} \varepsilon_{\alpha\gamma\beta} \omega_\gamma,$$

where $\varepsilon_{\alpha\gamma\beta} = \langle [e_\alpha, e_\gamma], e_\beta \rangle$ are the structure constants. Combining, the Hessian is

$$-H_{\text{SO}(3)} = \begin{bmatrix} \mathcal{L}_{E_x} \omega_x & \mathcal{L}_{E_x} \omega_y - \frac{1}{2} \omega_z & \mathcal{L}_{E_x} \omega_z + \frac{1}{2} \omega_y \\ \mathcal{L}_{E_y} \omega_x + \frac{1}{2} \omega_z & \mathcal{L}_{E_y} \omega_y & \mathcal{L}_{E_y} \omega_z - \frac{1}{2} \omega_x \\ \mathcal{L}_{E_z} \omega_x - \frac{1}{2} \omega_y & \mathcal{L}_{E_z} \omega_y + \frac{1}{2} \omega_x & \mathcal{L}_{E_z} \omega_z \end{bmatrix}, \quad (22)$$

which, although it is tedious, can be computed exactly.

Remark 9 *The reason that (22) returns the negative of the Hessian is because we are actually differentiating the inverse of the group element.*

Example 3 *To illustrate registration in this case, we will align two pictures of the globe. This will be done with a topographic map of the Earth using Matlab's `topo`. Both X and Z*

Parameters	Value
transformation convergence threshold ϵ	1e-4
gradient norm convergence threshold ϵ	5e-4
kernel characteristic length-scale ℓ	0.25
kernel characteristic length-scale ℓ (iteration > 3)	0.15
kernel characteristic length-scale ℓ (iteration > 10)	0.10
kernel characteristic length-scale ℓ (iteration > 20)	0.05
kernel signal variance σ	0.1
step length	0.1
color space inner product scale	1e-5
kernel sparsification threshold	1e-3

Table 2: Parameters used for evaluation of the globe registration. The kernel characteristic length-scale is chosen to be adaptive as the algorithm converges; intuitively, we prefer a large neighborhood of correlation for each point, but as the algorithm reaches the convergence reducing the local correlation neighborhood allows for faster convergence and better refinement.




Figure 4: Left (Reference): The image of the contour map of the Earth. Center (Perturbed): The plot of X along with $A \cdot Z$. Right (Aligned): The plot of X along with $T \cdot A \cdot Z$.

are initialized to be identical, however Z is perturbed by a transformation $A \in \text{SO}(3)$. The goal is to recover A^{-1} to realign Z with X . The algorithm reports $T \in \text{SO}(3)$ which should be the inverse of A . The values of A , T , $T \cdot A$, and $\|T \cdot A\|_F$ are shown in (23).

$$\begin{aligned}
 A &= \begin{bmatrix} 0.9386 & -0.3170 & 0.1361 \\ 0.3230 & 0.9461 & -0.0244 \\ -0.1210 & 0.0668 & 0.9904 \end{bmatrix}, & T &= \begin{bmatrix} 0.9391 & 0.3217 & -0.1205 \\ -0.3160 & 0.9466 & 0.0643 \\ 0.1347 & -0.0223 & 0.9906 \end{bmatrix}, \\
 T \cdot A &= \begin{bmatrix} 1.0000 & -0.0014 & 0.0007 \\ 0.0014 & 1.0000 & -0.0024 \\ -0.0007 & 0.0024 & 1.0000 \end{bmatrix}, & \|T \cdot A\|_F &= 0.0040.
 \end{aligned} \tag{23}$$

The parameters used in this registration are included in Table 2. The initial and final images are shown in Figure 4.

4. Limitations on Compact Riemannian Manifolds

In this section, we discuss an important problem regarding the positive-definiteness of the kernel on compact Riemannian manifolds (Jayasumana et al., 2015; Feragen et al., 2015; Borovitskiy et al., 2020). The key result is the following theorem.

Theorem 10 (Theorem 5 of Feragen et al. 2015, Bekka et al. 2008) *Let X be a topological space and let $\Psi : X \times X \rightarrow \mathbb{R}$ be a continuous kernel such that $\Psi(x, x) = 0$ and $\Psi(x, y) = \Psi(y, x)$ for all $x, y \in X$. Then the following are equivalent*

1. Ψ is conditionally of negative type, and
2. the kernel $\exp(t\Psi)$ is of positive type for every $t \geq 0$.

However, the issue is that for $\Psi = -d(x, y)^2$ to be conditionally negative, we need to be on Euclidean space. Fortunately, this is more than what we require in this work as we are not concerned if the positive-definiteness holds for all $t \geq 0$. Even though we are not on a Euclidean space, we do not know that there does not exist a $t \geq 0$ to make it positive definite. The following theorem ensures the existence of such a t .

Theorem 11 *Let $M \subset \mathbb{R}^m$ be some embedded manifold. Let d_c be the chordal distance on M (straight lines in \mathbb{R}^m) and d_g be the geodesic distance. For a finite collection of points $\{x_i\} \subset M$, there exists a $t \geq 0$ such that the matrix*

$$K_{ij} = [\exp(-t \cdot d_g(x_i, x_j)^2)]$$

is positive definite.

Proof It is clear that for all i , $K_{ii} = 1$. As long as the diagonal elements dominate, the matrix will be positive definite. Let

$$R_i = \sum_{j \neq i} \exp(-t \cdot d_g(x_i, x_j)).$$

Choose $t \geq 0$ large enough such that $R_i < 1$ for all i . Then by the Gershgorin circle theorem, all the eigenvalues of K are positive. A bound on t can be explicitly found via the following. Let

$$\Lambda := \min_{i \neq j} d_g(x_i, x_j).$$

Then

$$R_i \leq (N - 1) \cdot \exp(-t\Lambda),$$

where N is the number of points. This produces

$$t \geq \frac{\ln(N - 1)}{\Lambda}.$$

■

This is not a particularly useful result as the parameter t depends on the point cloud. Therefore it remains an open problem whether we can find a t (or an interval) that works on all possible point clouds systematically. This is a particularly attractive problem in machine learning as one ultimately wants to learn all hyperparameters. In particular, the derivation of useful bounds on t as well as algorithmic implementation of learning it are interesting future research directions.

Example 4 (The Sphere) *Let us see how the case of the sphere holds up. Here, we take the kernel*

$$k(x, y) = \exp\left(\frac{-d_{S^2}(x, y)^q}{2\ell^2}\right), \quad d_{S^2}(x, y) = \arccos\left(\frac{\langle x, y \rangle_3}{\|x\| \|y\|}\right). \quad (24)$$

In particular, $t = 1/(2\ell^2)$. For the usual case the exponent is $q = 2$. However, this may be changed as well. The idea is that k will not be positive definite for all ℓ , but it should be when ℓ is small enough. This will be numerically explored by taking a random sample of N points on S^2 and computing the matrix K . If all of its eigenvalues are positive, then it is positive definite. As such, we are interested in finding the smallest eigenvalue.

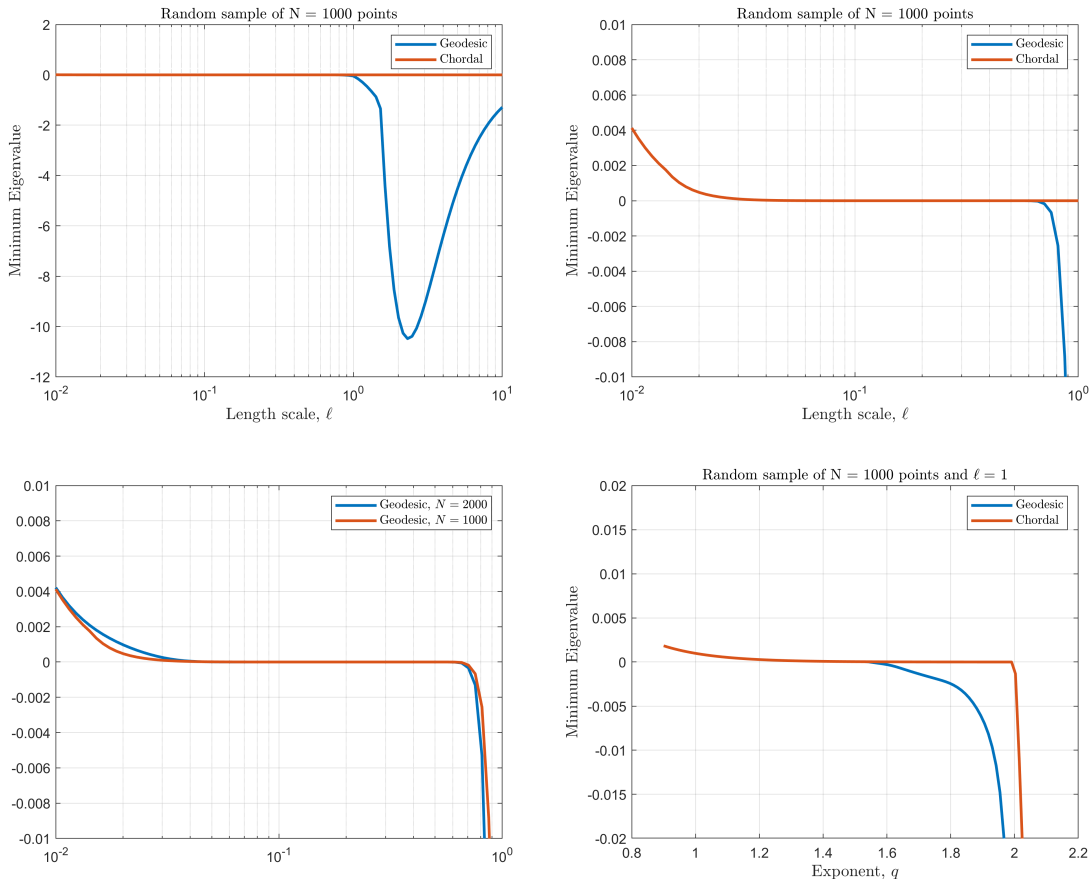


Figure 5: Top left: The geodesic kernel (24) fails to be positive definite as ℓ becomes large. As long as ℓ is small (which corresponds to t being large), the kernel appears to be positive definite. Top right: A zoomed version. Bottom left: The positive definiteness appears to not depend too heavily on the number of points. Bottom right: All of the previous figures were with the exponent fixed at $q = 2$. This figure varies the exponent. The chordal kernel appears to be positive-definite when $q \leq 2$ which is in agreement with the work of Feragen et al. (2015).

5. Special Euclidean Groups

We now move onto the more involved examples where $M = \mathbb{R}^n$ and $\mathcal{G} = \text{SE}(n)$, the special Euclidean group in n dimensions. \mathcal{G} acts on M in the standard fashion: let $(R, T) \in \text{SE}(n)$ where $R \in \text{SO}(n)$ and $T \in \mathbb{R}^n$,

$$(R, T).x = Rx + T, \quad x \in \mathbb{R}^n.$$

We will also choose the squared exponential kernel for $k : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$:

$$k(x, y) = \sigma^2 \exp\left(\frac{-\|x - y\|_n^2}{2\ell^2}\right), \tag{25}$$

for some fixed real parameters σ and ℓ , and $\|\cdot\|_n$ is the standard Euclidean norm on \mathbb{R}^n . In order to determine the gradient flow (10), we need to compute the infinitesimal generators of the action $\text{SE}(n) \curvearrowright \mathbb{R}^n$ as well as decide on a left-invariant metric for $\mathfrak{se}(n) = \text{Lie}(\text{SE}(n))$.

5.1 Infinitesimal Generator

For a fixed $\xi \in \mathfrak{se}(n)$, it has the form $\xi = (\omega, v)$ where $\omega \in \mathfrak{so}(n)$ and $v \in \mathbb{R}^n$. Because the infinitesimal generator map $\mathfrak{g} \rightarrow \mathfrak{X}(M)$ is a Lie algebra homomorphism (where $\mathfrak{X}(M)$ is the space of all vector fields over M , see §27 of Tu 2017), we see that $\xi_M = \omega_M + v_M$. A straight forward computation leads to

$$\xi_{\mathbb{R}^n} x = \hat{\omega} x + v, \quad \hat{\omega} \in \text{skew}(n), \quad v \in \mathbb{R}^n. \quad (26)$$

5.2 Metric

We need to choose a metric on $\mathfrak{se}(n)$ to turn dF_h into ∇F_h . For this example, we will take a multiple of the Killing form on $\mathfrak{so}(n)$ and the Euclidean norm on \mathbb{R}^n . That is,

$$\langle (\omega, v), (\eta, u) \rangle_{\mathfrak{se}(n)} = b^2 \cdot \langle v, u \rangle_n - a^2 \left(\frac{n-2}{2} \right) \cdot \text{Tr}(\omega\eta), \quad (27)$$

where $\langle \cdot, \cdot \rangle_n$ is the standard Euclidean inner product on \mathbb{R}^n (see Park and Brockett (1994) for a discussion in three dimensions), and a and b are tuning parameters. The reason for the $(2-n)/2$ term is because with this normalization (with $a = 1$) the skew matrices $E_{ij} - E_{ji}$ are orthonormal. Here E_{ij} denotes the matrix with all zeros except for a 1 in the (i, j) -coordinate.

5.3 Calculating the Gradient

Before we find the gradient, let us first determine its differential (at the identity for simplicity).

$$\begin{aligned} dF_e(\xi) &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d \left(\tilde{k}_{x_i} \right)_{z_j} (-\xi_M(z_j)) \\ &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (x_i - z_j), (-\hat{\omega} z_j - v) \rangle_n. \end{aligned}$$

To turn dF_e into ∇F_e , we will compute $\nabla_\omega F_e$ and $\nabla_v F_e$ separately:

$$\begin{aligned} -a^2 \left(\frac{n-2}{2} \right) \cdot \text{Tr} [(\nabla_\omega F_e) \hat{\omega}] &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (x_i - z_j), (-\hat{\omega} z_j) \rangle_n, \\ b^2 \langle (\nabla_v F_e), v \rangle_n &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (x_i - z_j), (-v) \rangle_n. \end{aligned}$$

To solve for this in coordinates, we will let $\{e^m\}_{m=1}^n$ be the standard orthonormal basis for $(\mathbb{R}^n, \langle \cdot, \cdot \rangle_n)$ and $\{J^{pq}\}_{p < q} := \{E_{pq} - E_{qp}\}_{p < q}$ be as above. Then, the gradient becomes:

$$\begin{aligned} (\nabla_\omega F_e)^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \langle (x_i - z_j), (-J^{pq} z_j) \rangle_n, \\ (\nabla_v F_e)^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \langle (x_i - z_j), (-e^m) \rangle_n. \end{aligned}$$

The above can be simplified by computing the inner product on the right hand side:

$$\begin{aligned} (\nabla_\omega F_e)^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \left(x_i^p z_j^q - x_i^q z_j^p \right), \\ (\nabla_v F_e)^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \left(z_j^m - x_i^m \right). \end{aligned}$$

Likewise, to translate away from the origin, we note that if $h = (R, T) \in \text{SE}(n)$, $(\ell_h)_*(\hat{\omega}, v) = (R\hat{\omega}, Rv)$. Then, if we express the gradient as $\nabla F_h = (\ell_h)_*(\hat{\omega}, v) = (R\hat{\omega}, Rv)$, we get the following expression for $(\hat{\omega}, v) \in \mathfrak{se}(n)$:

$$\begin{aligned} \hat{\omega}^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, \tilde{z}_j) \cdot \left(x_i^p \tilde{z}_j^q - x_i^q \tilde{z}_j^p \right), \\ v^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, \tilde{z}_j) \cdot \left(\tilde{z}_j^m - x_i^m \right), \end{aligned} \tag{28}$$

where $\tilde{z}_j = h^{-1} z_j = R^\top z_j - R^\top T$.

The Hessian calculation will be postponed until Section 8.

6. Analysis and Verification of Idea

It is important to take a moment and examine when solving Problem 2 actually causes the clouds to be best aligned. It is of course impossible to perfectly align two non-identical clouds. Presented below is a discussion of the question: when the two clouds *are* identical, when does the identity element in the group maximize (2)?

Suppose that $Z = X$ and $\ell_Z = \ell_X$. Then the identity ideally should be a fixed-point of (10). This leads to the following, which merely checks the consistency of the framework.

Proposition 12 *Assume that for all $h \in \mathcal{G}$ and $x \in M$, $k(hx, hx) = k(x, x)$. Then the identity is a global maximum of F .*

Proof We have that

$$F(h) = \langle f_X, h.f_X \rangle_{\mathcal{H}}, \quad F(e) = \|f_X\|_{\mathcal{H}}^2 \geq 0.$$

Then using the Cauchy-Schwarz inequality we obtain:

$$|\langle f_X, h.f_X \rangle_{\mathcal{H}}| \leq \|f_X\|_{\mathcal{H}} \cdot \|h.f_X\|_{\mathcal{H}},$$

which is less than $F(e)$ provided that $\|h.f_X\|_{\mathcal{H}} = \|f_X\|_{\mathcal{H}}$. Computing this, we see that

$$\begin{aligned} \|h.f_X\|_{\mathcal{H}}^2 &= \sum_{x_i, z_j \in X} c_{ij} \cdot k(h^{-1}x_i, h^{-1}z_j) \\ &= \sum_{x_i, z_j \in X} c_{ij} \cdot k(x_i, z_j) = \|f_X\|_{\mathcal{H}}^2. \end{aligned}$$

Combining everything, we get that

$$|F(h)| \leq \|f_X\|_{\mathcal{H}} \cdot \|h.f_X\|_{\mathcal{H}} \leq \|f_X\|_{\mathcal{H}}^2 = F(e). \quad \blacksquare$$

Corollary 13 *Suppose $k : M \times M \rightarrow \mathbb{R}$ is a stationary kernel (Williams and Rasmussen, 2006, Page 82), that is $k(x, y) = k(d(x, y))$ for some distance function d . If \mathcal{G} acts isometrically on M , then the identity is a global maximum of F .*

Corollary 14 *The identity is a global maximum for the $\text{SE}(n)$ case.*

Proof This follows from the fact that $\text{SE}(n)$ acts on \mathbb{R}^n isometrically, i.e. $\|hx - hy\|_n = \|x - y\|_n$. ■

Theorem 15 *The maximizer of Problem 2, minimizes the angle between f_X and f_Z .*

Proof Suppose $h^* \in \mathcal{G}$ is the maximizer of Problem 2. Then $\langle f_X, f_Z^* \rangle \geq \langle f_X, f_Z \rangle$ and $\|f_Z^*\|_{\mathcal{H}} \leq \|f_Z\|_{\mathcal{H}}$. Using Cauchy-Schwarz inequality:

$$0 \leq |\langle f_X, f_Z \rangle| \leq |\langle f_X, f_Z^* \rangle| \leq \|f_X\|_{\mathcal{H}} \|f_Z^*\|_{\mathcal{H}} \leq \|f_X\|_{\mathcal{H}} \|f_Z\|_{\mathcal{H}}$$

dividing by $\|f_X\|_{\mathcal{H}} \|f_Z\|_{\mathcal{H}}$ and replacing $\|f_Z\|_{\mathcal{H}}$ in the denominator by $\|f_Z^*\|_{\mathcal{H}}$:

$$0 \leq \cos(\theta) \leq \frac{|\langle f_X, f_Z^* \rangle|}{\|f_X\|_{\mathcal{H}} \|f_Z^*\|_{\mathcal{H}}} \leq \frac{\|f_Z^*\|_{\mathcal{H}}}{\|f_Z^*\|_{\mathcal{H}}} \leq 1$$

$$0 \leq \cos(\theta) \leq \cos(\theta^*) \leq 1$$

$$0 \leq \theta^* \leq \theta \leq \frac{\pi}{2}$$

where $\cos(\theta) = |\langle f_X, f_Z \rangle| / (\|f_X\|_{\mathcal{H}} \|f_Z\|_{\mathcal{H}})$. ■

7. Integrating the Flow for the Special Euclidean Group

Now that we know the direction for the flow, what remains is to determine a way to integrate the flow and to determine a reasonable step size. We integrate using the Lie exponential map to preserve the group structure and the step size is calculated using a 4th-order Taylor approximation in a line search algorithm.

7.1 Integrating

We will take care that in integrating (10), our trajectories will remain on $SE(n)$. This is slightly problematic because integrating is an additive process and $SE(n)$ is not closed under addition. To address this, we note that if in (10), $(\hat{\omega}, v)$ is constant in $\mathfrak{se}(n)$ (i.e. ∇F_h is a left-invariant vector field) the solution is merely

$$(R(t), T(t)) = (R_0, T_0) \exp(t(\hat{\omega}, v)),$$

where $\exp : \mathfrak{se}(n) \rightarrow SE(n)$ is the Lie exponential map (which is merely the matrix exponential). We will exploit this by assuming that $\hat{\omega}$ and v are constant over each time step.

$$\begin{aligned} \begin{bmatrix} R_{k+1} & T_{k+1} \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} R_k & T_k \\ 0 & 1 \end{bmatrix} \cdot \exp \begin{bmatrix} t\hat{\omega} & tv \\ 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} R_k & T_k \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta R & \Delta T \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} R_k \Delta R & R_k \Delta T + T_k \\ 0 & 1 \end{bmatrix}. \end{aligned}$$

Explicit formulas for ΔR and ΔT will be discussed in §7.4 for the special case where $n = 3$. Combining all of this, we get our integration step to be

$$\begin{aligned} R_{k+1} &= R_k \Delta R \\ T_{k+1} &= R_k \Delta T + T_k. \end{aligned} \tag{29}$$

7.2 Step Size

We can use (28) to point in the direction of maximal growth and (29) to find the updated element in $SE(n)$. However, we currently have no intelligent way of choosing t . We will proceed by a Taylor approximation of the solution curve. If we let $G(t) := F(h \exp(t\xi))$, then we want to find the value of t that maximizes G . We compute a 4th-order Taylor expansion of $G(t)$ about $t = 0$ and determine the value of t that maximizes this polynomial.

$$G(t) \approx \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot e^{\alpha_{ij}} \{g_{ij}^1 t + g_{ij}^2 t^2 + g_{ij}^3 t^3 + g_{ij}^4 t^4\}, \tag{30}$$




Figure 6: The images above show X and Z as used in Example 5. The cloud X contains 5282 points while Z contains 6865 points.

where

$$\begin{aligned}
 g_{ij}^1 &= \beta_{ij} \\
 g_{ij}^2 &= \gamma_{ij} + \frac{1}{2}\beta_{ij}^2 \\
 g_{ij}^3 &= \delta_{ij} + \beta_{ij}\gamma_{ij} + \frac{1}{6}\beta_{ij}^3 \\
 g_{ij}^4 &= \varepsilon_{ij} + \beta_{ij}\delta_{ij} + \frac{1}{2}\beta_{ij}^2\gamma_{ij} + \frac{1}{2}\gamma_{ij}^2 + \frac{1}{24}\beta_{ij}^4 \\
 \alpha_{ij} &= \frac{-1}{2\ell^2}\|x_i - z_j\|_n \\
 \beta_{ij} &= \frac{-1}{\ell^2}\langle \hat{\omega}z_j + v, x_i - z_j \rangle_n \\
 \gamma_{ij} &= \frac{-1}{2\ell^2} (\|\hat{\omega}z_j + v\|_n^2 + 2\langle \hat{\omega}^2z_j + v, x_i - z_j \rangle_n) \\
 \delta_{ij} &= \frac{1}{\ell^2} (\langle -\hat{\omega}z_j - v, \hat{\omega}^2z_j + \hat{\omega}v \rangle_n + \langle -\hat{\omega}^3z_j - \hat{\omega}^2v, x_i - z_j \rangle_n) \\
 \varepsilon_{ij} &= \frac{-1}{2\ell^2} (\|\hat{\omega}^2z_j + \hat{\omega}v\|_n^2 + 2\langle \hat{\omega}z_j + v, \hat{\omega}^3z_j + \hat{\omega}^2v \rangle_n + 2\langle \hat{\omega}^4z_j + \hat{\omega}^3v, x_i - z_j \rangle_n)
 \end{aligned}$$

The “optimal” step size is then taken to be the value of $t > 0$ that maximizes the quartic (30).




Figure 7: Left: The clouds X and $A \cdot Z$. Right: The clouds X and $T \cdot A \cdot Z$.

7.3 Special Case: $n = 2$

When we take the special case of $SE(2)$, the gradient (28) takes the special form:

$$\begin{aligned} \omega &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (x_i \tilde{\times} (h^{-1} z_j)) \\ v &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (h^{-1} z_j - x_i), \end{aligned} \quad (31)$$

where $\tilde{\times}$ is the “two-dimensional cross product,”

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \tilde{\times} \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = x_1 y_2 - x_2 y_1.$$

Finally, the exponential map for $SE(2)$ can be solved exactly.

$$\begin{aligned} \Delta R &= \begin{bmatrix} \cos \omega & -\sin \omega \\ \sin \omega & \cos \omega \end{bmatrix}, \\ \Delta T &= \frac{1}{\omega} \begin{bmatrix} \sin \omega & \cos \omega - 1 \\ 1 - \cos \omega & \sin \omega \end{bmatrix} v. \end{aligned} \quad (32)$$

Example 5 We perform a bird’s-eye view map of registration. For the purposes of this example, we choose X to be a contour plot of Matlab’s `peaks(100)` while Z is of `peaks(120)`. Specifically, X is made up of 40 contour lines while Z is made up of 43 lines. For both images, the coordinates are constrained in the following way: $x, y \in [-3, 3]$. The two images are shown in Figure 6.

As these images are initially aligned, we first perturb them and attempt to realign them. The initial perturbation is given by $A \in SE(2)$ while the algorithm determines $T \in SE(2)$. In this sense, we want $T \cdot A = \text{Id}$. The values of A , T , $T \cdot A$, and $\|T \cdot A\|_F$ are shown in (33).

Parameters	Value
transformation convergence threshold ϵ	1e-4
gradient norm convergence threshold ϵ	5e-4
kernel characteristic length-scale ℓ	0.25
kernel characteristic length-scale ℓ (iteration > 3)	0.15
kernel characteristic length-scale ℓ (iteration > 10)	0.10
kernel characteristic length-scale ℓ (iteration > 20)	0.05
kernel signal variance σ	1
minimum step length	0.2
color space inner product scale	1e-5
kernel sparsification threshold	1e-3

Table 3: Parameters used for evaluation of the bird’s eye map registration. The kernel characteristic length-scale is chosen to be adaptive as the algorithm converges; intuitively, we prefer a large neighborhood of correlation for each point, but as the algorithm reaches the convergence reducing the local correlation neighborhood allows for faster convergence and better refinement.

$$\begin{aligned}
 A &= \begin{bmatrix} 0.9654 & -0.2607 & 0.7250 \\ 0.2607 & 0.9654 & 0.6074 \\ 0 & 0 & 1.0000 \end{bmatrix}, \\
 T &= \begin{bmatrix} 0.9659 & 0.2591 & -0.8698 \\ -0.2591 & 0.9659 & -0.4049 \\ 0 & 0 & 1.0000 \end{bmatrix}, \\
 T \cdot A &= \begin{bmatrix} 1.0000 & -0.0017 & -0.0122 \\ 0.0017 & 1.0000 & -0.0060 \\ 0 & 0 & 1.0000 \end{bmatrix}, \\
 \|T \cdot A\|_F &= 0.0138.
 \end{aligned} \tag{33}$$

The parameters used in this registration are included in Table 3. The initial and final images are in Figure 7. For each cloud, the color information is generated based on the level set numbers.

7.4 Special Case: $n = 3$

When we restrict attention to SE(3), the gradient (28) takes a special form:

$$\begin{aligned}
 \omega &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (x_i \times (h^{-1} z_j)) \\
 v &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (h^{-1} z_j - x_i).
 \end{aligned} \tag{34}$$

Additionally, an explicit formula for the exponential map $\exp : \mathfrak{se}(3) \rightarrow \text{SE}(3)$ exists (see Ivancevic and Ivancevic (2011) and Rohan (2013) for example). This gives an exact way to solve (29).

$$\begin{aligned}\Delta R &= I + \left(\frac{\sin t\theta}{\theta}\right) \hat{\omega} + \left(\frac{1 - \cos t\theta}{\theta^2}\right) \hat{\omega}^2, \\ \Delta T &= \left[tI + \left(\frac{1 - \cos t\theta}{\theta^2}\right) \hat{\omega} + \left(\frac{t\theta - \sin t\theta}{\theta^3}\right) \hat{\omega}^2 \right] v.\end{aligned}\tag{35}$$

where $\theta = \|\omega\|_3$ with $\omega \in \mathbb{R}^3$ and t is taken to maximize $G(t)$ in equation (30).

8. Riemannian Geometry of the Special Euclidean Groups

This section deals with the geometry of $\text{SE}(2)$ and $\text{SE}(3)$. Specifically, the Riemannian exponential corresponding to the metric (27) (and the analogous version for $\text{SE}(2)$) as well as the Hessians will be computed. Additionally, there will be a discussion on using Newton's method as an alternate update rule.

8.1 $\text{SE}(2)$

We begin with the simpler case of $\text{SE}(2)$. We will compute the exponential first and the Hessian second.

8.1.1 THE RIEMANN EXPONENTIAL

To compute the Riemann exponential, we will use the Euler-Poincaré equations as described by Theorem 38. The Lagrangian will be $\mathcal{L} : \mathfrak{se}(2) \rightarrow \mathbb{R}$ where

$$\mathcal{L}(\omega; v_1, v_2) = \frac{1}{2}a^2\omega^2 + \frac{1}{2}b^2(v_1^2 + v_2^2).$$

The Euler-Poincaré equations are then

$$\begin{aligned}\dot{\omega} &= 0, \\ \dot{v}_1 &= \omega \cdot v_2, \\ \dot{v}_2 &= -\omega \cdot v_1.\end{aligned}$$

These can be integrated to get a path $(\omega(t); v_1(t), v_2(t)) \in \mathfrak{se}(2)$. However, this only defines a path in the Lie algebra and it needs to be lifted to the group. Doing so results in the following system of differential equations:

$$\begin{aligned}\dot{R}(t) &= R(t) \cdot \omega(t), \\ \dot{T}(t) &= R(t) \cdot \begin{bmatrix} v_1(t) \\ v_2(t) \end{bmatrix}.\end{aligned}$$

The rotation part gives the same answer as the Lie exponential while cancellation occurs in the translation part:

$$\dot{T} = R(t) \cdot \begin{bmatrix} v_1(t) \\ v_2(t) \end{bmatrix} = R(t) \cdot R(t)^{-1} \cdot \begin{bmatrix} v_1(0) \\ v_2(0) \end{bmatrix}.$$




Figure 8: Left: The images of the two point clouds. The blue stars represent X while the red are Z . Center: The vector field corresponding to the eigenvector of the Hessian with the weakest eigenvalue. Right: The vector field corresponding to the eigenvector of the Hessian with the strongest eigenvalue.

This provides us with the Riemann exponential (at the identity), $\text{Exp}_0 : \mathfrak{se}(2) \rightarrow \text{SE}(2)$

$$\text{Exp}_0 \left(\begin{bmatrix} 0 & -\omega & v_1 \\ \omega & 0 & v_1 \\ 0 & 0 & 0 \end{bmatrix} \right) = \begin{bmatrix} \cos \omega & -\sin \omega & v_1 \\ \sin \omega & \cos \omega & v_2 \\ 0 & 0 & 1 \end{bmatrix}.$$

Remark 16 *It is interesting to compare the Riemann and Lie exponential for $\text{SE}(2)$. When we restrict to $\text{SO}(2)$, both agree. This follows from the fact that $\text{SO}(2)$ is a compact Lie group and the metric restricted to this subgroup is bi-invariant. The second interesting phenomena is that the translation part is the identity for the Riemann case while the Lie case is more involved. The reason for this is that in the metric there are no cross terms intertwining rotation with translation, i.e. the Riemann exponential is the Lie exponential for the group $\text{SO}(2) \times \mathbb{R}^2$ as opposed to $\text{SE}(2) = \text{SO}(2) \ltimes \mathbb{R}^2$.*

8.1.2 THE HESSIAN

To calculate the Hessian for the $\text{SE}(2)$ case, we first choose an orthogonal basis $\{e_z, e_1, e_2\} \in \mathfrak{se}(2)$ where

$$e_z = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad e_1 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad e_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

Additionally, we will take $\{E_z, E_1, E_2\}$ to be the corresponding left-invariant vector fields. The connection is given by

$$\begin{aligned} \nabla_{E_z} E_1 &= E_2, \\ \nabla_{E_z} E_2 &= -E_1, \end{aligned}$$

while all other combinations are zero. The gradient in these coordinates is $\omega e_z + v_1 e_1 + v_2 e_2$. The Hessian in these coordinates is

$$-H_{\text{SE}(2)} = \begin{bmatrix} a^2 \mathcal{L}_{E_2} \omega & b^2 \mathcal{L}_{E_2} v^1 + b^2 v^2 & b^2 \mathcal{L}_{E_2} v^2 - b^2 v^1 \\ a^2 \mathcal{L}_{E_1} \omega & b^2 \mathcal{L}_{E_1} v^1 & b^2 \mathcal{L}_{E_1} v^2 \\ a^2 \mathcal{L}_{E_2} \omega & b^2 \mathcal{L}_{E_2} v^1 & b^2 \mathcal{L}_{E_2} v^2 \end{bmatrix}. \quad (36)$$

For the sake of simplicity, let the Hessian be

$$H_{\text{SE}(2)} = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, \tilde{z}_j) \mathcal{S}_{ij},$$

where

$$\mathcal{S}_{ij} = \begin{bmatrix} \frac{1}{\ell^2} (x_i \tilde{\times} \tilde{z}_j)^2 - \langle x_i, \tilde{z}_j \rangle & x_i^2 + \frac{1}{\ell^2} (\tilde{z}_j^1 - x_i^1)(x_i \tilde{\times} \tilde{z}_j) & -x_i^1 + \frac{1}{\ell^2} (\tilde{z}_j^2 - x_i^2)(x_i \tilde{\times} \tilde{z}_j) \\ x_i^2 + \frac{1}{\ell^2} (\tilde{z}_j^1 - x_i^1)(x_i \tilde{\times} \tilde{z}_j) & \frac{1}{\ell^2} (\tilde{z}_j^1 - x_i^1)^2 - 1 & \frac{1}{\ell^2} (x_i^1 - \tilde{z}_j^1)(x_i^2 - \tilde{z}_j^2) \\ -x_i^1 + \frac{1}{\ell^2} (\tilde{z}_j^2 - x_i^2)(x_i \tilde{\times} \tilde{z}_j) & \frac{1}{\ell^2} (x_i^1 - \tilde{z}_j^1)(x_i^2 - \tilde{z}_j^2) & \frac{1}{\ell^2} (\tilde{z}_j^2 - x_i^2)^2 - 1 \end{bmatrix}.$$

Example 6 (Hessian Eigenvalues) *As an initial thought experiment, what happens if each cloud contains only a single point? As we are interested in the registration problem, let us assume that each of these points is at the same location $x \in \mathbb{R}^2$. Then, the Hessian is*

$$H = \frac{1}{\ell^2} c_{11} \begin{bmatrix} -\|x\|^2 & x_2 & -x_1 \\ x_2 & -1 & 0 \\ -x_1 & 0 & -1 \end{bmatrix}.$$

It is interesting to note that this matrix is singular and its kernel is the span of $[-1, -x_2, x_1]^\top$. This Lie algebra element corresponds to a vector field on \mathbb{R}^2 (the infinitesimal generator, see Definition 5). The corresponding vector field is

$$\xi(v) = \begin{bmatrix} v_2 - x_2 \\ -v_1 + x_1 \end{bmatrix},$$

which is merely rotations about the point x . This rotation does not move either cloud and the degeneracy of H predicts this!

In the more general context, the eigenvalues/vectors of H carry much information about the structure and symmetry of the problem. At the solution, the Hessian is negative-definite and so the eigenvector corresponding to the largest eigenvalue (smallest in absolute value) reports on the direction of “greatest symmetry.” That is, suppose that $H\xi = \lambda\xi$ for $\xi \in \mathfrak{g}$ and $|\lambda|$ is “small.” Then moving Z along the vector field ξ_M is not noticeable under F . To see this, we consider the case where X and Z are images of a smiley face as in Figure 8. The Hessian for this example is

$$H = \begin{bmatrix} -422.1054 & 55.3179 & -75.8545 \\ 55.3179 & -16.1852 & 0.8360 \\ -75.8545 & 0.8360 & -23.3209 \end{bmatrix}.$$

The weakest eigenvalue is -0.1907 with a corresponding eigenvector of

$$\xi = [-0.2140, -0.7048, 0.6764]^\top.$$

The vector field for this eigenvector is shown in Figure 8. The impressive aspect of this is that it is possible to learn symmetries in images!

8.2 SE(3)

This section will closely mimic the previous; however it will be more involved.

8.2.1 THE RIEMANN EXPONENTIAL

The Euler-Poincaré equations for SE(3) are

$$\begin{aligned}\dot{\omega} &= 0, \\ \dot{v} &= -\omega \times v.\end{aligned}$$

Integrating these in conjunction with the reconstruction equations produce the Lie exponential for the rotation part and the identity for the translations; the same as the SE(2) case above. That is,

$$\text{Exp}_0(R, T) = (\exp(R), T).$$

Remark 16 holds in this case as well. Moreover, this holds for the higher-dimensional SE(n) as well.

8.2.2 THE HESSIAN

To compute the Hessian, we choose the following orthogonal coordinates for $\mathfrak{se}(3)$:

$$\begin{aligned}e_x &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, & e_y &= \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, & e_z &= \begin{bmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \\ e_1 &= \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, & e_2 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, & e_3 &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.\end{aligned}$$

Likewise, let $\{E_x, E_y, E_z, E_1, E_2, E_3\}$ be their corresponding left-invariant vector fields. With this, the connection is given by

$$\begin{aligned}\nabla_{E_x} E_y &= \frac{1}{2} E_z, \nabla_{E_x} E_z = -\frac{1}{2} E_y, \nabla_{E_y} E_z = \frac{1}{2} E_x, \\ \nabla_{E_y} E_x &= -\frac{1}{2} E_z, \nabla_{E_z} E_x = \frac{1}{2} E_y, \nabla_{E_z} E_y = -\frac{1}{2} E_x, \\ \nabla_{E_x} E_2 &= E_3, \nabla_{E_x} E_3 = -E_2, \nabla_{E_y} E_3 = E_1, \\ \nabla_{E_y} E_1 &= -E_3, \nabla_{E_z} E_1 = E_2, \nabla_{E_z} E_2 = -E_1,\end{aligned}\tag{37}$$

while all other combinations are zero.

Remark 17 Notice that the coefficients for the connection in (37) do not depend on either a or b ! This is to be expected because the Levi-Civita connection needs to be torsion free. This means that $\nabla_X Y - \nabla_Y X = [X, Y]$. Since the coordinates we are taking are the left-invariant vector fields, their bracket is just the bracket of the Lie algebra which does not depend on a nor b .

We are now ready to write down the Hessian matrix (we will assume that $a = b = 1$).

$$\begin{aligned}
 -H_{\text{SE}(3)} &= \\
 &\left[\begin{array}{ccc|ccc}
 \mathcal{L}_{E_x}\omega_x & \mathcal{L}_{E_x}\omega_y + \frac{1}{2}\omega_z & \mathcal{L}_{E_x}\omega_z - \frac{1}{2}\omega_y & \mathcal{L}_{E_x}v_1 & \mathcal{L}_{E_x}v_2 + v_3 & \mathcal{L}_{E_x}v_3 - v_2 \\
 \mathcal{L}_{E_y}\omega_x - \frac{1}{2}\omega_z & \mathcal{L}_{E_y}\omega_y & \mathcal{L}_{E_y}\omega_z + \frac{1}{2}\omega_x & \mathcal{L}_{E_y}v_1 - v_3 & \mathcal{L}_{E_y}v_2 & \mathcal{L}_{E_y}v_3 + v_1 \\
 \mathcal{L}_{E_z}\omega_x + \frac{1}{2}\omega_y & \mathcal{L}_{E_z}\omega_y - \frac{1}{2}\omega_x & \mathcal{L}_{E_z}\omega_z & \mathcal{L}_{E_z}v_1 + v_2 & \mathcal{L}_{E_z}v_2 - v_1 & \mathcal{L}_{E_z}v_3 \\
 \hline
 \mathcal{L}_{E_1}\omega_x & \mathcal{L}_{E_1}\omega_y & \mathcal{L}_{E_1}\omega_z & \mathcal{L}_{E_1}v_1 & \mathcal{L}_{E_1}v_2 & \mathcal{L}_{E_1}v_3 \\
 \mathcal{L}_{E_2}\omega_x & \mathcal{L}_{E_2}\omega_y & \mathcal{L}_{E_2}\omega_z & \mathcal{L}_{E_2}v_1 & \mathcal{L}_{E_2}v_2 & \mathcal{L}_{E_2}v_3 \\
 \mathcal{L}_{E_3}\omega_x & \mathcal{L}_{E_3}\omega_y & \mathcal{L}_{E_3}\omega_z & \mathcal{L}_{E_3}v_1 & \mathcal{L}_{E_3}v_2 & \mathcal{L}_{E_3}v_3
 \end{array} \right] \\
 &=: - \left[\begin{array}{c|c}
 \mathcal{A} & \mathcal{B} \\
 \hline
 \mathcal{C} & \mathcal{D}
 \end{array} \right]. \tag{38}
 \end{aligned}$$

Notice that the first 3 by 3 block, \mathcal{A} , is identical to (22). This reflects the copy of $\text{SO}(3)$ lying inside $\text{SE}(3)$.

For the sake of brevity, we will use the following abbreviation:

$$H_{\text{SE}(3)} = \frac{1}{\ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, \tilde{z}_j) \cdot \left[\begin{array}{c|c}
 \mathcal{A}_{ij} & \mathcal{B}_{ij} \\
 \hline
 \mathcal{C}_{ij} & \mathcal{D}_{ij}
 \end{array} \right].$$

Notice that by symmetry of H , $\mathcal{B}_{ij}^\top = \mathcal{C}_{ij}$. Therefore, we will only compute \mathcal{A}_{ij} , \mathcal{C}_{ij} , and \mathcal{D}_{ij} . To simplify expressions, we will call $v = [v_1, v_2, v_3]^\top := x_i \times \tilde{z}_j$ and $u = [u_1, u_2, u_3]^\top := \tilde{z}_j - x_i$.

$$\begin{aligned}
 \mathcal{A}_{ij} &= [\mathcal{A}_{ij}^1 \quad \mathcal{A}_{ij}^2 \quad \mathcal{A}_{ij}^3] \\
 \mathcal{A}_{ij}^1 &= \begin{bmatrix} \frac{1}{\ell^2}v_1^2 - (x_i^2\tilde{z}_j^2 + x_i^3\tilde{z}_j^3) \\ \frac{1}{\ell^2}v_1v_2 + \frac{1}{2}(x_i^1\tilde{z}_j^2 + x_i^2\tilde{z}_j^1) \\ \frac{1}{\ell^2}v_1v_3 + \frac{1}{2}(x_i^1\tilde{z}_j^3 + x_i^3\tilde{z}_j^1) \end{bmatrix} \\
 \mathcal{A}_{ij}^2 &= \begin{bmatrix} \frac{1}{\ell^2}v_1v_2 + \frac{1}{2}(x_i^1\tilde{z}_j^2 + x_i^2\tilde{z}_j^1) \\ \frac{1}{\ell^2}v_2^2 - (x_i^1\tilde{z}_j^1 + x_i^3\tilde{z}_j^3) \\ \frac{1}{\ell^2}v_2v_3 + \frac{1}{2}(x_i^2\tilde{z}_j^3 + x_i^3\tilde{z}_j^2) \end{bmatrix} \\
 \mathcal{A}_{ij}^3 &= \begin{bmatrix} \frac{1}{\ell^2}v_1v_3 + \frac{1}{2}(x_i^1\tilde{z}_j^3 + x_i^3\tilde{z}_j^1) \\ \frac{1}{\ell^2}v_2v_3 + \frac{1}{2}(x_i^2\tilde{z}_j^3 + x_i^3\tilde{z}_j^2) \\ \frac{1}{\ell^2}v_3^2 - (x_i^1\tilde{z}_j^1 + x_i^2\tilde{z}_j^2) \end{bmatrix} \\
 \mathcal{C}_{ij} &= \begin{bmatrix} \frac{1}{\ell^2}u_1v_1 & -x_i^3 + \frac{1}{\ell^2}u_1v_2 & x_i^2 + \frac{1}{\ell^2}u_1v_3 \\ x_i^3 + \frac{1}{\ell^2}u_2v_1 & \frac{1}{\ell^2}u_2v_2 & -x_i^1 + \frac{1}{\ell^2}u_2v_3 \\ -x_i^2 + \frac{1}{\ell^2}u_3v_1 & x_i^1 + \frac{1}{\ell^2}u_3v_2 & \frac{1}{\ell^2}u_3v_3 \end{bmatrix} \\
 \mathcal{D}_{ij} &= \begin{bmatrix} \frac{1}{\ell^2}u_1^2 - 1 & \frac{1}{\ell^2}u_1u_2 & \frac{1}{\ell^2}u_1u_3 \\ \frac{1}{\ell^2}u_1u_2 & \frac{1}{\ell^2}u_2^2 - 1 & \frac{1}{\ell^2}u_2u_3 \\ \frac{1}{\ell^2}u_1u_3 & \frac{1}{\ell^2}u_2u_3 & \frac{1}{\ell^2}u_3^2 - 1 \end{bmatrix}.
 \end{aligned}$$

Example 7 (Hessian Eigenvalues, Cont.) *The same analysis as presented in Example 6 can be carried out for SE(3) as well. Again, we will use the Hessian to identify symmetries within the point clouds. Specifically, we will consider the case where both X and Z are a sphere living in \mathbb{R}^3 . The spheres can be aligned via translations but rotations do nothing. This should manifest in the Hessian; the block \mathcal{D} in (38) should dominate the Hessian while A should be the smallest.*

In what follows, X contains 1200 points on $S^2 \in \mathbb{R}^3$ and Z contains 1500 points (both of these clouds are already on S^2 so the registration problem is already solved) and all $c_{ij} = 3$. The remaining parameters are $\ell = 0.25$ and $\sigma = 10^{-4}$. The Hessian is

$$H = \begin{bmatrix} 1.3995 & -4.0561 & -1.4875 & 2.1207 & 0.9970 & 1.2481 \\ -4.0561 & -6.3054 & -0.3125 & 0.7368 & -1.8921 & 0.8737 \\ -1.4875 & -0.3125 & 0.0554 & 1.5803 & -0.3814 & -0.2286 \\ \hline 2.1207 & 0.7368 & 1.5803 & -594.3227 & 1.9310 & 23.2125 \\ 0.9970 & -1.8921 & -0.3814 & 1.9310 & -603.8121 & -13.3898 \\ 1.2481 & 0.8737 & -0.2286 & 23.2125 & -13.3898 & -561.8939 \end{bmatrix}. \quad (39)$$

It can be seen that the translation component is on the order of 100 times larger than the rotation component. More formally, we can examine the eigenvalues/vectors of H .

Table 4 shows that the eigenvalues corresponding to rotations are around 100 times smaller than the eigenvalues corresponding to translations. As a result, the Hessian can report on the regularity of the solution / identify infinitesimal symmetries in the pictures.

9. Experimental Results: RGB-D Visual Odometry

The RGB-D visual odometry problem is illustrated in Figure 9. In the context of visual odometry, we call our method Continuous Visual Odometry (CVO), which has appeared in our earlier work (Ghaffari et al., 2019; Lin et al., 2019). We compare CVO with the state-of-the-art direct (and dense) RGB-D visual odometry (DVO) (Kerl et al., 2013a,b). Since the original DVO source code requires outdated ROS dependency (Kerl, 2013), we reproduced DVO results using the version provided by Matthieu Pizenberg (Pizenberg, 2019), which only removes the dependency for ROS while maintaining the DVO core source code unchanged.

To improve the computational efficiency, we adopted a similar approach to Direct Sparse Visual Odometry (DSO) by Engel et al. (2014) to create a semi-dense point cloud (around

Eigenvalue	-613.9730	-598.4636	-547.6175	-8.1304	-0.2898	3.5950
Eigenvector	$\begin{bmatrix} 0.0000 \\ 0.0022 \\ 0.0021 \\ -0.5827 \\ 0.6858 \\ 0.4360 \end{bmatrix}$	$\begin{bmatrix} -0.0032 \\ 0.0016 \\ -0.0015 \\ 0.6897 \\ 0.7011 \\ -0.1810 \end{bmatrix}$	$\begin{bmatrix} 0.0033 \\ 0.0027 \\ 0.0010 \\ -0.4298 \\ 0.1952 \\ -0.8815 \end{bmatrix}$	$\begin{bmatrix} 0.4029 \\ 0.9089 \\ 0.1073 \\ 0.0030 \\ -0.0023 \\ 0.0025 \end{bmatrix}$	$\begin{bmatrix} 0.2685 \\ -0.2295 \\ 0.9355 \\ 0.0032 \\ 0.0006 \\ -0.0000 \end{bmatrix}$	$\begin{bmatrix} -0.8750 \\ 0.3481 \\ 0.3365 \\ -0.0019 \\ -0.0027 \\ -0.0015 \end{bmatrix}$

Table 4: The eigenvalue/vectors of H in (39).




Figure 9: Visual odometry using depth cameras is the problem of finding a rigid body transformation between two colored point clouds.

Parameters	Symbol	Value
Transformation convergence threshold	ϵ_1	$1e-5$
Gradient norm convergence threshold	ϵ_2	$5e-5$
Minimum step length		0.2
Kernel sparsification threshold		$8.315e-3$
Spatial kernel initial length-scale	l_{init}	0.1
Spatial kernel signal variance	σ	0.1
Color kernel length-scale	l_c	0.1
Color kernel signal variance	σ_c	1

Table 5: Parameters used for evaluation using TUM RGB-D Benchmark, similar values are chosen for all experiments. The kernel characteristic length-scale is chosen to be adaptive as the algorithm converges (Ghaffari et al., 2019); intuitively, we prefer a large neighborhood of correlation for each point, but as the algorithm reaches the convergence reducing the local correlation neighborhood allows for faster convergence and better refinement.

3000 points) for each scan. To prevent insufficient points being selected in environments that lack rich visual information, we also used a Canny edge detector (Canny, 1987) from OpenCV (Bradski, 2000). When the number of points selected by the DSO point selector is less than one-third of the desired number of points, more points will be selected by downsampling the pixels highlighted by the Canny detector. While generating the point cloud, RGB values are first transformed into the HSV colormap and normalized. The normalized HSV values are then combined with the normalized intensity gradients and used as the labels of the selected points in the color space. For all experiments, we used the same set of parameters, which are listed in Table 5.

All experiments are performed on a Dell XPS15 9750 laptop with Intel i7-8750H CPU (6 cores with 2.20 GHz each) and 32GB RAM. The source code is implemented in C++

Sequence	Training				Validation			
	CVO		DVO		CVO		DVO	
	Trans.	Rot.	Trans.	Rot.	Trans.	Rot.	Trans.	Rot.
fr1/desk	0.0486	2.4860	0.0387	2.3589	0.0401	2.0148	0.0371	2.0645
fr1/desk2	0.0535	3.0383	0.0583	3.6529	0.0225	1.7691	0.0208	1.7416
fr1/room	0.0560	2.4566	0.0518	2.8686	0.0446	3.9183	0.2699	7.4144
fr1/360	0.0991	3.0025	0.1602	4.4407	0.1420	3.0746	0.2811	7.0876
fr1/teddy	0.0671	4.8089	0.0948	2.5495	n/a	n/a	n/a	n/a
fr1/floor	0.0825	2.3745	0.0635	2.2805	n/a	n/a	n/a	n/a
fr1/xyz	0.0240	1.1703	0.0327	1.8751	0.0154	1.3872	0.0453	3.0061
fr1/rpy	0.0457	3.3073	0.0336	2.6701	0.1138	3.6423	0.3607	7.9991
fr1/plant	0.0316	1.9973	0.0272	1.5523	0.0630	4.9185	0.0660	2.5865
Average	0.0561	2.7380	0.0623	2.6943	0.0631	2.9607	0.1544	4.5571

Table 6: The RMSE of Relative Pose Error (RPE) for `fr1` sequences. The trans. columns show the RMSE of the translational drift in m/sec and the rot. columns show the RMSE of the rotational error in deg/sec. There’s no corresponding validation data sets for `fr1/teddy` and `fr1/floor`.

and compiled with the Intel Compiler. The kernel computations are parallelized using the Intel Threading Building Blocks (TBB) (Intel Corporation, 2019). Using compiler auto-vectorization and the parallelization, the average time for frame-to-frame registration is 0.5 sec. Software is available for download at <https://github.com/MaaniGhaffari/cvo-rgbd>.

9.1 TUM RGB-D Benchmark

We performed experiments on two parts of RGB-D SLAM data set and benchmark by the Technical University of Munich (Sturm et al., 2012). This data set was collected indoors with a Microsoft Kinect using a motion capture system as a proxy for ground truth trajectory. For all tracking experiments, the entire images were used sequentially without any skipping, i.e., at full frame rate. We evaluated CVO and DVO on the training and validation sets for all the `fr1` sequences and the structure versus texture sequences. RGB-D benchmark tools (Sturm et al., 2012) were then used to evaluate the Relative Pose Error (RPE) of both methods.

Table 6 shows the Root-Mean-Squared Error (RMSE) of the RPE for `fr1` sequences. The Trans. columns show the RMSE of the translational drift in m/sec and the Rot. columns show the RMSE of the rotational drift in deg/sec. There are no corresponding validation sequences for `fr1/teddy` and `fr1/floor`. On the validation set, CVO has improved performance compared with DVO which shows it can generalize across different scenarios better. From the results, we can see that CVO is intrinsically robust. The next experiment will further reveal that CVO has the advantage of performing well in extreme environments that lack rich structure or texture.

Sequence			Training				Validation			
			CVO		DVO		CVO		DVO	
structure-texture-dist.			Trans.	Rot.	Trans.	Rot.	Trans.	Rot.	Trans.	Rot.
×	✓	near	0.0279	1.3470	0.0563	1.7560	0.0310	1.6367	0.0315	1.1498
×	✓	far	0.0609	1.2342	0.1612	3.4135	0.1374	2.3929	0.5351	8.2529
✓	×	near	0.0221	1.3689	0.1906	10.6424	0.0465	2.0359	0.1449	4.9022
✓	×	far	0.0372	1.3061	0.1171	2.4044	0.0603	1.8142	0.1375	2.2728
✓	✓	near	0.0236	1.2972	0.0175	0.9315	0.0306	1.8694	0.0217	1.2653
✓	✓	far	0.0409	1.1640	0.0171	0.5717	0.0616	1.4760	0.0230	0.6312
×	×	near	0.2119	9.7944	0.3506	13.3127	0.1729	5.8674	0.1747	6.0443
×	×	far	0.0799	3.0978	0.1983	6.8419	0.0899	2.6199	0.2000	6.5192
Average			0.0631	2.5762	0.1386	4.9843	0.0787	2.4640	0.1586	3.8797

Table 7: The RMSE of Relative Pose Error (RPE) for the structure v.s texture sequence. The Trans. columns show the RMSE of the translational drift in m/sec and the Rot. columns show the RMSE of the rotational error in deg/sec. The ✓ means the sequence has structure/texture and × means the sequence does not have structure/texture. The results show while DVO performs better in structure and texture case, CVO has significantly better performance in the environments that lack structure and texture.

9.2 Structure vs. Texture Sequences

Table 7 shows the RMSE of RPE for the structure vs. texture sequences. This data set contains image sequences in structure/nostructure and texture/nottexture environments. As elaborated by Ghaffari et al. (2019), by treating point clouds as points in the function space (RKHS), CVO is inherently robust to the lack of features in the environment. CVO shows the best performance on cases where either structure or texture is not rich in the environment. This reinforces the claim that CVO is robust to such scenes.

We also note that DVO has the best performance on the case where the environment contains rich texture and structure information. This could be due to two factors: 1) CVO adopts a semi-dense point cloud construction from DSO (Engel et al., 2018), while DVO uses the entire dense image without subsampling. Although the semi-dense tracking approach of Engel et al. (2014, 2018) is computationally attractive and we advocate it, the semi-dense point cloud construction process used in this work is a heuristic process and might not necessarily capture the relevant information in each frame optimally; 2) DVO uses a motion prior as regularizer whereas CVO solely depends on the camera information with no regularizer. We conjecture this latter is the reason DVO, relative to the training set, does not perform well on validation sequences. The motion prior is a useful assumption when it is true. It can help to tune the method better on the training sets but if the assumption gets violated can lead to poor performance. The addition of an IMU sensor of course can improve the performance of both methods and is an interesting future research direction.




Figure 10: Left: The images of the two point clouds. The blue stars represent X while the red are Z . Center: A contour plot of $F : T^2 \rightarrow \mathbb{R}$ where $\ell = 0.5$. Right: A contour plot of $F : T^2 \rightarrow \mathbb{R}$ where $\ell = 0.2$. A total of 33 local maxima are found in this picture. A video of the effect of varying ℓ is available at <https://youtu.be/ETr6-c0VapQ>.




Figure 11: The images above display the number of local maxima of the function $F : T^2 \rightarrow \mathbb{R}$ as a function of the length-scale. The black curve shows the exact number of local maxima as reported by Matlab’s “`imregionalmax`” while the blue curve is a plot of $N = 1.16602 \cdot \ell^{-2}$.

10. Discussions and Further Considerations

In this section, we present discussions on a number of topics that are worth investigating.

10.1 Locality of Solutions

It is important to stress that the registration problem laid out here is only a local solver. However, how detrimental is this? For the purposes of gaining intuition, we will compute $F : T^2 \rightarrow \mathbb{R}$ for the T^2 case, see Section 3.2. For the point clouds, we will use the same clouds as in Example 6 (although viewing them in $[0, 2\pi)^2 \sim T^2$ rather than \mathbb{R}^2). The torus

is discretized into a 500 by 500 grid and length-scales are chosen in the range $[0.01, 0.5]$. For each value of ℓ , a contour plot of $F : T^2 \rightarrow \mathbb{R}$ is constructed and the number of local maxima is recorded (with Matlab’s “`imregionalmax`” function). Two contour plots corresponding to $\ell = 0.5$ and $\ell = 0.2$ are shown in Figure 10.

It is important to note that there exists a plethora of local maxima. This illustrates how important it is to have a good initialization, or else the gradient flow will end up getting trapped at a false solution. However, the number of local solutions depends heavily on ℓ . This can be used in practice by having ℓ start off large and shrink between iterations, even though care must be taken if one deals with positive-definite kernel design and hyperparameter learning on curved spaces as the Gaussian kernel might not be positive-definite for all ℓ (Feragen et al., 2015; Borovitskiy et al., 2020).

Another observation is that the number of local solutions seems to grow according to a power law of the length-scale, as shown in Figure 11. Controlling the number of modes is an important problem because if we can control the number of modes, we can initially generate a convex problem and guide the algorithm to the global solution. This approach can be seen as the continuous analogue of the image pyramid method in computer vision (Szeliski, 2010) where low-resolution images are processed first, and the next level has a higher resolution initialized by the previous step.

A heuristic explanation can be that we have more modes of the Gaussian kernel at lower bandwidths since a higher bandwidth makes the bumps around points coalesce. Assuming an inversely proportional relation to the bandwidth, then Figure 11 is a reasonable verification of the number of maxima. Nevertheless, a systematic approach for learning the bandwidth and controlling the complexity of the problem remains an open problem and is an interesting future study.

10.2 Reduction to the Weighted Least Squares

The standard solution to sensor registration is typically formulated as a maximum likelihood or least-squares problem and solved via numerical optimization techniques. In the maximum likelihood setting, the measurement noise distribution significantly affects the robustness of the solver to outliers. Hence, in practice, the problem is necessarily solved using a robust estimator such as the iteratively reweighted least-squares algorithm or truncated least squares. The central shortcoming of this approach is the reliance on point-wise correspondences between two measurement sets. As a result, the implicit assumption is that the perceived data points by digital sensors are from the same points in reality. Of course, this is seldom true as the nature of digital sensing means a discretization of the reality is captured at best. Sparse and uniform (feature-scarce) measurement sets are where this shortcoming can become problematic. A prime example is the lack of robustness of the visual odometry algorithms to texture and structure-less scenes.

Without loss of generality, we will use the Gaussian (Squared Exponential) kernel based on a distance function, $d(\cdot, \cdot)$:

$$k(x, y) = \sigma^2 \exp(-d(x, y)^2).$$

Further, let us define $w_{ij} := \sigma^2 \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}}$ and $d_{ij} := d(x_i, z_j)$. Then we have

$$\langle f_X, f_Z \rangle_{\mathcal{H}} = \sum_{\substack{x_i \in X \\ z_j \in Z}} w_{ij} \cdot \exp(-d_{ij}^2). \quad (40)$$

Now we explicitly apply the three assumptions that are used in the least squares problem.

Assumption 1 *The assumptions used in the least squares problem are:*

1. *The point clouds X and Z have the same number of points.*
2. *The point cloud Z is ordered such that the matching indices are corresponding measurements.*
3. *The residuals are only computed between corresponding measurements.*

By applying the above-mentioned assumptions to (40), we get

$$\langle f_X, f_Z \rangle_{\mathcal{H}} \approx \sum_i w_i \cdot \exp(-d_i^2),$$

where $w_i = w_{ii}$ and $d_i = d_{ii}$. and it is reduced to a single sum over matching indices of the double sum in (40). Finally, the exp can be approximated using a Taylor expansion as

$$\exp(-d_i^2) \approx 1 - d_i^2 + \dots$$

Problem 3 *The problem of aligning the point clouds can now be rephrased as the following weighted least squares form*

$$\arg \min_{h \in \mathcal{G}} S(h), \quad S(h) := \sum_i w_i \cdot d(x_i, h z_i)^2.$$

10.3 Feature Embedding via Tensor Product Representation

For completeness, we include an extension of the proposed framework developed by Zhang et al. (2020) that is of particular interest to practitioners. The framework contains a specialized solver implemented using GPU that exploits the double sum’s sparse structure in the objective function. In addition, it incorporates deep learning-based features such as semantic segmentation labels by extending the feature space to a hierarchical distributed representation. For extensive experimental results and open source software that validates presented topics in this work, please see Zhang et al. (2020).

Let (V_1, V_2, \dots) be different inner product spaces describing different types of non geometric features of a point, such as color, intensity, and semantics. Their tensor product, $V_1 \otimes V_2 \otimes \dots$ is also an inner product space. For any $x \in X, z \in Z$ with features $\ell_X(x) = (u_1, u_2, \dots)$ and $\ell_Z(z) = (v_1, v_2, \dots)$, with $u_1, v_1 \in V_1, u_2, v_2 \in V_2, \dots$, we have

$$\langle \ell_X(x), \ell_Z(z) \rangle_{\mathcal{I}} = \langle u_1 \otimes u_2 \otimes \dots, v_1 \otimes v_2 \otimes \dots \rangle = \langle u_1, v_1 \rangle \cdot \langle u_2, v_2 \rangle \cdot \dots \quad (41)$$

By substituting (41) into (4), we obtain

$$\langle f_X, f_{h^{-1}Z} \rangle_{\mathcal{H}} = \sum_{x_i \in X, z_j \in Z} \langle u_{1i}, v_{1j} \rangle \cdot \langle u_{2i}, v_{2j} \rangle \dots k(x_i, h^{-1}z_j)$$

After applying the kernel trick (Bishop, 2006; Williams and Rasmussen, 2006; Murphy, 2012) we arrive at

$$\langle f_X, f_{h^{-1}Z} \rangle_{\mathcal{H}} = \sum_{x_i \in X, z_j \in Z} k(x_i, h^{-1}z_j) \cdot \prod_k k_{V_k}(u_{ki}, v_{kj}) := \sum_{x_i \in X, z_j \in Z} k(x_i, h^{-1}z_j) \cdot c_{ij}. \quad (42)$$

Equation (42) describes the full geometric and non-geometric relationship between the two point clouds. Each c_{ij} does not depend on the relative transformation; thus, it will be a constant when computing the gradient and the step size for solving registration problems. Interestingly, the double sum in (42) is sparse because a point $x_i \in X$ is far away from the majority of the points $z_j \in Z$, either in the spatial (geometry) space or one of the feature (semantic) spaces.

This formulation can be further simplified to a purely geometric model, if we let the label functions $\ell_X(x_i) = \ell_Z(z_j) = 1$. Then (42) becomes

$$\langle f_X, f_{h^{-1}Z} \rangle_{\mathcal{H}} = \sum_{x_i \in X, z_j \in Z} k(x_i, h^{-1}z_j). \quad (43)$$

Through (43), the proposed method can register point clouds that do not have appearance measurements. It is worth noting that, when choosing the squared exponential kernel, (43) has the same formulation as Kernel Correlation (Tsin and Kanade, 2004).

10.4 Non-Isotropic Kernel

For the 3D space, the kernel used throughout is

$$k(x, y) = \sigma^2 \exp\left(-\frac{\|x - y\|^2}{2\ell^2}\right).$$

where $\ell \in \mathbb{R}$. What happens if we change so $\ell \in \mathbb{R}^3$ so the length-scale changes depending on the direction? This makes the kernel into

$$k_{\Lambda}(x, y) = \sigma^2 \exp(-\langle x - y, \Lambda(x - y) \rangle), \quad (44)$$

where Λ is positive-definite and symmetric. Going through the same calculations as before, we get that

$$\begin{aligned} \omega &= \frac{1}{a^2} \Lambda \sum_{ij} c_{ij} k_{\Lambda}(x_i, h^{-1}z_j) (x_i \times (h^{-1}z_j)), \\ v &= \frac{1}{b^2} \Lambda \sum_{ij} c_{ij} k_{\Lambda}(x_i, h^{-1}z_j) (x_i - h^{-1}z_j). \end{aligned} \quad (45)$$

Note that in the original case,

$$\Lambda = \begin{bmatrix} \frac{1}{2\ell^2} & 0 & 0 \\ 0 & \frac{1}{2\ell^2} & 0 \\ 0 & 0 & \frac{1}{2\ell^2} \end{bmatrix}.$$

In terms of applications, modeling directional similarities has been successful in the Generalized-ICP framework developed by Segal et al. (2009). In particular, it was shown that in structured environments, e.g., flat walls, ceiling, and floor, defining a norm weighted by local empirical covariances of the two point clouds effectively reduces the pose estimation error. If such a distance metric is used inside the kernel, the kernel can be considered to be anisotropic. In terms of modeling, this is similar to a covariance function with automatic relevance determination (Neal, 1996).

10.5 Matérn Kernel

Rather than using the Gaussian (squared exponential) kernel, we will use the Matérn kernel (Williams and Rasmussen, 2006). The following derivations are only valid, i.e., positive definite, for all hyperparameters in Euclidean spaces. For the analogue of squared exponential and Matérn kernels on compact Riemannian manifold please see the work of Borovitskiy et al. (2020).

This is an infinite family of kernels. The Matérn Kernel, for points a distance d apart is

$$C_\nu(d) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{d}{\rho} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right), \quad (46)$$

where Γ is the gamma function, K_ν is the Bessel function of the second kind, and ρ and ν are positive parameters. We have the following derivative property:

$$\frac{d}{dt} K_\nu(t) = -K_{\nu-1}(t) - \frac{\nu}{t} K_\nu(t).$$

With half-integers, $\nu = p + 1/2$, $p \in \mathbb{N}$

$$C_{p+1/2}(d) = \sigma^2 \exp \left(-\frac{\sqrt{2p+1}d}{\rho} \right) \frac{p!}{(2p)!} \sum_{i=0}^p \frac{(p+i)!}{i!(p-i)!} \left(\frac{2\sqrt{2p+1}d}{\rho} \right)^{p-i}.$$

Differentiating the kernel with respect to d , we see that

$$\begin{aligned} \frac{\partial}{\partial d} C_\nu(d) &= \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left[\nu \sqrt{2\nu} \left(\sqrt{2\nu} \frac{d}{\rho} \right)^{\nu-1} K_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right) \right. \\ &\quad \left. + \left(\sqrt{2\nu} \frac{d}{\rho} \right)^\nu \frac{\sqrt{2\nu}}{\rho} \left[-K_{\nu-1} \left(\sqrt{2\nu} \frac{d}{\rho} \right) - \frac{\nu\rho}{\sqrt{2\nu}d} K_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right) \right] \right] \end{aligned}$$

Working with this does not seem appealing, so let us work with half-integer values instead. Below, the kernels and their corresponding gradients are listed. In what follows, $d_{ij} = \|x_i - h^{-1}z_j\|$.

10.5.1 $\nu = 1/2$

$$C_{1/2}(d) = \sigma^2 \exp\left(-\frac{d}{\rho}\right) \quad (47)$$

$$\begin{aligned} \omega &= \frac{1}{a^2\rho} \sum_{ij} c_{ij} \exp\left(-\frac{d_{ij}}{\rho}\right) (x_i \times h^{-1}z_j), \\ v &= \frac{1}{b^2\rho} \sum_{ij} c_{ij} \exp\left(-\frac{d_{ij}}{\rho}\right) (x_i - h^{-1}z_j). \end{aligned} \quad (48)$$

10.5.2 $\nu = 3/2$

$$C_{3/2}(d) = \sigma^2 \left(1 + \frac{\sqrt{3}d}{\rho}\right) \exp\left(-\frac{\sqrt{3}d}{\rho}\right) \quad (49)$$

$$\begin{aligned} \omega &= \frac{3}{a^2\rho^2} \sum_{ij} c_{ij}d_{ij} \exp\left(-\frac{\sqrt{3}d_{ij}}{\rho}\right) (x_i \times h^{-1}z_j), \\ v &= \frac{3}{b^2\rho^2} \sum_{ij} c_{ij}d_{ij} \exp\left(-\frac{\sqrt{3}d_{ij}}{\rho}\right) (x_i - h^{-1}z_j). \end{aligned} \quad (50)$$

10.5.3 $\nu = 5/2$

$$C_{5/2}(d) = \sigma^2 \left(1 + \frac{\sqrt{5}d}{\rho} + \frac{5d^2}{3\rho^2}\right) \exp\left(-\frac{\sqrt{5}d}{\rho}\right) \quad (51)$$

$$\begin{aligned} \omega &= \frac{5}{3a^2\rho^3} \sum_{ij} c_{ij}d_{ij} (\rho + \sqrt{5}d_{ij}) \exp\left(-\frac{\sqrt{5}d_{ij}}{\rho}\right) (x_i \times h^{-1}z_j), \\ v &= \frac{5}{3b^2\rho^3} \sum_{ij} c_{ij}d_{ij} (\rho + \sqrt{5}d_{ij}) \exp\left(-\frac{\sqrt{5}d_{ij}}{\rho}\right) (x_i - h^{-1}z_j). \end{aligned} \quad (52)$$

10.6 Relationship with Deep Learning

Point clouds obtained by modern sensors such as RGB-D cameras, stereo cameras, and LIDARs contain up to 300,000 points per scan at 10 – 60 Hz and rich color and intensity (reflectivity of a material sensed by an active light beam) measurements besides the geometric information. In addition, deep learning (LeCun et al., 2015) can provide semantic attributes of the scene as measurements (Long et al., 2015; Chen et al., 2017; Zhu et al., 2019).

Representation learning provides a way to perform *semi-supervised learning* (Chapelle et al., 2006), *self-supervised learning* (Dahlkamp et al., 2006; Sofman et al., 2006; Doersch and Zisserman, 2017; Pathak et al., 2017), and *unsupervised learning* (Goodfellow et al., 2016). Although the motivation behind learning a “good” representation from data varies across different fields and research communities, its mathematical foundation is moving towards a more unifying direction (Bengio et al., 2013; Litjens et al., 2017) where exploiting

well-established classical tools and modern data-driven approaches such as deep learning seem inevitable (Girshick et al., 2014). Perhaps the main undesirable aspect of modern deep learning-based methods is the heavy dependency on large labeled data, i.e., *supervised learning*.

An intrinsic metric available in our framework is the angle, θ , between two functions. This indicator can be computed to track the optimization progress. The cosine of the angle is defined as

$$\cos(\theta) = \frac{\langle f_X, f_Z \rangle_{\mathcal{H}}}{\|f_X\| \cdot \|f_Z\|}. \quad (53)$$

However, calculating $\|f_X\|$ and $\|f_Z\|$ is time-consuming as it requires evaluating the double sum for each of the two point clouds. To approximate (53), we use the following result.

Remark 18 *Suppose $k(x_i, x_j) = \delta_{ij}$ and $c_{ii} = 1$, where δ_{ij} is the Kronecker delta, then $\|f_X\| = \sqrt{|X|}$.*

Corollary 19 *Using the previous assumption, we define the following alignment indicator.*

$$i_{\theta} := \frac{1}{\sqrt{|X| \cdot |Z|}} \sum_{x_i \in X, z_j \in Z} c_{ij} \cdot k(x_i, z_j). \quad (54)$$

The behavior of the alignment indicator with respect to the rotation and translation errors is studied by Zhang et al. (2020).

As a connection with deep neural networks, OverlapNet (Chen et al., 2020) uses a neural network to predict a similar metric and detect loop closures. The cosine of the angle in (53) or the indicator in (54) provide such a metric for *self-supervised* learning while taking into account the semantic information. Given the promising results of the work of Chen et al. (2020), the combination of our metric with deep learning is an interesting future research direction. The expected outcomes will be a scalable diagnostic tool for monitoring the quality of point cloud alignments and representation learning for the model in (42).

A more direct approach will be the development of differentiable loss functions for deep learning tasks. Zhu et al. (2020) developed a new continuous 3D loss function for monocular single-view depth prediction based on our function representation. The proposed loss function addresses the gap between dense image prediction and sparse LIDAR supervision. By simply adding this new loss function to existing network architectures, the accuracy and geometric consistency of depth predictions are improved significantly on all three state-of-the-art baseline networks tested by Zhu et al. (2020).

11. Conclusion

We developed a new mathematical framework for sensor registration that enables nonparametric joint semantic/appearance and geometric representation of continuous functions using data. The continuous functions allow the registration to be independent of a specific signal resolution, and the framework is fully analytical with a closed-form derivation of the Riemannian gradient and Hessian. We studied a restriction of the framework where the Lie group acts on functions isometrically. This restriction has a variety of applications, and its high-dimensional derivation for the special Euclidean group acting on the Euclidean space

showcases the point cloud registration and bird’s-eye view map registration abilities. We derived low-dimensional cases with numerical examples to show the generality of the proposed framework. A specific implementation of this framework for RGB-D cameras performs well in texture and structure-scarce environments and is inherently robust to the input noise.

The general problem where the Lie group does not act isometrically on the functions has interesting applications such as image registration and nonrigid structure from motion. Although, the Lie algebra of the affine matrix group does not possess any particular structure, i.e., it includes all square matrices, due to its mentioned applications it is an attractive future research direction. The development of a globally optimal solver is, of course, another challenging and attractive future research direction.

The developed framework is a data-driven analytical model that possesses interesting intrinsic mathematical structures simultaneously, i.e., equivariance, and computational structures such as sparsity and parallel processing. Autonomy via computational intelligence is a multifaceted research domain which nicely integrates mathematics, computer science, and engineering and can have enormous impacts on our future and improve our quality of life. We developed a novel framework that can spark new research directions in robotics via sensor registration and mapping applications, computer vision via 3D scene reconstruction applications, medical imaging via nonrigid 3D reconstructions and registration, machine learning via Lie group machine learning, and in general artificial intelligence via an equivariant cognitive model.

Acknowledgments

A. Bloch and W. Clark were supported in part by NSF grant DMS-1613819 and AFOSR grant FA 0550-18-0028. A. Bloch was also supported in part by NSF grant DMS-2103026. Toyota Research Institute provided funds to support this work. Funding for M. Ghaffari was in part provided by NSF Award No. 2118818. The authors thank the anonymous reviewers for many constructive feedback and in-depth review of this work.

Appendix A. Required Mathematical Machinery and Notation

We review some mathematical preliminaries to establish the notation.

A.1 Matrix Lie Group of Motion in \mathbb{R}^n

The general linear group of degree n , denoted $\text{GL}_n(\mathbb{R})$, is the set of all $n \times n$ real invertible matrices, where the group binary operation is the ordinary matrix multiplication. The n -dimensional special orthogonal group, denoted

$$\text{SO}(n) = \{R \in \text{GL}_n(\mathbb{R}) \mid RR^T = I_n, \det R = +1\},$$

is the rotation group on \mathbb{R}^n . The n -dimensional special Euclidean group, denoted

$$\text{SE}(n) = \left\{ h = \begin{bmatrix} R & T \\ 0 & 1 \end{bmatrix} \in \text{GL}_{n+1}(\mathbb{R}) \mid R \in \text{SO}(n), T \in \mathbb{R}^n \right\},$$

is the group of rigid transformations, i.e., *direct isometries*, on \mathbb{R}^n . A transformation such as $h \in \text{SE}(3)$ is the parameter space in many sensor registration problems which consists of a rotation and a translation components. Let $\bar{h} \in \text{SE}(3)$ be an estimate of h . We compute the rotational and translational distances using $\|\log(\bar{R}R^\top)\|_F$ and $\|\bar{T} - \bar{R}R^\top T\|$, respectively, where $\log(\cdot)$ is the Lie logarithm map which is, here, the matrix logarithm. These definitions are consistent with the transformation distance that can be directly computed using $\|\log(\bar{h}h^{-1})\|_F$. Here, $\|A\|_F^2 = \text{Tr}(A^\top A)$ is the Frobenius norm.

This paper studies not only matrix Lie groups, but also their actions on manifolds. We have the following definition:

Definition 20 *Let \mathcal{G} be a group and X a set. A (left) group action of \mathcal{G} on X , denoted as $\mathcal{G} \curvearrowright X$, is a group homomorphism $\mathcal{G} \rightarrow \text{Aut}(X)$ (automorphism of X). If X is a smooth manifold, the action is smooth if $\mathcal{G} \rightarrow \text{Diff}(X)$ (diffeomorphism of X).*

Remark 21 *A group action can similarly be viewed as a function $\varphi : \mathcal{G} \times X \rightarrow X$ satisfying two conditions (where $\varphi(g, x)$ will be denoted by $g.x$)*

1. *Identity: If $e \in \mathcal{G}$ is the identity element, then $e.x = x$ for all $x \in X$.*
2. *Compatibility: $(gh).x = g.(h.x)$ for all $g, h \in \mathcal{G}$ and $x \in X$.*

Remark 22 *The standard action of $\text{SE}(n)$ on \mathbb{R}^n is given by $(R, T).x = Rx + T$ for $R \in \text{SO}(n)$ and $T \in \mathbb{R}^n$.*

A.2 Hilbert Space

Let V be a finite-dimensional vector space over the field of real numbers \mathbb{R} . An inner product on V is a function $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$ that is bilinear, symmetric, and positive definite. The pair $(V, \langle \cdot, \cdot \rangle)$ is called an inner product space. The inner product induces a norm that measures the magnitude or length of a vector: $\|v\| = \sqrt{\langle v, v \rangle}$. The norm in turn induces a metric that allows for calculating the distance between two vectors: $d(v, w) = \|v - w\| = \sqrt{\langle v - w, v - w \rangle}$. Such a metric is homogeneous; for $a \in \mathbb{R}$, $d(av, aw) = |a|d(v, w)$, and translation invariant; $d(v + x, w + x) = d(v, w)$. These properties are inherited from the induced norm. The distance metric is positive definite, symmetric, and satisfies the triangle inequality. In addition to measuring distances, it is important to be able to understand limits. This leads to the definition of a Cauchy sequence and completeness.

Definition 23 (Cauchy Sequence) *A Cauchy sequence is a sequence $\{x_i\}_{i=1}^\infty$ such that for any real number $\epsilon > 0$ there exists a natural number $\bar{n} \in \mathbb{N}$ such that for some distance metric $d(x_n, x_m) < \epsilon$ for all $n, m > \bar{n}$.*

Definition 24 (Completeness) *A metric space (M, d) is complete if every Cauchy sequence in M converges in M , i.e., to a limit that is in M .*

Such a metric space contains all its limit points. Note that completeness is with respect to the metric d and not the topology of the space. Now, we can give a definition for a Hilbert space.

Definition 25 (Hilbert Space) *A Hilbert space, \mathcal{H} , is a complete inner product space; that is, any Cauchy sequence in \mathcal{H} , using the metric induced by the inner product, converges to an element in \mathcal{H} .*

The typical example for a Hilbert space is the space of square-integrable functions on \mathbb{R} , i.e., $\mathcal{H} = L^2(\mathbb{R}, \mu)$ where μ is the Lebesgue measure on \mathbb{R} . The inner product is defined as:

$$\langle f, g \rangle_{\mathcal{H}} := \int f(x)g(x)d\mu(x). \quad (55)$$

Similarly, the induced norm by the inner product is $\|f\|_{\mathcal{H}} = \sqrt{\langle f, f \rangle_{\mathcal{H}}}$. The Hilbert space of functions can be thought of as an infinite-dimensional counterpart of the finite-dimensional vector spaces discussed earlier. However, the Hilbert spaces of interest in this work will be reproducing kernel Hilbert spaces discussed below.

A.3 Representation and Reproducing Kernel Hilbert Space

We now move to a more special type of Hilbert Space called Reproducing Kernel Hilbert Space (RKHS) (Berlinet and Thomas-Agnan, 2004) which we will use in this work.

Definition 26 (Kernel) *Let $x, x' \in \mathcal{X}$ be a pair of inputs for a function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ known as the kernel. A kernel is symmetric if $k(x, x') = k(x', x)$, and is positive definite if for any nonzero $f \in \mathcal{H}$ (or $L^2(\mathcal{X}, \mu)$):*

$$\int k(x, x')f(x)f(x')d\mu(x)d\mu(x') > 0.$$

Definition 27 (Reproducing Kernel Hilbert Space) *Let \mathcal{H} be a real-valued Hilbert space on a non-empty set \mathcal{X} . A function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a reproducing kernel of the Hilbert space \mathcal{H} iff:*

1. $\forall x \in \mathcal{X}, \quad k(\cdot, x) \in \mathcal{H}$,
2. $\forall x \in \mathcal{X}, \quad \forall f \in \mathcal{H} \quad \langle f, k(\cdot, x) \rangle = f(x)$.

The Hilbert space \mathcal{H} (RKHS) which possesses a reproducing kernel k is called a Reproducing Kernel Hilbert Space or a proper Hilbert space.

The second property is called *the reproducing property*; that is using the inner product of f with $k(\cdot, x)$, the value of function f is reproduced at point x . Also, using both conditions we have: $\forall x, z \in \mathcal{X}, k(x, z) = \langle k(\cdot, x), k(\cdot, z) \rangle$.

Lemma 28 *Any reproducing kernel is a positive definite function (Berlinet and Thomas-Agnan, 2004).*

Finding a reproducing kernel of an RKHS might seem difficult, but fortunately, there is a one-to-one relation between a reproducing kernel and its associated RKHS, and such a reproducing kernel is unique. Therefore, our problem reduces to finding an appropriate kernel.

Theorem 29 (Moore-Aronszajn Theorem, Berlinet and Thomas-Agnan, 2004) *Let k be a positive definite function on $\mathcal{X} \times \mathcal{X}$. There exists only one Hilbert space \mathcal{H} of functions on \mathcal{X} with k as reproducing kernel. The subspace \mathcal{H}_0 of \mathcal{H} spanned by the function $k(\cdot, x), x \in \mathcal{X}$ is dense¹ in \mathcal{H} and \mathcal{H} is the set of functions on \mathcal{X} which are point-wise limits of Cauchy sequence in \mathcal{H}_0 with the inner product*

$$\langle f, g \rangle_{\mathcal{H}_0} = \sum_{i=1}^n \sum_{j=1}^m \alpha_i \beta_j k(z_j, x_i), \quad (56)$$

where $f = \sum_{i=1}^n \alpha_i k(\cdot, x_i)$ and $g = \sum_{j=1}^m \beta_j k(\cdot, z_j)$.

The important property while working in an RKHS is that the convergence in norm implies point-wise convergence; the converse need not be true. In other words, if two functions in an RKHS are close in the norm sense, they are also close point-wise. We will rely on this property to solve the problem discussed in this paper. In Theorem 29, f and g are defined only in \mathcal{H}_0 . The following theorem known as the representer theorem ensures that the solution of minimizing the regularized risk functional admits such a representation.

Theorem 30 (Nonparametric Representer Theorem, Schölkopf et al., 2001) *Let \mathcal{X} be a nonempty set and \mathcal{H}_k be an RKHS with reproducing kernel k on $\mathcal{X} \times \mathcal{X}$. Suppose we are given a training sample $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{X} \times \mathbb{R}$, a strictly monotonically increasing real-valued function h on $[0, \infty)$, an arbitrary cost function $c : (\mathcal{X} \times \mathbb{R})^m \rightarrow \mathbb{R} \cup \{\infty\}$, and a class of functions²*

$$\mathcal{F} = \{f \in \mathbb{R}^{\mathcal{X}} \mid f(\cdot) = \sum_{i=1}^{\infty} \beta_i k(\cdot, z_i), \beta_i \in \mathbb{R}, z_i \in \mathcal{X}, \|f\|_{\mathcal{H}_k} < \infty\},$$

where $\|\cdot\|_{\mathcal{H}_k}$ is the induced norm of the RKHS \mathcal{H}_k associated with k . Then any $f \in \mathcal{F}$ minimizing the regularized risk functional

$$c((x_1, y_1, f(x_1)), \dots, (x_m, y_m, f(x_m))) + h(\|f\|_{\mathcal{H}_k})$$

admits a representation of the form

$$f(\cdot) = \sum_{i=1}^m \alpha_i k(\cdot, x_i). \quad (57)$$

A.4 Riemannian Geometry

Let M be a smooth manifold. This paper will be concerned with the problem of finding the maximum of a function $f : M \rightarrow \mathbb{R}$. This will be accomplished (locally) via gradient ascent. Due to the fact that M will not generally be Euclidean, we will be using tools from Riemannian geometry.

Recall that a vector field $X : M \rightarrow TM$ is a section of the tangent bundle, i.e. for $\pi : TM \rightarrow M$ we have $\pi \circ X = \text{Id}$. The set of all (smooth) vector fields on M will be

-
1. A dense subset of M implies the closure of the subset X equals M .
 2. $\mathbb{R}^{\mathcal{X}}$ is the space of functions mapping \mathcal{X} to \mathbb{R} .

denoted by $\mathfrak{X}(M)$. In contrast to vector fields, 1-forms are sections of the cotangent bundle, $M \rightarrow T^*M$ and the set of all 1-forms are denoted by $\Omega^1(M)$.

For a function $f : M \rightarrow \mathbb{R}$, the differential is naturally a covector rather than a vector: $df_x : T_x M \rightarrow \mathbb{R}$

$$df_x(v) = \left. \frac{d}{dt} \right|_{t=0} f(\gamma(t)), \quad \gamma(0) = x, \quad \dot{\gamma}(0) = v.$$

In order to determine the gradient, $\nabla f \in \mathfrak{X}(M)$, we need an identification $\Omega^1(M) \rightarrow \mathfrak{X}(M)$ which is where the Riemannian metric will be used.

Definition 31 *A Riemannian metric on a manifold M is the assignment to each point $p \in M$, of an inner product $\langle \cdot, \cdot \rangle_p$ on $T_p M$ which is smooth in the following sense: for $X, Y \in \mathfrak{X}(M)$, the map $p \mapsto \langle X(p), Y(p) \rangle_p$ is a smooth function on M .*

This gives a way to turn df into ∇f as follows:

$$df = \langle \nabla f, \cdot \rangle.$$

On a general manifold, if we wish to differentiate a vector field Y in the direction of X (i.e. a directional derivative), we cannot compute it in the traditional sense since different tangent vectors lie in different tangent spaces so we cannot take their differences. A way around this issue is to use a connection (see, e.g., §6 of Tu 2017):

Definition 32 *An affine connection on a manifold, M , is a \mathbb{R} -bilinear map $\nabla : \mathfrak{X}(M) \times \mathfrak{X}(M) \rightarrow \mathfrak{X}(M)$ such that for all $f \in C^\infty(M)$ and $X, Y \in \mathfrak{X}(M)$:*

- i. $\nabla_{fX} Y = f \nabla_X Y$,
- ii. $\nabla_X fY = \mathcal{L}_X f \cdot Y + f \nabla_X Y$,

where \mathcal{L} is the usual Lie derivative (Tu, 2011, Section 20).

For a Riemannian metric there exists a unique affine connection, called the Levi-Civita connection that is both torsion-free and compatible with the metric. A connection is compatible with the metric if for all $X, Y, Z \in \mathfrak{X}(M)$,

$$Z \langle X, Y \rangle = \langle \nabla_Z X, Y \rangle + \langle X, \nabla_Z Y \rangle,$$

and torsion-free if

$$\nabla_X Y - \nabla_Y X = [X, Y].$$

This connection can be constructed in the following way:

$$\begin{aligned} 2 \langle \nabla_X Y, Z \rangle &= X \langle Y, Z \rangle + Y \langle Z, X \rangle - Z \langle X, Y \rangle \\ &\quad - \langle X, [Y, Z] \rangle + \langle Y, [Z, X] \rangle + \langle Z, [X, Y] \rangle. \end{aligned} \tag{58}$$

Remark 33 *Suppose that \mathcal{G} is a commutative Lie group and we choose an orthogonal basis, $\{e_i\}$, of its Lie algebra \mathfrak{g} . If we let $\{E_i\}$ be the corresponding left-invariant vector fields, then the connection vanishes in these coordinates: $\nabla_{E_i} E_j = 0$ for all i, j . This will be useful in Sections 3.1 and 3.2 where we compute the registration problem for $\mathcal{G} = S^1$ and $\mathcal{G} = T^2$.*

The Levi-Civita connection is used to define a geodesic on a Riemannian manifold, i.e. a “straight line.”

Definition 34 A geodesic is a curve $\gamma : I \rightarrow M$ where I is an interval in \mathbb{R} and

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

where $\nabla_{\dot{\gamma}} \dot{\gamma}$ is the covariant derivative of the velocity vector field $\dot{\gamma}(t) = \frac{d}{dt} \gamma(t)$.

Theorem 35 Let $\gamma : I \rightarrow M$ be a curve. Then the following are equivalent:

1. γ is a geodesic,
2. In coordinates $\gamma = (y^1, \dots, y^n)$, the curve satisfies the following system of differential equations

$$\ddot{y}^k + \sum_{ij} \Gamma_{ij}^k \dot{y}^i \dot{y}^j = 0,$$

where Γ_{ij}^k are the Christoffel symbols (Tu, 2017, Sec. 13.3, p. 99), and

3. γ satisfies the Euler-Lagrange equations with Lagrangian $L(y, \dot{y}) = \frac{1}{2} \langle \dot{y}, \dot{y} \rangle_y$.

Example 8 Consider the case where $M = \mathbb{R}^2$ and the metric is given by $\langle u, v \rangle_p = u^\top m(p)v$ for a symmetric, positive-definite matrix m which depends on the point $p = (x, y) \in \mathbb{R}^2$. The Euler-Lagrange equations are

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{x}} - \frac{\partial L}{\partial x} = 0, \quad \frac{d}{dt} \frac{\partial L}{\partial \dot{y}} - \frac{\partial L}{\partial y} = 0.$$

Computing the first equation results with:

$$\begin{aligned} \frac{d}{dt} \frac{\partial L}{\partial \dot{x}} &= \frac{d}{dt} (m_{11} \dot{x} + m_{12} \dot{y}) = m_{11} \ddot{x} + m_{12} \ddot{y} + \frac{\partial m_{11}}{\partial x} \dot{x}^2 + \frac{\partial m_{12}}{\partial y} \dot{y}^2 + \left(\frac{\partial m_{11}}{\partial y} + \frac{\partial m_{12}}{\partial x} \right) \dot{x} \dot{y}, \\ \frac{\partial L}{\partial x} &= \frac{1}{2} \frac{\partial m_{11}}{\partial x} \dot{x}^2 + \frac{\partial m_{12}}{\partial x} \dot{x} \dot{y} + \frac{1}{2} \frac{\partial m_{22}}{\partial x} \dot{y}^2. \end{aligned}$$

Combining these (as well as performing the analogous computation for the y equation) we obtain the Euler-Lagrange equations.

$$\begin{aligned} m_{11} \ddot{x} + m_{12} \ddot{y} + \frac{1}{2} \frac{\partial m_{11}}{\partial x} \dot{x}^2 + \left(\frac{\partial m_{12}}{\partial y} - \frac{1}{2} \frac{\partial m_{22}}{\partial x} \right) \dot{y}^2 + \frac{\partial m_{11}}{\partial y} \dot{x} \dot{y} &= 0, \\ m_{12} \ddot{x} + m_{22} \ddot{y} + \left(\frac{\partial m_{12}}{\partial x} - \frac{1}{2} \frac{\partial m_{11}}{\partial y} \right) \dot{x}^2 + \frac{1}{2} \frac{\partial m_{22}}{\partial y} \dot{y}^2 + \frac{\partial m_{22}}{\partial x} \dot{x} \dot{y} &= 0. \end{aligned}$$

Decoupling the acceleration terms produces

$$\begin{aligned} \ddot{x} + \Gamma_{xx}^x \dot{x}^2 + \Gamma_{xy}^x \dot{x} \dot{y} + \Gamma_{yy}^x \dot{y}^2 &= 0, \\ \ddot{y} + \Gamma_{xx}^y \dot{x}^2 + \Gamma_{xy}^y \dot{x} \dot{y} + \Gamma_{yy}^y \dot{y}^2 &= 0, \end{aligned}$$

where

$$\begin{aligned}
 \Gamma_{xx}^x &= \frac{1}{\det(m)} \left(\frac{1}{2} m_{22} \frac{\partial m_{11}}{\partial x} - m_{12} \frac{\partial m_{12}}{\partial x} + \frac{1}{2} m_{12} \frac{\partial m_{11}}{\partial y} \right), \\
 \Gamma_{xy}^x &= \frac{1}{\det(m)} \left(m_{22} \frac{\partial m_{11}}{\partial y} - m_{12} \frac{\partial m_{22}}{\partial x} \right), \\
 \Gamma_{yy}^x &= \frac{1}{\det(m)} \left(m_{22} \frac{\partial m_{12}}{\partial y} - \frac{1}{2} m_{22} \frac{\partial m_{22}}{\partial x} - \frac{1}{2} m_{12} \frac{\partial m_{22}}{\partial y} \right), \\
 \Gamma_{xx}^y &= \frac{1}{\det(m)} \left(m_{11} \frac{\partial m_{12}}{\partial x} - \frac{1}{2} m_{12} \frac{\partial m_{11}}{\partial x} - \frac{1}{2} m_{11} \frac{\partial m_{11}}{\partial x} \right), \\
 \Gamma_{xy}^y &= \frac{1}{\det(m)} \left(m_{11} \frac{\partial m_{22}}{\partial x} - m_{12} \frac{\partial m_{11}}{\partial y} \right), \\
 \Gamma_{yy}^y &= \frac{1}{\det(m)} \left(\frac{1}{2} m_{12} \frac{\partial m_{22}}{\partial x} - m_{12} \frac{\partial m_{12}}{\partial y} + \frac{1}{2} m_{11} \frac{\partial m_{22}}{\partial y} \right).
 \end{aligned}$$

In particular, due to the existence and uniqueness of solutions to ordinary differential equations, geodesics always exist locally. Moreover, when the manifold is a Lie group and the metric is invariant under left-translations, the geodesic equations can be described via the Euler-Poincaré equations (Bloch et al., 1996).

Definition 36 Let \mathcal{G} be a Lie group. The function $\ell_g : \mathcal{G} \rightarrow \mathcal{G}$, $h \mapsto gh$ for $g \in \mathcal{G}$ is called left-translation by g . Its tangent lift, $(\ell_g)_* : T\mathcal{G} \rightarrow T\mathcal{G}$ is defined as follows: let $v \in T_h\mathcal{G}$ and $\gamma : (-\varepsilon : \varepsilon) \rightarrow \mathcal{G}$ such that $\gamma(0) = h$ and $\dot{\gamma}(0) = v$. Then, we have

$$(\ell_g)_* v := \left. \frac{d}{dt} \right|_{t=0} \ell_g \circ \gamma(t).$$

Definition 37 For a Lie group, \mathcal{G} , a function $L : T\mathcal{G} \rightarrow \mathbb{R}$ is said to be a left-invariant Lagrangian if $(\ell_g)^* L = L$ for all $g \in \mathcal{G}$. That is,

$$L(x, v) = L(gx, (\ell_g)_* v).$$

Theorem 38 Let \mathcal{G} be a Lie group and $L : T\mathcal{G} \rightarrow \mathbb{R}$ be a left-invariant Lagrangian. Let $\mathcal{L} : \mathfrak{g} \rightarrow \mathbb{R}$ be its restriction to the tangent space of \mathcal{G} at the identity (the corresponding Lie algebra). For a curve $g(t) \in \mathcal{G}$, let

$$\xi(t) = g(t)^{-1} \cdot \dot{g}(t); \quad \text{i.e.,} \quad \xi(t) = (\ell_{g^{-1}})_* \dot{g}(t).$$

Then the following are equivalent:

1. $g(t)$ satisfies the Euler-Lagrange equations for L on \mathcal{G} ,
2. The Euler-Poincaré equations hold:

$$\frac{d}{dt} \frac{\delta \mathcal{L}}{\delta \xi} = \text{ad}_\xi^* \frac{\delta \mathcal{L}}{\delta \xi}. \quad (59)$$

Proof See Theorem 5.2 of Bloch et al. (1996). ■

Remark 39 *The reasoning behind using $g^{-1}\dot{g}$ rather than just \dot{g} is because $\dot{g} \in T_g\mathcal{G}$ and $g^{-1} : T_g\mathcal{G} \rightarrow T_e\mathcal{G} = \mathfrak{g}$ and therefore $g^{-1}\dot{g} \in \mathfrak{g}$.*

In coordinates, let $\xi = \sum_i \xi^i e_i \in \mathfrak{g}$ where $\{e_i\}$ forms a basis. Then equation (59) takes the form

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \xi^j} = c_{ij}^k \frac{\partial \mathcal{L}}{\partial \xi^k} \xi^i, \quad [e_i, e_j] = \sum_k c_{ij}^k e_k. \quad (60)$$

The idea of a geodesic is useful because it provides a map $T_x M \rightarrow M$ in the following way:

$$v \mapsto \gamma(1), \quad \gamma(0) = x, \quad \dot{\gamma}(0) = v$$

where γ is a geodesic. This map is called the Riemannian exponential map and is denoted by $\text{Exp}_x : T_x M \rightarrow M$.

Remark 40 *There are two important aspects of the Riemannian exponential that need to be discussed. (1) In general, Exp_x is only defined on a neighborhood of $0 \in T_x M$ and not the whole tangent space. (2) When $M = \mathcal{G}$ is a Lie group, the Lie exponential map $\exp : \mathfrak{g} \rightarrow \mathcal{G}$ will usually not agree with the Riemannian exponential although they have matching domain and codomain.*

The last tool from Riemannian geometry we will use is the Hessian. This is a generalization of the second derivative. For a function $f : M \rightarrow \mathbb{R}$, the Hessian is a second-order tensor given by

$$H(f) = \sum_{ij} \left(\frac{\partial^2 f}{\partial x^i \partial x^j} - \sum_k \Gamma_{ij}^k \frac{\partial f}{\partial x^k} \right) dx^i \otimes dx^j.$$

A different, equivalent, definition of the Hessian is:

$$H(f)(X, Y) = \langle \nabla_X \text{grad}(f), Y \rangle = X(Yf) - df(\nabla_X Y).$$

This object can be used to calculate the covariance as well as be an ingredient in Newton's root finding method.

References

- Timothy D Barfoot. *State Estimation for Robotics*. Cambridge University Press, 2017.
- Bachir Bekka, Pierre de La Harpe, and Alain Valette. *Kazhdan's property (T)*. Cambridge university press, 2008.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- Youcef Bentoutou, Nasreddine Taleb, Kidiyo Kpalma, and Joseph Ronsin. An automatic image registration for applications in remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 43(9):2127–2137, 2005.

- Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Kluwer Academic, 2004.
- Rolf Berndt and Michael Klucznik. *An Introduction to Symplectic Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2001. ISBN 9780821820568.
- Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- Anthony Bloch, P. S. Krishnaprasad, Jerrold E. Marsden, and Tudor S. Ratiu. The Euler-Poincaré equations and double bracket dissipation. *Communications in Mathematical Physics*, 175(1):1–42, Jan 1996.
- Anthony M. Bloch, Roger W. Brockett, and Tudor S. Ratiu. Completely integrable gradient flows. *Communications in Mathematical Physics*, 147(1):57–74, Jun 1992.
- Silvère Bonnabel, Martin Barczyk, and François Goulette. On the covariance of icp-based scan-matching techniques. In *Proceedings of the American Control Conference*, pages 5498–5503. IEEE, 2016.
- Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Peter Deisenroth. Matérn gaussian processes on riemannian manifolds. *arXiv preprint arXiv:2006.10160*, 2020.
- Gary Bradski. The OpenCV Library. *Dr. Dobbs’s Journal of Software Tools*, 2000.
- Martin Brossard, Silvère Bonnabel, and Axel Barrau. A new approach to 3d icp covariance estimation. *IEEE Robotics and Automation Letters*, 5(2):744–751, 2020.
- John Canny. A computational approach to edge detection. In *Readings in Computer Vision*, pages 184–203. Elsevier, 1987.
- Andrea Censi. An accurate closed-form estimate of icp’s covariance. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3167–3172. IEEE, 2007.
- Andrea Censi. An ICP variant using a point-to-line metric. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 19–25. IEEE, 2008.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-supervised Learning*. Adaptive computation and machine learning. MIT Press, 2006. ISBN 9780262033589. URL <https://books.google.com/books?id=kfqvQgAACAAJ>.
- Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017.
- Xieyuanli Chen, Thomas Läbe, Andres Milioto, Timo Röhling, Olga Vysotska, Alexandre Haag, Jens Behley, Cyrill Stachniss, and FKIE Fraunhofer. OverlapNet: Loop closing for LiDAR-based SLAM. In *Proceedings of the Robotics: Science and Systems Conference*, 2020.

- Yang Chen and Gérard Medioni. Object modeling by registration of multiple range images. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2724–2729. IEEE, 1991.
- Hendrik Dahlkamp, Adrian Kaehler, David Stavens, Sebastian Thrun, and Gary R. Bradski. Self-supervised monocular road detection in desert terrain. In *Proceedings of the Robotics: Science and Systems Conference*, Philadelphia, USA, August 2006. doi: 10.15607/RSS.2006.II.005.
- Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2051–2060, 2017.
- Jakob Engel, Thomas Schöps, and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *European Conference on Computer Vision*, pages 834–849. Springer, 2014.
- Jakob Engel, Jörg Stückler, and Daniel Cremers. Large-scale direct SLAM with stereo cameras. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1935–1942. IEEE, 2015.
- Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2018.
- Aasa Feragen, Francois Lauze, and Soren Hauberg. Geodesic exponential kernels: When curvature and linearity conflict. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3032–3042, 2015.
- Friedrich Fraundorfer and Davide Scaramuzza. Visual odometry: Part ii: Matching, robustness, optimization, and applications. *IEEE Robotics & Automation Magazine*, 19(2): 78–90, 2012.
- Maani Ghaffari, William Clark, Anthony Bloch, Ryan M. Eustice, and Jessy W. Grizzle. Continuous direct sparse visual odometry from RGB-D images. In *Proceedings of the Robotics: Science and Systems Conference*, Freiburg, Germany, June 2019.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2014.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- Intel Corporation. Official Threading Building Blocks (TBB) GitHub repository. <https://github.com/intel/tbb>, 2019.
- Vladimir G. Ivancevic and Tijana T. Ivancevic. Lecture notes in Lie groups, 2011.

- Sadeep Jayasumana, Richard Hartley, Mathieu Salzmann, Hongdong Li, and Mehrtash Harandi. Kernel methods on riemannian manifolds with gaussian rbf kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(12):2464–2477, 2015.
- Christian Kerl. Dense Visual Odometry (dvo). <https://github.com/tum-vision/dvo>, 2013.
- Christian Kerl, Jürgen Sturm, and Daniel Cremers. Dense visual SLAM for RGB-D cameras. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2100–2106. IEEE, 2013a.
- Christian Kerl, Jürgen Sturm, and Daniel Cremers. Robust odometry estimation for RGB-D cameras. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3748–3754. IEEE, 2013b.
- Taejung Kim and Yong-Jo Im. Automatic satellite image registration by combination of matching and random sample consensus. *IEEE Transactions on Geoscience and Remote Sensing*, 41(5):1111–1117, 2003.
- David Landry, François Pomerleau, and Philippe Giguere. Cello-3d: Estimating the covariance of icp in the real world. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 8190–8196. IEEE, 2019.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- Tzu-Yuan Lin, William Clark, Ryan M Eustice, Jessy W Grizzle, Anthony Bloch, and Maani Ghaffari. Adaptive continuous visual odometry from rgb-d images. *arXiv preprint arXiv:1910.00713*, 2019.
- Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghahfoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- Edward M Mikhail, James S Bethel, and J Chris McGlone. *Introduction to modern photogrammetry*. Wiley, 2001.
- Kevin P Murphy. *Machine learning: a probabilistic perspective*. The MIT Press, 2012.
- Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer New York, 1996.
- David Nistér, Oleg Naroditsky, and James Bergen. Visual odometry for ground vehicle applications. *Journal of Field Robotics*, 23(1):3–20, 2006.

- Frank C Park and Roger W Brockett. Kinematic dexterity of robotic mechanisms. *International Journal of Robotics Research*, 13(1):1–15, 1994.
- Steven A Parkison, Lu Gan, Maani Ghaffari Jadidi, and Ryan M. Eustice. Semantic iterative closest point through expectation-maximization. In *Proceedings of the British Machine Vision Conference*, pages 1–17, Newcastle, UK, September 2018.
- Steven A Parkison, Maani Ghaffari, Lu Gan, Ray Zhang, Arash K Ushani, and Ryan M Eustice. Boosting shape registration algorithms via reproducing kernel Hilbert space regularizers. *IEEE Robotics and Automation Letters*, 4(4):4563–4570, 2019.
- Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- Matthieu Pizenberg. DVO core (without ROS dependency). <https://github.com/mpizenberg/dvo/tree/76f65f0c9b438675997f595471d39863901556a9>, 2019.
- Ramona-Andreea Rohan. Some remarks on the exponential map on the groups $SO(n)$ and $SE(n)$. In *Proceedings of the Fourteenth International Conference on Geometry, Integrability and Quantization*, pages 160–175, Sofia, Bulgaria, 2013. Avangard Prima.
- Antoni Rosinol, Marcus Abate, Yun Chang, and Luca Carlone. Kimera: an open-source library for real-time metric-semantic localization and mapping. *arXiv preprint arXiv:1910.02490*, 2019.
- Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry. *IEEE robotics & automation magazine*, 18(4):80–92, 2011.
- Bernhard Schölkopf, Ralf Herbrich, and Alex Smola. A generalized representer theorem. In *Computational learning theory*, pages 416–426. Springer, 2001.
- Aleksandr Segal, Dirk Haehnel, and Sebastian Thrun. Generalized-ICP. In *Proceedings of the Robotics: Science and Systems Conference*, Seattle, USA, June 2009. doi: 10.15607/RSS.2009.V.021.
- James Servos and Steven L Waslander. Multi-Channel Generalized-ICP: A robust framework for multi-channel scan registration. *Robotics and Autonomous Systems*, 87:247–257, 2017.
- Boris Sofman, Ellie Lin, J Andrew Bagnell, John Cole, Nicolas Vandapel, and Anthony Stentz. Improving robot navigation through self-supervised online learning. *Journal of Field Robotics*, 23(11-12):1059–1075, 2006.
- Todor Stoyanov, Martin Magnusson, Henrik Andreasson, and Achim J Lilienthal. Fast and accurate scan registration through minimization of the distance between compact 3D NDT representations. *International Journal of Robotics Research*, 31(12):1377–1393, 2012.

- Hauke Strasdat. *Local accuracy and global consistency for efficient visual SLAM*. PhD thesis, Imperial College London, 2012.
- Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of RGB-D SLAM systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 573–580. IEEE, 2012.
- Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
- Yanghai Tsin and Takeo Kanade. A correlation-based approach to robust point set registration. In *Proceedings of the European Conference on Computer Vision*, pages 558–569. Springer, 2004.
- Loring W. Tu. *An introduction to manifolds*. New York, US: Springer, 2011.
- Loring W. Tu. *Differential Geometry*. Graduate Texts in Mathematics. Springer, 2017. ISBN 9783319660824.
- Christopher K. Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 1. MIT press, 2006.
- Anestis Zaganidis, Li Sun, Tom Duckett, and Grzegorz Cielniak. Integrating deep semantic segmentation into 3-d point cloud registration. *IEEE Robotics and Automation Letters*, 3(4):2942–2949, October 2018.
- Ray Zhang, Tzu-Yuan Lin, Chien Erh Lin, Steven A Parkison, William Clark, Jessy W Grizzle, Ryan M Eustice, and Maani Ghaffari. A new framework for registration of semantic point clouds from stereo and RGB-D cameras. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2020.
- Minghan Zhu, Maani Ghaffari, Yuanxin Zhong, Pingping Lu, Zhong Cao, Ryan M Eustice, and Huei Peng. Monocular depth prediction through continuous 3D loss. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- Minghan Zhu, Maani Ghaffari, and Huei Peng. Correspondence-free point cloud registration with $SO(3)$ -equivariant implicit shape representations. In *Proceedings of the Conference on Robot Learning*, 2021. URL <https://openreview.net/forum?id=K0q9qDgn-Ta>.
- Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019.