

Continuous Direct Sparse Visual Odometry from RGB-D Images

Maani Ghaffari*, William Clark*, Anthony Bloch, Ryan M. Eustice, and Jessy W. Grizzle
University of Michigan, Ann Arbor, MI, USA

{maanigj, wiclark, abloch, eustice, grizzle}@umich.edu

*The authors contributed equally to this work.

Abstract—This paper reports on a novel formulation and evaluation of visual odometry from RGB-D images. Assuming a static scene, the developed theoretical framework generalizes the widely used direct energy formulation (photometric error minimization) technique for obtaining a rigid body transformation that aligns two overlapping RGB-D images to a continuous formulation. The continuity is achieved through functional treatment of the problem and representing the process models over RGB-D images in a reproducing kernel Hilbert space; consequently, the registration is not limited to the specific image resolution and the framework is fully analytical with a closed-form derivation of the gradient. We solve the problem by maximizing the inner product between two functions defined over RGB-D images, while the continuous action of the rigid body motion Lie group is captured through the integration of the flow in the corresponding Lie algebra. Energy-based approaches have been extremely successful and the developed framework in this paper shares many of their desired properties such as the parallel structure on both CPUs and GPUs, sparsity, semi-dense tracking, avoiding explicit data association which is computationally expensive, and possible extensions to the simultaneous localization and mapping frameworks. The evaluations on experimental data and comparison with the equivalent energy-based formulation of the problem confirm the effectiveness of the proposed technique, especially, when the lack of structure and texture in the environment is evident.

I. INTRODUCTION

Sensor registration is a fundamental task in robotic perception. Modern algorithms for sensor registration rely on optimization techniques to solve often nonlinear and large-scale problems [9, 23]. In particular, estimating the rigid body transformation of a moving sensor such as a camera (or in combination with other sensors) is of major importance in robotics and computer vision [8, 12, 15, 17, 30]. Cameras in both *monocular* and *depth* versions are rich sources of information acquisition and are well-suited for varieties of real-world applications of autonomous systems that usually involves solving the Simultaneous Localization and Mapping (SLAM) problem [10, 13, 21].

A key strength of the successful state-of-the-art visual perception algorithms is the ability to track the camera pose accurately with the lowest possible drift. Then, visual place-recognition algorithms such as the work of Gálvez-López and Tardos [16] are used to remove the drift by providing place-revisiting measurements known as *loop-closures*. Direct visual odometry methods minimize the *photometric error* using the photometric measurements, i.e., image intensity values, provided by the camera. As opposed to *indirect* methods,




Fig. 1: The registration of frames 1 and 2 of fr1/desk sequence from RGB-D SLAM dataset and benchmark [31] using the proposed continuous sensor registration framework. In the right, the figure shows the sparsity pattern of the corresponding kernel matrix of the inner product structure. The computation of this sparse matrix can be done entirely in parallel. Geometrically, the matrix shows the correlation of each measurement in frame 2 after applying the final rigid body transformation with other measurements in frame 1.

the information contained in an image is not abstracted into a set of sparse *keypoints* or *features*, e.g., corners; hence, direct methods can achieve superior performance in tracking and *dense* or *semi-dense* mappings, given a well-calibrated camera [3, 12, 22]. In addition, direct visual odometry *front-end systems* also possess a *sparse Hessian structure* which is very similar to the general SLAM structure, leading to real-time performance. In contrast, indirect methods model the geometric error and are much more robust to geometric noise such as a poor intrinsic camera calibration or a rolling shutter [13].

In this paper, inspired by the fact that (colored) point cloud measurements constitute discrete sensing information from the continuous motion of the camera and the idea of avoiding costly data association between measurements, we formulate the registration problem in a continuous form that directly operates on a mapping from 3D space to an abstract information space such as intensity surface. The explicit representation between color information and 2D/3D geometry (image or Euclidean space coordinates) is not directly available; hence, the current direct methods use numerical differentiation for computing the gradient and are limited to fixed image size and resolution, given the camera model and measurements for re-projection of the 3D points. As illustrated in Fig. 1, we show that the proposed continuous representation can be

sparse or sparsified which means, geometrically speaking, the local correlation is sufficiently rich to capture the information available in data. This work has the following contributions:

- 1) We develop a fundamentally novel formulation of the sensor registration problem that is continuous and models the action of an arbitrary Lie group on any smooth manifold. Our solution uses the integration of the flow in the Lie algebra by maximizing the inner product between two functions defined over the fixed and moving measurements sets. The continuity is achieved through functional treatment of the problem and representing the functions in a reproducing kernel Hilbert space.
- 2) We apply the developed framework to the particular, and commonly used, case of RGB-D images (depth camera) and $SE(n)$ matrix Lie group. As a result, the registration is not limited to the specific image resolution.
- 3) As opposed to the current direct energy formulation, which involves computation of the numerical image intensity gradient to be used in conjunction with the analytical Jacobian of the pose via the chain rule, our framework is fully analytical and the gradient has a complete closed-form derivation.
- 4) We evaluate the proposed algorithm using publicly available RGB-D benchmark in Sturm et al. [31] and provide the open-source implementation available at: <https://github.com/MaaniGhaffari/cvo-rgbd>

The remainder of this paper is organized as follows. In §II, the required mathematical preliminaries and notation are briefly explained. The main theoretical result for the general problem formulation on any smooth manifold and any Lie group acting on it is given in §III. The specialized form of the problem formulation for the special Euclidean group is derived in §IV. A brief theoretical analysis for the verification of the idea is provided in §V. The integration of the flow for $SE(n)$ and for the special case of 3D space, $SE(3)$, to obtain the solution is explained in §VI. Experimental evaluations of the proposed method for registration and tracking using RGB-D images are presented in §VII. Finally, §VIII concludes the paper and provides suggestions as future work.

II. MATHEMATICAL PRELIMINARIES AND NOTATION

We first review some mathematical preliminaries to establish the notation and give the paper better flow before stating the main results.

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

The general linear group of degree n , denoted $GL_n(\mathbb{R})$, 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

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

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

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

is the group of rigid transformations, i.e., *direct isometries*, on \mathbb{R}^n . A transformation such as $h \in SE(3)$ is the parameter space in many sensor registration problems which consists of a rotation and a translation components. Let $\bar{h} \in SE(3)$ be an estimate of h . We compute the rotational and translational distances using $\|\log(\bar{R}R^T)\|_F$ and $\|\bar{T} - \bar{R}R^T T\|$, 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, $\|A\|_F^2 = \text{Tr}(A^T A)$ is the Frobenius norm.

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

Definition 1. 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 1. 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 2. This paper will be primarily focused on the standard action of $SE(n)$ on \mathbb{R}^n given by $(R, T).x = Rx + T$ for $R \in SO(n)$ and $T \in \mathbb{R}^n$.

B. Hilbert Space

Let V be a 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 translational 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 2 (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 3 (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 4 (Hilbert Space). *A Hilbert space, \mathcal{H} , is a complete inner product space; that is any Cauchy sequence, using the metric induced by the inner product, converges to an element in \mathcal{H} .*

The reviewed definitions and properties are also valid for a vector space of functions. Let $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$ be a real Hilbert space of functions with the inner product between any two square-integrable functions $f, g \in \mathcal{H}$ (or $f, g \in L^2(\mathbb{R}, \mu)$) defined as:

$$\langle f, g \rangle_{\mathcal{H}} := \int f(x)g(x)d\mu(x), \quad (1)$$

where μ is the Lebesgue measure on \mathbb{R} . 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.

C. Representation and Reproducing Kernel Hilbert Space

We now move to a more special type called Reproducing Kernel Hilbert Space (RKHS) which we will use in this work.

Definition 5 (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 6 (Reproducing Kernel Hilbert Space [4]). *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 1. *Any reproducing kernel is a positive definite function [4].*

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 2 (Moore-Aronszajn Theorem [4]). *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

$$\langle f, g \rangle_{\mathcal{H}_0} = \sum_{i=1}^n \sum_{j=1}^m \alpha_i \beta_j k(z_j, x_i), \quad (2)$$

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 2, 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 3 (Nonparametric Representer Theorem [28]). *Let \mathcal{X} be a nonempty set and \mathcal{H} 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*

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

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

$$f(\cdot) = \sum_{i=1}^m \alpha_i k(\cdot, x_i). \quad (3)$$

III. 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 the two point clouds X and $hZ = \{hz_j\}$ the ‘‘best.’’ 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 3. *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 $h.f_Z := f_{h^{-1}Z}$.

¹A 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} .

Problem 1. 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

$$\arg \max_{h \in \mathcal{G}} F(h), \quad F(h) := \langle f_X, h.f_Z \rangle_{\mathcal{H}}. \quad (4)$$

A. 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

$$\begin{aligned} f_X(\cdot) &:= \sum_{x_i \in X} \ell_X(x_i) k(\cdot, x_i), \\ f_Z(\cdot) &:= \sum_{z_j \in Z} \ell_Z(z_j) k(\cdot, z_j). \end{aligned} \quad (5)$$

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

$$\langle f_X, f_Z \rangle_{\mathcal{H}} := \sum_{\substack{x_i \in X \\ z_j \in Z}} \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}} \cdot k(x_i, z_j). \quad (6)$$

Remark 4. We note two advantages of measuring the alignment of X and Z by (6). 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 !

B. Building the Gradient Flow

In order to (at least locally) solve (4), we will construct a gradient flow: $\dot{h} = \nabla F(h)$. Before we can do this, 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 [5]).

Definition 7. Suppose that a Lie Group \mathcal{G} acts diffeomorphically on a smooth manifold M via φ ; that is

$$\begin{aligned} \varphi : \mathcal{G} &\rightarrow \text{Diff}(M) \\ g &\mapsto \varphi_g. \end{aligned} \quad (7)$$

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:

$$df_x(\xi_M) := \left. \frac{d}{dt} \right|_{t=0} f(\varphi_{\exp(t\xi)}(x)), \quad f \in C^\infty(M). \quad (8)$$

This lets us compute the differential, dF .

Remark 5. To make notation more concise for the remainder of this paper, we will denote $c_{ij} := \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}}$.

Theorem 4. Suppose that $F(h) = \langle f_X, h.f_Z \rangle_{\mathcal{H}}$ as described above. Then

$$dF_e(\xi) = \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j}(-\xi_M(z_j)), \quad (9)$$

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

Remark 6. The notation for the differential of a function used throughout this paper is $df_x(v)$, where $x \in M$ and $v \in T_x M$:

$$df_x(v) = \left. \frac{d}{dt} \right|_{t=0} f(c(t)), \quad c(0) = x, \quad c'(0) = v. \quad (10)$$

Proof. This follows from a straight-forward application of the chain rule.

$$\begin{aligned} dF_e(\xi) &= \left. \frac{d}{dt} \right|_{t=0} \langle f_X, \exp(t\xi).f_Z \rangle_{\mathcal{H}} \\ &= \sum c_{ij} \cdot \left. \frac{d}{dt} \right|_{t=0} k(x_i, \exp(-t\xi)z_j) \\ &= \sum c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j} \cdot \left. \frac{d}{dt} \right|_{t=0} \exp(-t\xi)z_j \\ &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j}(-\xi_M(z_j)). \end{aligned} \quad (11)$$

Which matches equation (9). \square

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

$$\begin{aligned} dF_h((\ell_h)_* \xi) &= \\ &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{h^{-1}z_j}(-\xi_M(h^{-1}z_j)), \end{aligned} \quad (12)$$

where $c_{ij} = \langle \ell_X(x_i), \ell_Z(z_j) \rangle_{\mathcal{I}}$.

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

$$\langle \nabla F_h, (\ell_h)_* \xi \rangle_{T_h \mathcal{G}} = dF_h((\ell_h)_* \xi). \quad (13)$$

This allows for a way to obtain a (local) solution to (4) by following

$$\dot{h} = \nabla F(h). \quad (14)$$

IV. SPECIAL EUCLIDEAN GROUP

We will specialize the above treatment for the case where $\mathcal{G} = \text{SE}(n)$, the special Euclidean group in n dimensions. We will likewise let $M = \mathbb{R}^n$ on which $\text{SE}(n)$ acts in the usual way: 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 $k : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$:

$$k(x, y) = \sigma^2 \exp\left(\frac{-\|x - y\|_n^2}{2\ell^2}\right), \quad (15)$$

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 (14), 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 $\mathfrak{se}(n) = \text{Lie}(\text{SE}(n))$.

A. 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 [32]), we see that $\xi_M = \omega_M + v_M$. A straight forward computation leads to

$$\xi_{\mathbb{R}^n} x = \hat{\omega} x + v, \quad \hat{\omega} \in \text{Skew}(n), \quad v \in \mathbb{R}^n. \quad (16)$$

B. 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,

$$\langle (\omega, v), (\eta, u) \rangle_{\mathfrak{se}(n)} = b^2 \cdot \langle v, u \rangle_n - a^2 \left(\frac{n-2}{2} \right) \cdot \text{Tr}(\omega\eta), \quad (17)$$

where $\langle \cdot, \cdot \rangle_n$ is the standard Euclidean inner product on \mathbb{R}^n (see [25] for a discussion in three dimensions), as well as 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 only zeros except for a 1 in the (i, j) -coordinate.

C. Calculating the Gradient

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

$$\begin{aligned} dF_e(\xi) &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot d\left(\tilde{k}_{x_i}\right)_{z_j} (-\xi_M(z_j)) \\ &= \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (z_j - x_i), (-\hat{\omega} z_j - v) \rangle_n. \end{aligned} \quad (18)$$

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

$$\begin{aligned} &-a^2 \left(\frac{n-2}{2} \right) \cdot \text{Tr}[(\nabla_{\omega} F_e) \hat{\omega}] = \\ &\sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (z_j - x_i), (-\hat{\omega} z_j) \rangle_n, \\ &b^2 \langle (\nabla_v F_e), v \rangle_n = \\ &\sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot \frac{1}{\ell^2} k(x_i, z_j) \cdot \langle (z_j - x_i), (-v) \rangle_n. \end{aligned} \quad (19)$$

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:

$$\begin{aligned} (\nabla_{\omega} F_e)^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \langle (z_j - x_i), (-J^{pq} z_j) \rangle_n, \\ (\nabla_v F_e)^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot \langle (z_j - x_i), (-e^m) \rangle_n. \end{aligned} \quad (20)$$

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

$$\begin{aligned} (\nabla_{\omega} F_e)^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot (x_i^p z_j^q - x_i^q z_j^p), \\ (\nabla_v F_e)^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, z_j) \cdot (x_i^m - z_j^m). \end{aligned} \quad (21)$$

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 $(\hat{\omega}, v) \in \mathfrak{se}(n)$:

$$\begin{aligned} \hat{\omega}^{pq} &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, \tilde{z}_j) \cdot (x_i^p \tilde{z}_j^q - x_i^q \tilde{z}_j^p), \\ v^m &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot k(x_i, \tilde{z}_j) \cdot (x_i^m - \tilde{z}_j^m), \end{aligned} \quad (22)$$

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

V. ANALYSIS AND VERIFICATION OF IDEA

It is important to take a moment and examine when solving Problem 1 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 when the two clouds are identical, when does the identity element in the group maximize (4)?

Suppose that $Z = X$ and $\ell_Z = \ell_X$. Then the identity ideally should be a fixed-point of (14). This leads to the following:

Theorem 6. *Assume that for all $h \in \mathcal{G}$ and $x \in M$, $k(hx, hx) \leq k(x, x)$. Then the identity is a global maximum of F .*

Proof. We have that

$$F(h) = \langle f_X, h.f_X \rangle_{\mathcal{H}}, \quad F(e) = \|f_X\|_{\mathcal{H}}^2 \geq 0. \quad (23)$$

Then using the Cauchy-Schwarz inequality we obtain:

$$|\langle f_X, h.f_X \rangle_{\mathcal{H}}| \leq \|f_X\|_{\mathcal{H}} \cdot \|h.f_X\|_{\mathcal{H}}, \quad (24)$$

which is less than $F(e)$ provided that $\|h.f_X\|_{\mathcal{H}} \leq \|f_X\|_{\mathcal{H}}$.

Computing this, we see that

$$\begin{aligned} \|h \cdot f_X\|_{\mathcal{H}}^2 &= \sum_{x_i, z_j \in X} c_{ij} \cdot k(h^{-1}x_i, h^{-1}z_j) \\ &\leq \sum_{x_i, z_j \in X} c_{ij} \cdot k(x_i, z_j) = \|f_X\|_{\mathcal{H}}^2. \end{aligned} \quad (25)$$

Combining everything, we get that

$$|F(h)| \leq \|f_X\|_{\mathcal{H}} \cdot \|h \cdot f_X\|_{\mathcal{H}} \leq \|f_X\|_{\mathcal{H}}^2 = F(e). \quad (26)$$

□

Corollary 7. *Suppose $k : M \times M \rightarrow \mathbb{R}$ is a stationary kernel [26, 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 8. *The identity is a global maximum for the $SE(n)$ case.*

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

Theorem 9. *The maximizer of Problem 1, minimizes the angle between f_X and f_Z .*

Proof. Suppose $h^* \in \mathcal{G}$ is the maximizer of Problem 1. 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 $\cos(\theta) = |\langle f_X, f_Z \rangle| / (\|f_X\|_{\mathcal{H}} \|f_Z\|_{\mathcal{H}})$. □

VI. 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.

A. Integrating

We will take care that in integrating (14), 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 (14), $(\hat{\omega}, v)$ is constant in $\mathfrak{se}(n)$ (i.e. ∇F_h is a left-invariant vector field) the solution is merely

$$(R(t), T(t)) = (R_0, T_0) \exp(t(\hat{\omega}, v)), \quad (27)$$

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.

$$\begin{bmatrix} R_{k+1} & T_{k+1} \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} R_k & T_k \\ 0 & 1 \end{bmatrix} \cdot \exp \begin{bmatrix} t\hat{\omega} & tv \\ 0 & 0 \end{bmatrix} \quad (28)$$

$$= \begin{bmatrix} R_k & T_k \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta R & \Delta T \\ 0 & 1 \end{bmatrix} \quad (29)$$

$$= \begin{bmatrix} R_k \Delta R & R_k \Delta T + T_k \\ 0 & 1 \end{bmatrix} \quad (30)$$

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

$$R_{k+1} = R_k \Delta R \quad (31)$$

$$T_{k+1} = R_k \Delta T + T_k. \quad (32)$$

B. Step size

We can use (22) to point in the direction of maximal growth and (31) 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.

$$G(t) \approx \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} \cdot e^{\alpha_{ij}} \{g_{ij}^1 t + g_{ij}^2 t^2 + g_{ij}^3 t^3 + g_{ij}^4 t^4\}, \quad (33)$$

where

$$g_{ij}^1 = \beta_{ij}$$

$$g_{ij}^2 = \gamma_{ij} + \frac{1}{2} \beta_{ij}^2$$

$$g_{ij}^3 = \delta_{ij} + \beta_{ij} \gamma_{ij} + \frac{1}{6} \beta_{ij}^3$$

$$g_{ij}^4 = \varepsilon_{ij} + \beta_{ij} \delta_{ij} + \frac{1}{2} \beta_{ij}^2 \gamma_{ij} + \frac{1}{2} \gamma_{ij}^2 + \frac{1}{24} \beta_{ij}^4$$

$$\alpha_{ij} = \frac{-1}{2\ell^2} \|x_i - z_j\|_n$$

$$\beta_{ij} = \frac{-1}{\ell^2} \langle \hat{\omega} z_j + v, x_i - z_j \rangle_n$$

$$\gamma_{ij} = \frac{-1}{2\ell^2} (\|\hat{\omega} z_j + v\|_n^2 + 2\langle \hat{\omega}^2 z_j + v, x_i - z_j \rangle_n)$$

$$\delta_{ij} = \frac{1}{\ell^2} (\langle -\hat{\omega} z_j - v, \hat{\omega}^2 z_j + \hat{\omega} v \rangle_n + \langle -\hat{\omega}^3 z_j - \hat{\omega}^2 v, x_i - z_j \rangle_n)$$

$$\varepsilon_{ij} = \frac{-1}{2\ell^2} (\|\hat{\omega}^2 z_j + \hat{\omega} v\|_n^2 + 2\langle \hat{\omega} z_j + v, \hat{\omega}^3 z_j + \hat{\omega}^2 v \rangle_n + 2\langle \hat{\omega}^4 z_j + \hat{\omega}^3 v, x_i - z_j \rangle_n)$$

(34)

The “optimal” step size is then taken to be the value of $t > 0$ that maximizes the quartic (33).

C. Special Case: $n = 3$

When we restrict attention to $SE(3)$, the gradient (22) takes a special form:

$$\begin{aligned}\omega &= \frac{1}{a^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (x_i \times (h^{-1} z_j)) \\ v &= \frac{1}{b^2 \ell^2} \sum_{\substack{x_i \in X \\ z_j \in Z}} c_{ij} k(x_i, h^{-1} z_j) (x_i - h^{-1} z_j).\end{aligned}\quad (35)$$

Additionally, an explicit formula for the exponential map $\exp : \mathfrak{se}(3) \rightarrow SE(3)$ exists (see [20] and [27] for example). This gives an exact way to solve (31).

$$\begin{aligned}\Delta R &= I + \left(\frac{\sin t\theta}{\theta}\right) \hat{\omega} + \left(\frac{1 - \cos t\theta}{\theta^2}\right) \hat{\omega}^2, \\ \Delta T &= \left[tI + \left(\frac{1 - \cos t\theta}{\theta^2}\right) \hat{\omega} + \left(\frac{t\theta - \sin t\theta}{\theta^3}\right) \hat{\omega}^2 \right] v.\end{aligned}\quad (36)$$

where $\theta = \|\omega\|_3$ with $\omega \in \mathbb{R}^3$ and t is taken to maximize $G(t)$ in equation (33).

VII. EVALUATION AND DISCUSSION

First, we provide a numerical example to justify our choice of 4th-order Taylor expansion for the step length computation which in practice proven to be important. Next, we present the experimental evaluation using RGB-D data collected using the Microsoft Kinect sensor.

A. Taylor Expansion

This subsection explains (numerically) the reason for taking the extra complexity in computing the 4th-order approximation as opposed to a 2nd-order approximation. The following numerical example uses the point clouds supplied in MATLAB [1]. Here, X is the first picture and Z is the second. The parameter values used are $\sigma = 0.1$, $\ell = 0.15$, and $a^2 = b^2 = 7$.

With these chosen set of parameters, the computed gradient gives

$$\hat{\omega} = \begin{bmatrix} 0 & 3.6452 & -38.0800 \\ -3.6452 & 0 & -1.2370 \\ 38.0800 & 1.2370 & 0 \end{bmatrix}, \quad v = \begin{bmatrix} -9.3792 \\ 1.5252 \\ -8.9388 \end{bmatrix}.$$

Figure 2 contains the plots comparing the exact function, $G(t)$, along with its 2nd- and 4th-order Taylor approximation. As can be seen from the figure, although the 2nd-order approximation is a reasonable approximation for $t < 0.001$, it does not give a meaningful approximation for the maximum of $G(t)$. However, the 4th-order does.

B. TUM RGBD Benchmark

We performed experiments on four RGB-D image sequences of RGB-D SLAM dataset and benchmark [31]. This dataset 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 framerate. Also, the same set of parameters were used for all experiments




Fig. 2: Comparing the exact values of $G(t)$ with its Taylor approximations. It is clear that the 2nd-order approximation cannot capture the function behavior and, despite extra complexity, the approximation up to the 4th-order is necessary to maintain the integration of the flow smooth.

TABLE I: Parameters used for evaluation using TUM RGBD Benchmark, similar values were chosen for all experiments. 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.

Parameters	Value
transformation convergence threshold ϵ	1e-5
gradient norm convergence threshold ϵ	5e-5
kernel characteristic length-scale ℓ	0.15
kernel characteristic length-scale ℓ (iteration > 3)	0.10
kernel characteristic length-scale ℓ (iteration > 10)	0.06
kernel characteristic length-scale ℓ (iteration > 20)	0.03
kernel signal variance σ	0.1
minimum step length	0.2
color space inner product scale	10e-5
kernel sparsification threshold	1e-3

which are listed in Table I. To improve the computational efficiency, we adopted a similar approach to Engel et al. [12] by maintaining a semi-dense point cloud for each scan. In this work, we simply extracted edges on 2D images and exploited the one-to-one correspondences in the organized point cloud.

Figure 3 shows the cumulative distribution functions of position and orientation relative per frame errors. The figure also shows the top view of the accumulate trajectories in which the absolute accumulated error can be seen. The initial transformation always was set to the identity and the kernel evaluations were sparsified (an example is shown in Fig. 1) by setting any value less than 1e-3 to zero. The baseline for comparison is the open source implementation of the direct RGB-D visual odometry method in Steinbrücker et al. [29] which is available through OpenCV [2]. The proposed method performs better than or similar to the baseline, suggesting its potential for improving the visual front-end systems used for odometry or SLAM solutions.

Remark 7. Due to the fact that the images are in color, each point in the cloud has a color associated with it. This is encoded as a point in \mathbb{R}^3 . Therefore, we will take $\mathcal{I} = \mathbb{R}^3$ and its inner product to be

$$\langle \cdot, \cdot \rangle_{\mathcal{I}} = \frac{1}{10^5} \langle \cdot, \cdot \rangle_3. \quad (37)$$




Fig. 3: Experimental evaluation of the proposed Continuous Visual Odometry (CVO) framework using the RGB-D SLAM dataset and benchmark [31]. The baseline for comparison is the open source implementation of the Steinbrücker et al. [29] which is available through OpenCV [2]. The cumulative distribution functions of position and orientation errors shows performance of each method where the relative errors between two registered frames are computed using metrics defined earlier in II-A. The evaluation is done over the entire trajectories. For CVO, to reduce the number of points and improve the computational efficiency, a similar approach to Engel et al. [13] is adopted by only maintaining the high gradient points. In this work, we extracted edges on 2D images and exploited the one-to-one correspondences in the Microsoft Kinect organized point cloud. The proposed method almost always performs better than or similar to the baseline; however, it is particularly interesting that the difference is significant when the lack of both structure and texture in the environment is evident (first column from left).

VIII. CONCLUSION AND FUTURE WORK

We developed a continuous and direct formulation and solution for the RGB-D visual odometry problem. The proposed solution, given a calibrated camera, directly works on 3D data and possesses a sparse (or sparsified) structure as well as naturally parallelizable structures. Furthermore, due to the continuous representation, we neither require the association between two measurements sets nor the same number of measurements within each set. Theoretically speaking, given the promising results presented in this paper, the proposed method is an alternative to the core module of many modern visual odometry and tracking systems [13, 14, 24, 33].

Building a visual tracking system such as the work of Engel et al. [13] is an interesting future direction. Another interesting idea is guiding the gradient flow of $SE(3)$ using an Inertial Measurement Unit (IMU) [11, 15, 19] or combined contact-IMU process models [18, 19]. The propagation of the flow can be done between any camera frames as the visual data has a typically lower frequency.

A similar approach to the problem presented in this paper was studied in [6] and [7]. The authors show that it is possible to view the differential equation, $\dot{H} = [H, [H, N]]$, where H and N are elements of some Lie algebra, as a gradient flow on an adjoint orbit of the corresponding group *so long as*

the group is compact. This can then be used to solve linear programming problems where the data is encoded in $H(0)$ and N . Information about the solution to the linear programming problem can then be learned from the double bracket structure. The function which generates the gradient is

$$F(\theta) = \kappa(Q, \text{Ad}_\theta N), \quad \theta \in \mathcal{G}, \quad Q, N \in \mathfrak{g}, \quad (38)$$

where κ is the Killing form and Ad is the adjoint action. When the Lie group, \mathcal{G} , is compact, the Killing form is negative definite, i.e. $-\kappa$ is a nondegenerate, Ad -invariant inner product on \mathfrak{g} . Equation (38) is closely related to Equation (4); however the group studied here is $SE(n)$ which is noncompact (and furthermore, there do not exist any Ad -invariant inner products on $\mathfrak{se}(n)$). In future work we shall discuss how this framework can be used to solve Problem 1 when the group is restricted to be $SO(n)$.

ACKNOWLEDGMENTS

Funding for M. Ghaffari is given by the Toyota Research Institute (TRI), partly under award number N021515, however this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. Funding for J. Grizzle was in part provided by TRI and in part by NSF Award No. 1808051. Funding for W. Clark and A. Bloch was provided by NSF DMS 1613819 and AFSOR FA9550-18-10028.

REFERENCES

- [1] MATLAB 3-D point cloud registration and stitching. <https://www.mathworks.com/help/vision/examples/3-d-point-cloud-registration-and-stitching.html>. Accessed: 2018-10-18. 7
- [2] Open source computer vision; RgbOdometry class reference. https://docs.opencv.org/3.4/d0/d60/classcv_1_1RgbdOdometry.html. Accessed: 2019-01-01. 7, 8
- [3] Cédric Audras, A Comport, Maxime Meilland, and Patrick Rives. Real-time dense appearance-based SLAM for RGB-D sensors. In *Australasian Conference on Robotics and Automation*, 2011. 1
- [4] Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Kluwer Academic, 2004. 3
- [5] R. Berndt and M. Klucznik. *An Introduction to Symplectic Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2001. ISBN 9780821820568. 4
- [6] Anthony M. Bloch, Roger W. Brockett, and Tudor S. Ratiu. Completely integrable gradient flows. *Communications in Mathematical Physics*, 147(1):57–74, Jun 1992. 8
- [7] R.W. Brockett. Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems. *Linear Algebra and its Applications*, 146:79 – 91, 1991. ISSN 0024-3795. 8
- [8] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. MonoSLAM: Real-time single camera SLAM. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (6): 1052–1067, 2007. 1
- [9] Frank Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical report, Georgia Institute of Technology, 2012. 1
- [10] Frank Dellaert and Michael Kaess. Square root SAM: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research*, 25 (12):1181–1203, 2006. 1
- [11] Kevin Eickenhoff, Patrick Geneva, and Guoquan Huang. Closed-form preintegration methods for graph-based visual-inertial navigation. *The International Journal of Robotics Research*, 38(5):563–586, 2019. 8
- [12] 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. 1, 7
- [13] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2018. 1, 8
- [14] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. SVO: Fast semi-direct monocular visual odometry. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 15–22. IEEE, 2014. 8
- [15] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Transactions on Robotics*, 33(1):1–21, 2017. 1, 8
- [16] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, 2012. 1
- [17] Ross Hartley, Maani Ghaffari Jadidi, Lu Gan, Jiunn-Kai Huang, Jessy W. Grizzle, and Ryan M. Eustice. Hybrid contact preintegration for visual-inertial-contact state estimation within factor graphs. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3783–3790, Madrid, Spain, October 2018. 1
- [18] Ross Hartley, Maani Ghaffari Jadidi, Jessy W. Grizzle, and Ryan M. Eustice. Contact-aided invariant extended Kalman filtering for legged robot state estimation. In *Proceedings of the Robotics: Science and Systems Conference*, pages 1–9, Pittsburgh, PA, USA, June 2018. 8
- [19] Ross Hartley, Maani Ghaffari, Ryan M Eustice, and Jessy W Grizzle. Contact-aided invariant extended Kalman filtering for robot state estimation. *arXiv preprint arXiv:1904.09251*, 2019. 8
- [20] Vladimir G. Ivancevic and Tijana T. Ivancevic. Lecture notes in Lie groups, 2011. 7
- [21] Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John J Leonard, and Frank Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *The International Journal of Robotics Research*, 31(2):216–235, 2012. 1
- [22] 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, 2013. 1
- [23] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. g2o: A general framework for graph optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3607–3613. IEEE, 2011. 1
- [24] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2320–2327. IEEE, 2011. 8
- [25] Frank C Park and Roger W Brockett. Kinematic dexterity of robotic mechanisms. *International Journal of Robotics Research*, 13(1):1–15, 1994. 5
- [26] C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*, volume 1. MIT press, 2006. 6
- [27] 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. 7
- [28] Bernhard Schölkopf, Ralf Herbrich, and Alex Smola. A generalized representer theorem. In *Computational learning theory*, pages 416–426. Springer, 2001. 3
- [29] Frank Steinbrücker, Jürgen Sturm, and Daniel Cremers. Real-time visual odometry from dense RGB-D images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 719–722. IEEE, 2011. 7, 8
- [30] Hauke Strasdat. *Local accuracy and global consistency for efficient visual SLAM*. PhD thesis, Imperial College London, 2012. 1
- [31] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012. 1, 2, 7, 8
- [32] L. Tu. *Differential Geometry*. Graduate Texts in Mathematics. Springer, 2017. ISBN 9783319660824. 5
- [33] Rui Wang, Martin Schworer, and Daniel Cremers. Stereo DSO: Large-scale direct sparse visual odometry with stereo cameras. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3903–3911, 2017. 8