

Simultaneous Localization and Mapping: Through the Lens of Nonlinear Optimization

Amay Saxena*, *Student Member, IEEE*, Chih-Yuan Chiu*, *Student Member, IEEE*, Ritika Shrivastava, *Student Member, IEEE*, Joseph Menke, *Student Member, IEEE*, and Shankar Sastry, *Fellow, IEEE*

Abstract—Simultaneous Localization and Mapping (SLAM) algorithms perform visual-inertial estimation via filtering or batch optimization methods. Empirical evidence suggests that filtering algorithms are computationally faster, while optimization methods are more accurate. This work presents an optimization-based framework that unifies these approaches, and allows users to flexibly implement different design choices, e.g., the number and types of variables maintained in the algorithm at each time. We prove that filtering methods correspond to specific design choices in our generalized framework. We then reformulate the Multi-State Constrained Kalman Filter (MSCKF) and contrast its performance with that of sliding-window based filters. Our approach modularizes state-of-the-art SLAM algorithms to allow for adaptation to various scenarios. Experiments on the EuRoC MAV dataset verify that our implementations of these algorithms are competitive with the performance of off-the-shelf implementations in the literature. Using these results, we explain the relative performance characteristics of filtering and batch-optimization based algorithms in the context of our framework. We illustrate that under different design choices, our empirical performance interpolates between those of state-of-the-art approaches.

Index Terms—SLAM, estimation, control, computer vision

I. INTRODUCTION

In Simultaneous Localization and Mapping (SLAM), a robotic agent maps its uncharted environment while locating itself in the constructed map [1]. Applications include military map construction, search-and-rescue missions, augmented and virtual reality, and 3D scene capture [2–4].

Typical modern SLAM algorithms consist of *front* and *back ends*. The front end performs feature extraction, data association, and outlier rejection on raw sensor data. The back end then uses dynamics and measurement models for inference over the processed data, and produce compatible state estimates. Back end algorithms are often considered one of two classes—*Gaussian filtering* or *batch optimization* based. Filtering methods iteratively refine the distribution of recent states under a Gaussian prior [5–7], while optimization methods iteratively estimate states as solutions to an optimization problem, with objective constructed from inertial measurement unit (IMU) and image reprojection error terms. In particular, factor graph-based approaches efficiently solve optimization problems over past variables via factorization schemes that

maintain the sparsity of the underlying least squares problem [8–11]. Keyframe-based methods are optimization-based approaches that retain only a small subset of maximally informative frames (“keyframes”) in the optimization window arbitrarily spaced apart in time, while dropping all other poses [12]. Empirically, both classes of algorithms attain state-of-the-art performance, though the latter often attain higher accuracy at the cost of longer compute times [2, 12, 13].

Prior literature contrasted theoretical and empirical properties of filtering and batch optimization algorithms. Scaramuzza and Fraundorfer compared filtering and bundle adjustment-based methods for visual odometry [14, 15]. Frese et al. surveyed the use of grid-based and pose graph-based SLAM algorithms from a practitioner’s perspective [16]. Huang and Dissanyake conducted a theoretical study of the consistency, accuracy, and computational speed of filtering, optimization-based, and pose-graph SLAM [17]. Khosoussi et al. exploited sparsity in SLAM problems by conditioning on estimates of robot orientations [18]. Strasdat et al. conducted Monte Carlo experiments on visual SLAM algorithms [19], revealed that including more features in the back end increased accuracy more (compared to including more frames), and concluded that bundle adjustment outperforms filtering, since its computation time increases less drastically with the number of features.

In this work, we build upon prior literature by formulating a unified optimization-based framework for the SLAM back end that encompasses a large class of existing, state-of-the-art SLAM algorithms. We use this unified framework to recast the Extended Kalman Filter (EKF), Multi-State Constrained Kalman Filter (MSCKF), and Open Keyframe Visual-Inertial SLAM algorithm (OKVIS) as optimization-based back-end algorithms, and compare the empirical performance of the reformulated MSCKF with that of sliding-window optimization-based back-end algorithms, including the keyframe-based approach of Open Keyframe Visual-Inertial SLAM [12]. Somewhat surprisingly, the MSCKF outperforms sliding window filters of comparable sizes on several datasets, despite not performing multiple Gauss-Newton updates. We use our generalized framework to analyze these empirical findings.

II. SLAM: FORMULATION ON EUCLIDEAN SPACES

A. SLAM on Euclidean Spaces

SLAM estimates two types of variables: *states* and *features*. The state at each time t , denoted $x_t \in \mathbb{R}^{d_x}$, encodes information describing the robot, e.g., camera positions and orientations (poses). Feature positions available at time t in

Research supported by NSF Grant DMS 2013985 THEORINet: Transferable, Hierarchical, Expressive, Optimal, Robust and Interpretable Networks and U.S. Office of Naval Research MURI grant N00014-16-1-2710.

*Equal contribution.

All authors are with the Department of EECS at the University of California, Berkeley, CA 94720 USA (emails: {amaysaxena, chihyuan_chiu, joemenke, ritishri, sastry} at berkeley dot edu.)

a global frame, denoted $\{f_{t,j}|j = 1, \dots, p\} \subset \mathbb{R}^{d_f}$, can be obtained by analyzing information from image measurements $\{z_{t,j}|j = 1, \dots, p\} \subset \mathbb{R}^{d_z}$ and state estimates; these describe the relative position of the robot in its environment. States and features are described by an infinitely differentiable (i.e., C^∞) dynamics map $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ and a C^∞ measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$, via additive noise models:

$$x_{t+1} = g(x_t) + w_t, \quad w_t \sim \mathcal{N}(0, \Sigma_w), \quad (1)$$

$$z_{t,j} = h(x_t, f_{t,j}) + v_{t,j}, \quad v_{t,j} \sim \mathcal{N}(0, \Sigma_v), \quad (2)$$

where $\Sigma_w \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_w \succeq 0$ and $\Sigma_v \in \mathbb{R}^{d_z \times d_z}$, $\Sigma_v \succeq 0$.

For localization and mapping, SLAM algorithms maintain a *full state (vector)* $\bar{x}_t \in \mathbb{R}^d$, in which a number of past states and feature positions are concatenated. The exact number and time stamps of these states and features vary with the design choice of each SLAM algorithm. For example, sliding-window filters define the full state $\bar{x}_t := (x_{t-n+1}, \dots, x_t, f_{t,p-q+1}, \dots, f_{t,p}) \in \mathbb{R}^d$, with $d := d_x n + d_f q$, to be a sliding window of the most recent n states, consisting of one pose each, and the most recent estimates, at time t , of a collection of q features [5, 6]. Batch optimization methods, on the other hand, maintain all states and features encountered in the problem up to the current time [8–11].

Equations (1) and (2) do not involve overparameterized state variables, e.g., quaternion representations for poses, which are discussed in Section III.

B. SLAM as an Optimization Problem on Euclidean Spaces

SLAM estimates state and feature positions that best enforce constraints posed by given dynamics and measurement models, as well as noisy state and feature measurements collected over time. This is formulated as the minimization of the sum of weighted residual terms representing these constraints. For example, weighted residuals associated with the prior distribution over $\bar{x}_t \in \mathbb{R}^d$, the dynamics constraints between states $x_i, x_{i+1} \in \mathbb{R}^{d_x}$, and the reprojection error of feature $f_j \in \mathbb{R}^{d_f}$ corresponding to the state $x_i \in \mathbb{R}^{d_x}$ and image measurement $z_{t,j} \in \mathbb{R}^{d_z}$, may be given by $\Sigma_0^{-1/2}(\bar{x}_t - \mu_0) \in \mathbb{R}^d$, $\Sigma_w^{-1/2}(x_{i+1} - g(x_i)) \in \mathbb{R}^{d_x}$, and $\Sigma_v^{-1/2}(z_{i,j} - h(x_i, f_{t,j})) \in \mathbb{R}^{d_z}$, respectively (here, $1 \leq i \leq n-1$, $1 \leq j \leq q$). We define the running cost, $c : \mathbb{R}^{d_x n + d_f q} \rightarrow \mathbb{R}$, as the sum of weighted norm squares of these residuals. For example, for a sliding-window filtering algorithm for SLAM:

$$c(\bar{x}_t) := \|\bar{x}_t - \mu_0\|_{\Sigma_0}^2 + \sum_{i=t-n+1}^{t-1} \|x_{i+1} - g(x_i)\|_{\Sigma_w}^2 \quad (3)$$

$$+ \sum_{j=p-q+1}^p \sum_{i=t-n+1}^t \|z_{i,j} - h(x_i, f_{t,j})\|_{\Sigma_v}^2,$$

where $\|v\|_A^2 := v^\top A v$ for any real vector v and real matrix A of compatible dimension.

To formulate SLAM as a nonlinear least-squares problem, we stack all residual terms into one residual vector $C(\bar{x}_t)$. For example, for the sliding-window filter given above:

$$C(\bar{x}_t) := \left[(\Sigma_0^{-1/2}(\bar{x}_t - \mu_0))^\top \right.$$

$$\left. (\Sigma_w^{-1/2}(x_{t-n+1} - g(x_{t-n})))^\top \cdots (\Sigma_w^{-1/2}(x_t - g(x_{t-1})))^\top \right.$$

$$\left. (\Sigma_v^{-1/2}(z_{t-n+1,p-q+1} - h(x_{t-n+1}, f_{t,p-q+1})))^\top \cdots \right.$$

$$\left. (\Sigma_v^{-1/2}(z_{t-n+1,p} - h(x_{t-n+1}, f_{t,p})))^\top \cdots \right.$$

$$\left. (\Sigma_v^{-1/2}(z_{t,p-q+1} - h(x_t, f_{t,p-q+1})))^\top \cdots \right.$$

$$\left. (\Sigma_v^{-1/2}(z_{t,p} - h(x_t, f_{t,p})))^\top \right]^\top \in \mathbb{R}^{(2n-1)d_x + nq d_z}.$$

Thus, $c(\bar{x}_t) = C(\bar{x}_t)^\top C(\bar{x}_t)$, and the SLAM problem is now reduced to the nonlinear least squares problem below:

$$\min_{\bar{x}_t} c(\bar{x}_t) = \min_{\bar{x}_t} C(\bar{x}_t)^\top C(\bar{x}_t) \quad (4)$$

Section IV introduces the main algorithmic submodules used to find an approximate solution to (4).

III. SLAM: FORMULATION ON MANIFOLDS

Here, we generalize the SLAM formulation in Section II to the case where dynamical states are defined on smooth manifolds rather than Euclidean spaces. SLAM often involves the orientations of rigid bodies, which evolve on a smooth manifold embedded in an ambient space, e.g., rotation matrices expressed as unit quaternions. In such situations, we use *boxplus* (\boxplus) and *boxminus* (\boxminus) operators, defined below, to perform composition and difference operations in the iterative algorithm presented in Section IV, while enforcing constraints imposed by the manifold's geometric structure.

Suppose the full state x evolves on a smooth manifold \mathcal{M} , with $\dim(\mathcal{M}) = n$. For each $x \in \mathcal{M}$, let $\pi_x : U_x \rightarrow V_x$ be a diffeomorphic chart from an open neighborhood $U_x \subset \mathcal{M}$ of $x \in \mathcal{M}$ to an open neighborhood $V_x \subset \mathbb{R}^n$ of $0 \in \mathbb{R}^n$. Without loss of generality, suppose $\pi_x(x) = 0$. The operators $\boxplus : U_x \times V_x \rightarrow U_x$ and $\boxminus : U_x \times U_x \rightarrow V_x$ are defined by:

$$x \boxplus \delta = \pi_x^{-1}(\delta) \quad (5)$$

$$y \boxminus x = \pi_x(y) \quad (6)$$

In essence, \boxplus adds a perturbation $\delta \in \mathbb{R}^n$, in local coordinates, to a state $x \in \mathcal{M}$, while \boxminus extracts the difference $\delta \in \mathbb{R}^n$, in local coordinates, between states $x, x' \in \mathcal{M}$ covered by the same chart. Below, “ δ ” often describes an error or increment to a nominal state on the manifold.

A. Manifold Examples

This subsection gives examples of the \boxplus , \boxminus and π operators for manifolds that occur widely in SLAM: the set of unit quaternions, \mathbb{H}_u , and the set of rotation matrices, $SO(3)$.

Each $q \in \mathbb{H}_u$ is expressed as $q = (q_u, \vec{q}_v)$ where $q_u \in \mathbb{R}$ and $\vec{q}_v \in \mathbb{R}^3$ denote the scalar and vector (imaginary) parts, respectively, with $\|q\| = \sqrt{q_u^2 + \|\vec{q}_v\|^2} = 1$ (JPL convention). Here, the coordinate map $\pi : \mathbb{H}_u \rightarrow \mathbb{R}^3$ is defined as the Log map on \mathbb{H}_u ; its inverse π^{-1} is the Exp map. Specifically, we write each $q \in \mathbb{H}_u$ as $q = (\cos(\frac{\theta}{2}), \sin(\frac{\theta}{2})\vec{\omega})$ for some $\theta \in [0, \pi]$, $\vec{\omega} \in \mathbb{R}^3$ with $\|\vec{\omega}\| = 1$, i.e., the quaternion q implements a rotation about the axis $\vec{\omega}$ by θ radians counterclockwise. Then, $\pi : \mathbb{H}_u \rightarrow \mathbb{R}^3$ and $\pi^{-1} : B_\pi(0) \rightarrow \mathbb{H}_u$ are defined by: $(B_\pi(0) := \{x \in \mathbb{R}^3 : \|x\|_2 < \pi\})$ denotes the image of π

$$\pi(q) = \text{Log}(q) = \theta \vec{\omega},$$

$$\pi^{-1}(\theta\vec{\omega}) = \text{Exp}(\theta\vec{\omega}) = (\cos(\theta/2), \sin(\theta/2)\vec{\omega}).$$

The \boxplus and \boxminus maps are then implemented via the standard quaternion product $\star: \mathbb{H}_u \times \mathbb{H}_u \rightarrow \mathbb{H}_u$:

$$\begin{aligned} q_a \boxplus \vec{\omega} &= q_a \star \text{Exp}(\vec{\omega}) \\ q_a \boxminus q_b &= \text{Log}(q_b^{-1} \star q_a) \end{aligned}$$

For $SO(3)$, we define \boxplus and \boxminus similarly, i.e.,

$$\begin{aligned} R_a \boxplus \vec{\omega} &= R_a \text{Exp}(\vec{\omega}) \\ R_a \boxminus R_b &= \text{Log}(R_b^T R_a) \end{aligned}$$

Often, the full state in a SLAM problem exists in the Cartesian product of a finite collection of manifolds, since it contains poses and features on their own manifolds. For a product manifold $\mathcal{M}_1 \times \mathcal{M}_2$, with projection, increment, and difference maps already defined on \mathcal{M}_1 and \mathcal{M}_2 , we define \boxplus and \boxminus on $\mathcal{M}_1 \times \mathcal{M}_2$ by:

$$\begin{aligned} (g_1, g_2) \boxplus (\xi_1, \xi_2) &= (g_1 \boxplus \xi_1, g_2 \boxplus \xi_2) \\ (g_1, g_2) \boxminus (h_1, h_2) &= (g_1 \boxminus h_1, g_2 \boxminus h_2) \end{aligned}$$

B. SLAM as an Optimization Problem on Manifolds

The SLAM problem can be formulated on manifolds using modified cost functions, where plus and minus operations are replaced with \boxplus and \boxminus when necessary. (Appendix A1).

IV. MAIN ALGORITHM

A. Algorithm Overview

This section details submodules for a general SLAM algorithm, using state variables and cost terms defined in Sections II and III. We first introduce a formulation on Euclidean spaces. Below, denote the state and concatenated cost vector by $\bar{x}_t \in \mathbb{R}^d$ and $C: \mathbb{R}^d \rightarrow \mathbb{R}^{d_C}$, respectively. (e.g., the sliding window filter in Section II would correspond to $d = d_x n + d_f q$ and $d_C = (2n-1)d_x + nq d_z$). The SLAM problem is then equivalent to solving the nonlinear least-squares problem (4), reproduced below:

$$\min_{\bar{x}_t} .c(\bar{x}_t) = \min_{\bar{x}_t} .\|C(\bar{x}_t)\|_2^2.$$

B. Gauss-Newton Descent

Gauss-Newton descent minimizes $c(\bar{x}_t)$ via Gauss-Newton steps (Alg. 3), by iteratively approximating $c(\bar{x}_t)$ about a given linearization point \bar{x}_t^* as a linear least-squares cost term, i.e.,

$$\min_{\bar{x}_t} .c(\bar{x}_t) = \min_{\bar{x}_t} .\|\bar{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + o(\bar{x}_t - \bar{x}_t^*) \quad (7)$$

for some $\mu_t \in \mathbb{R}^d$ and $\Sigma_t \in \mathbb{R}^{d \times d}$. The theorem below describes the linearization procedure required to obtain $\mu_t \in \mathbb{R}^d$ and $\Sigma_t \in \mathbb{R}^{d \times d}$, and the approximation involved.

Theorem 4.1: (Gauss-Newton Step) Let $\bar{x}_t^* \in \mathbb{R}^d$ be a given linearization point, and suppose $J := \frac{\partial C}{\partial \bar{x}_t} \in \mathbb{R}^{d_C \times d}$ has full column rank. Applying a Gauss-Newton step (Alg. 3) to the cost $c(\bar{x}_t)$, about $\bar{x}_t^* \in \mathbb{R}^d$ yields the new cost:

$$c(\bar{x}_t) = \|\bar{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + o(\bar{x}_t - \bar{x}_t^*),$$

where $\mu_t \in \mathbb{R}^d$ and $\Sigma_t \in \mathbb{R}^{d \times d}$ are given by:

$$\begin{aligned} \Sigma_t &\leftarrow (J^\top J)^{-1}, \\ \mu_t &\leftarrow \bar{x}_t^* - (J^\top J)^{-1} J^\top C(\bar{x}_t^*). \end{aligned}$$

Proof: See Appendix (Section B2). ■

C. Marginalization of States

The marginalization step (Alg. 4) reduces the size of the SLAM problem by removing states that are no longer relevant, thus improving computational efficiency. First, we partition the overall state $\bar{x}_t \in \mathbb{R}^d = \mathbb{R}^{d_M + d_K}$ into a *marginalized component* $\bar{x}_{t,M} \in \mathbb{R}^{d_M}$, to be discarded from \bar{x}_t , and a *non-marginalized component* $\bar{x}_{t,K} \in \mathbb{R}^{d_K}$, to be kept. Then, we partition $c(\bar{x}_t)$ into two cost terms: $c_1(\bar{x}_{t,K})$, which depends only on non-marginalized state components, and $c_2(\bar{x}_{t,K}, \bar{x}_{t,M})$ which depends on both marginalized and non-marginalized state components:

$$\begin{aligned} c(\bar{x}_t) &= c(\bar{x}_K, \bar{x}_M) = c_1(\bar{x}_K) + c_2(\bar{x}_K, \bar{x}_M) \\ &= \|C_1(\bar{x}_K)\|_2^2 + \|C_2(\bar{x}_K, \bar{x}_M)\|_2^2. \end{aligned}$$

Here, $C_1(\bar{x}_K) \in \mathbb{R}^{d_{C,1}}$ and $C_2(\bar{x}_K, \bar{x}_M) \in \mathbb{R}^{d_{C,2}}$ denote the concatenation of residuals associated with $c_1(\bar{x}_K)$ and $c_2(\bar{x}_K, \bar{x}_M)$ (with $d_C = d_{C,1} + d_{C,2}$). To remove $\bar{x}_{t,M} \in \mathbb{R}^{d_M}$ from the optimization problem, observe that:

$$\begin{aligned} \min_{\bar{x}_t} c(\bar{x}_t) &= \min_{\bar{x}_{t,K}, \bar{x}_{t,M}} \left(c_1(\bar{x}_{t,K}) + c_2(\bar{x}_{t,K}, \bar{x}_{t,M}) \right) \\ &= \min_{\bar{x}_{t,K}} \left(\|C_1(\bar{x}_{t,K})\|_2^2 + \min_{\bar{x}_{t,M}} \|C_2(\bar{x}_{t,K}, \bar{x}_{t,M})\|_2^2 \right). \end{aligned}$$

To remove $\bar{x}_{t,M}$, we approximate the solution to the inner minimization problem by a linear least-squares cost, i.e.:

$$\min_{\bar{x}_{t,M}} \|C_2(\bar{x}_{t,K}, \bar{x}_{t,M})\|_2^2 \approx \|\bar{x}_{t,K} - \bar{\mu}_{t,K}\|_{\bar{\Sigma}_{t,K}^{-1}}^2$$

for some $\bar{\mu}_{t,K} \in \mathbb{R}^{d_K}$ and $\bar{\Sigma}_{t,K} \in \mathbb{R}^{d_K \times d_K}$. Since $\|C_2(\bar{x}_{t,K}, \bar{x}_{t,M})\|_2^2$ is in general non-convex, we obtain $\bar{\mu}_{t,K}$ and $\bar{\Sigma}_{t,K}$ by minimizing the first-order Taylor expansion of $\|C_2(\bar{x}_{t,K}, \bar{x}_{t,M})\|_2^2$ about some linearization point. Below, Theorem 4.2 details the derivation of $\bar{\mu}_{t,K}$ and $\bar{\Sigma}_{t,K}$. (For the proof, see Appendix B3).

Theorem 4.2 (Marginalization Step): Let $\bar{x}_t^* \in \mathbb{R}^d$ be a given linearization point, and suppose $J := \frac{\partial C}{\partial \bar{x}_t} \in \mathbb{R}^{d_C \times d}$ has full column rank. Define $J_K := \frac{\partial C}{\partial \bar{x}_{t,K}} \in \mathbb{R}^{d_C \times d_K}$, $J_M := \frac{\partial C}{\partial \bar{x}_{t,M}} \in \mathbb{R}^{d_C \times d_M}$. If $C(\bar{x}_{t,M}, \bar{x}_{t,K})$ were a linear function of $\bar{x}_t = (\bar{x}_{t,M}, \bar{x}_{t,K})$, then applying a Marginalization step (Alg. 4) to the cost $c(\bar{x}_t)$, about the linearization point $\bar{x}_t^* = (\bar{x}_{t,K}^*, \bar{x}_{t,M}^*) \in \mathbb{R}^d$, yields:

$$\min_{\bar{x}_t} c(\bar{x}_t) = \min_{\bar{x}_{t,K}} \left(c_1(\bar{x}_{t,K}) + \min_{\bar{x}_{t,M}} c_2(\bar{x}_{t,K}, \bar{x}_{t,M}) \right), \quad (8)$$

where $\Sigma_{t,K} \in \mathbb{R}^{d_K \times d_K}$ and $\mu_{t,K} \in \mathbb{R}^{d_K}$ are given by:

$$\begin{aligned} \Sigma_{t,K} &:= \left(J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K \right)^{-1}, \\ \mu_{t,K} &:= \bar{x}_{t,K}^* - \Sigma_{t,K} J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}_t^*). \end{aligned} \quad (9)$$

(10)

D. Main Algorithm on Manifolds

The Euclidean-space framework above can be directly extended to a formulation on manifolds, by using concepts in Section III to modify the dynamics and measurement maps in Section II, as well as the cost functions, Gauss-Newton steps, and marginalization steps in Sections IV-B, IV-C. When appropriate, plus and minus operations must be replaced with \boxplus and \boxminus (Appendix A2).

V. EQUIVALENCE OF FILTERING AND OPTIMIZATION APPROACHES

Here, we demonstrate the equivalence of filtering and batch optimization-based SLAM algorithms, using the Extended Kalman Filter (EKF, in Section V-A) and Multi-State Constrained Kalman Filter (MSCKF, in Section V-B), as examples. Although similar results exist in the optimization literature, they do not analyze algorithmic submodules unique to SLAM, e.g., feature incorporation and discarding [20]. For an introduction to the classical formulations of EKF and MSCKF SLAM, please see Appendices C2, C4.

A. Extended Kalman Filter (EKF), on Euclidean Spaces

At each time t , the EKF SLAM algorithm on Euclidean spaces maintains the full state vector $\tilde{x}_t := (x_t, f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{d_x+pd_f}$, consisting of the most recent state $x_t \in \mathbb{R}^{d_x}$ and feature position estimates $f_{t,1}, \dots, f_{t,p} \in \mathbb{R}^{d_f}$. At initialization ($t = 0$), no feature has been detected ($p = 0$), and the EKF full state is simply the initial state $\tilde{x}_0 = x_0 \in \mathbb{R}^{d_x}$, with mean $\mu_0 \in \mathbb{R}^{d_x}$ and covariance $\Sigma_0 \in \mathbb{R}^{d_x \times d_x}$. Suppose, at the current time t , the running cost $c_{EKF,t,0} : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}^{d_x+pd_f}$ is:

$$c_{EKF,t,0} = \|\tilde{x}_t - \mu_t\|_{\Sigma_t}^2,$$

where $\tilde{x}_t := (x_t, f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{d_x+pd_f}$ denotes the EKF full state at time t , with mean $\mu_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$. First, the *feature augmentation step* appends position estimates of new features $f_{p+1}, \dots, f_{p+p'} \in \mathbb{R}^{d_f}$ to the EKF full state \tilde{x}_t , and updates its mean and covariance. In particular, feature measurements $z_{t,p+1}, \dots, z_{t,p+p'} \in \mathbb{R}^{d_z}$ are incorporated by adding measurement residual terms to the current running cost $c_{EKF,t,0}$, creating a new cost $c_{EKF,t,1} : \mathbb{R}^{d_x+(p+p')d_f} \rightarrow \mathbb{R}$:

$$\begin{aligned} c_{EKF,t,1}(\tilde{x}_t, f_{t,p+1}, \dots, f_{t,p+p'}) \\ := \|\tilde{x}_t - \mu_t\|_{\Sigma_t}^2 + \sum_{k=p+1}^{p+p'} \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v}^2. \end{aligned}$$

In effect, $c_{EKF,t,1}$ appends positions of new features to \tilde{x}_t , and constrains it using feature measurements residuals.

Next, the *feature update* step uses measurements of features contained in \tilde{x}_t to update the mean and covariance of \tilde{x}_t . More precisely, feature measurements $z_{t,1:p} := (z_{t,1}, \dots, z_{t,p}) \in \mathbb{R}^{pd_z}$, of the p features f_1, \dots, f_p included in \tilde{x}_t , are introduced by incorporating associated measurement residuals

to the running cost $c_{EKF,t,0}$, creating a new cost $c_{EKF,t,2} : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}$:

$$c_{EKF,t,2}(\tilde{x}_t) := \|\tilde{x}_t - \mu_t\|_{\Sigma_t}^2 + \sum_{k=1}^p \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v}^2.$$

A Gauss-Newton step then constructs an updated mean $\bar{\mu}_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\bar{\Sigma}_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$ for \tilde{x}_t , creating a new cost $c_{EKF,t,3} : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}$:

$$c_{EKF,t,3}(\tilde{x}_t) := \|\tilde{x}_t - \bar{\mu}_t\|_{\bar{\Sigma}_t}^2,$$

which returns the running cost to the form of $c_{EKF,t,0}$.

The *state propagation* step propagates the EKF full state forward by one time step, via the EKF state propagation map $g : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}^{d_x+pd_f}$. To propagate \tilde{x}_t forward in time, we incorporate the dynamics residual to the running cost $c_{EKF,t,0}$ to create a new cost $c_{EKF,t,4} : \mathbb{R}^{2d_x+pd_f} \rightarrow \mathbb{R}$:

$$c_{EKF,t,4}(\tilde{x}_t, x_{t+1}) := \|\tilde{x}_t - \bar{\mu}_t\|_{\bar{\Sigma}_t}^2 + \|x_{t+1} - g(x_t)\|_{\Sigma_w}^2.$$

In effect, $c_{EKF,t,4}$ appends the new state $x_{t+1} \in \mathbb{R}^{d_x}$ to \tilde{x}_t , while adding a new constraint posed by the dynamics residuals. A marginalization step, with $\tilde{x}_{t,K} := (x_{t+1}, f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{d_x+pd_f}$ and $\tilde{x}_{t,M} := x_t \in \mathbb{R}^{d_x}$, then removes the previous state $x_t \in \mathbb{R}^{d_x}$ from the running cost. This step produces a mean $\mu_{t+1} \in \mathbb{R}^{d_x+pd_f}$ and a covariance $\Sigma_{t+1} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$ for the new EKF full state, $\tilde{x}_{t+1} := \tilde{x}_{t,K}$. The running cost is updated to $c_{EKF,t+1,0} : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}$:

$$c_{EKF,t+1,0}(\tilde{x}_{t+1}) := \|\tilde{x}_{t+1} - \mu_{t+1}\|_{\Sigma_{t+1}}^2,$$

which returns the running cost to the form of $c_{EKF,t,0}$.

The theorems below establish that the feature augmentation, feature update, and state propagation steps of the EKF, presented above in our optimization framework, correspond precisely to those presented in the standard EKF SLAM algorithm (Alg. 5) [5, 21]. (For proofs, see Appendix C3).

Theorem 5.1: The feature augmentation step of standard EKF SLAM (Alg. 6) is equivalent to applying a Gauss-Newton step to $c_{EKF,t,1} : \mathbb{R}^{(d_x+pd_f)+p'd_f} \rightarrow \mathbb{R}$, with:

$$\begin{aligned} c_{EKF,t,1}(\tilde{x}_t, f_{t,p+1}, \dots, f_{t,p+p'}) \\ = \|\tilde{x}_t - \mu_t\|_{\Sigma_t}^2 + \sum_{k=p+1}^{p+p'} \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v}^2. \end{aligned}$$

Theorem 5.2: The feature update step of standard EKF SLAM (Alg. 7) is equivalent to applying a Gauss-Newton step on $c_{EKF,t,2} : \mathbb{R}^{d_x+pd_f} \rightarrow \mathbb{R}$, with:

$$c_{EKF,t,2}(\tilde{x}_t) := \|\tilde{x}_t - \mu_t\|_{\Sigma_t}^2 + \sum_{k=1}^p \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v}^2.$$

Theorem 5.3: The state propagation step of standard EKF SLAM (Alg. 8) is equivalent to applying a Marginalization step to $c_{EKF,t,4} : \mathbb{R}^{2d_x+pd_f} \rightarrow \mathbb{R}$, with:

$$c_{EKF,t,4}(\tilde{x}_t, x_{t+1}) := \|\tilde{x}_t - \bar{\mu}_t\|_{\bar{\Sigma}_t}^2 + \|x_{t+1} - g(x_t)\|_{\Sigma_w}^2.$$

where $\tilde{x}_{t,K} := (x_{t+1}, f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{d_x+pd_f}$ and $\tilde{x}_{t,M} = x_t \in \mathbb{R}^{d_x}$.

Remark 5.1: In practice, Gauss-Newton steps for pose augmentation can be delayed and done with feature updates.

Algorithm 1: EKF SLAM on Euclidean spaces, as an iterative optimization problem.

Data: Prior $\mathcal{N}(\mu_0, \Sigma_0)$ on $x_0 \in \mathbb{R}^{d_x}$, noise covariances Σ_w, Σ_v , dynamics map g , measurement map h , time horizon T .

Result: Estimates $\hat{x}_t \in \mathbb{R}^{d_x}, \forall t \in \{1, \dots, T\}$.

```

1  $f_0(x) \leftarrow \|x_0 - \mu_0\|_{\Sigma_0^{-1}}^2$ 
2  $p \leftarrow 0$ .
3 for  $t = 0, 1, \dots, T$  do
4    $(z_{t,p+1}, \dots, z_{t,p+p'}) \leftarrow$  Measurements of new features.
5    $\text{cost}_t \leftarrow \text{cost}_t + \sum_{k=p+1}^{p+p'} \|z_{t,k} - h(x_t, f_k)\|_{\Sigma_v^{-1}}^2$ 
6    $\bar{\mu}_t \leftarrow (\bar{\mu}_t, \ell(x_t, z_{t,p+1}), \dots, \ell(x_t, z_{t,p+p'})) \in$ 
    $\mathbb{R}^{d_x + (p+p')d_f}$ .
7    $\bar{\mu}_t, \bar{\Sigma}_t, \text{cost}_t \leftarrow$  1 Gauss-Newton step on  $\text{cost}_t$ , about  $\bar{\mu}_t$ 
   (Alg. 3).
8    $p \leftarrow p + p'$ 
9    $(z_{t,1}, \dots, z_{t,p}) \leftarrow$  Measurements of existing features.
10   $\text{cost}_t \leftarrow \text{cost}_t + \sum_{k=1}^p \|z_{t,k} - h(x_t, f_k)\|_{\Sigma_v^{-1}}^2$ 
11   $\bar{\mu}_t, \bar{\Sigma}_t, \text{cost}_t \leftarrow$  1 Gauss-Newton step on  $\text{cost}_t$ , about  $\mu_t$ ,
   (Alg. 3).
12   $\hat{x}_t \leftarrow \bar{\mu}_t \in \mathbb{R}^{d_x + pd_f}$ .
13  if  $t < T$  then
14     $\text{cost}_t \leftarrow \text{cost}_t + \|x_{t+1} - g(x_t)\|_{\Sigma_w^{-1}}^2$ 
15     $\mu_{t+1}, \Sigma_{t+1}, \text{cost}_t \leftarrow$  1 Marginalization step on
      $\text{cost}_{t+1}$  with  $x_M = x_t$ , about  $(\bar{\mu}_t, g(\bar{\mu}_t))$  (Alg. 4).
16     $\text{cost}_{t+1} \leftarrow \|x_{t+1} - \mu_{t+1}\|_{\Sigma_{t+1}^{-1}}^2$ 
17  end
18 end
19 return  $\hat{x}_0, \dots, \hat{x}_T$ 

```

B. Multi-State Constrained Kalman Filter (MSCKF), on Manifolds

The MSCKF algorithm maintains a full state, $\tilde{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, containing the most recent IMU state, $x_{\text{IMU}} \in \mathcal{X}_{\text{IMU}}$ and n recent poses, $(x_1, \dots, x_n) \in (\mathcal{X}_p)^n$:

$$\tilde{x}_t := (x_{t,\text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n,$$

with mean $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and covariance $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}} + nd_x) \times (d_{\text{IMU}} + nd_x)}$. As new poses are introduced, old poses are discarded, and features are marginalized to update \tilde{x}_t , the mean μ_t , covariance Σ_t , and $n \in \mathbb{N}$ accordingly.

At initialization ($t = 0$), no pose has yet been recorded ($n = 0$), and the full state \tilde{x}_0 is the initial IMU state $\tilde{x}_{0,\text{IMU}} \in \mathcal{X}_{\text{IMU}}$, with mean $\mu_0 \in \mathcal{X}_{\text{IMU}}$ and covariance $\Sigma_0 \in \mathbb{R}^{d_{\text{IMU}} \times d_{\text{IMU}}}$. Thus, $\tilde{x}_0 = \mu_0$ optimizes the initial running cost $c_{\text{MSCKF},0} : \mathcal{X}_{\text{IMU}} \rightarrow \mathbb{R}$ in our algorithm:

$$c_{\text{MSCKF},0,0}(\tilde{x}_0) = \|\tilde{x}_0 \boxminus \mu_0\|_{\Sigma_0^{-1}}^2.$$

Suppose that, at the current time t , the running cost $c_{\text{MSCKF},t,0} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathbb{R}$ is:

$$c_{\text{MSCKF},t,0}(\tilde{x}_t) = \|\tilde{x}_t \boxminus \mu_t\|_{\Sigma_t^{-1}}^2,$$

where $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}} + nd_x) \times (d_{\text{IMU}} + nd_x)}$ denote the mean and covariance of the full state $\tilde{x}_t := (x_{t,\text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ at time t , consisting of the current IMU state and n poses. When a new image is received, the *pose augmentation step* adds a new pose

$x_{n+1} \in \mathcal{X}_p$ (global frame) to \tilde{x}_t , derived from $x_{n+1}^{\text{IMU}} \in \mathcal{X}_{\text{IMU}}$, the IMU position estimate in the global frame, via the map $\psi : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathcal{X}_{\text{IMU}} \rightarrow \mathcal{X}_p$, i.e.,

$$x_{n+1} := \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}}) \in \mathcal{X}_p.$$

The *feature update step* uses features measurements to update the mean and covariance of \tilde{x}_t . In MSCKF, features are discarded if (A) unobserved in the current pose, or (B) $n \geq N_{\text{max}}$, a specified upper bound, in which case $\lfloor N_{\text{max}} \rfloor / 3$ of the n poses, evenly spaced in time, are dropped after features common to these poses are marginalized. Let $S_{z,1}$ and $S_{z,2}$ denote sets of pose-feature pairs (x_i, f_j) from cases (A) and (B) above, respectively, and let S_f denote the set of features to be marginalized (Alg. 9). These constraints are then incorporated into the running cost, creating a new cost $c_{\text{MSCKF},t,2} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathbb{R}$:

$$c_{\text{MSCKF},t,2}(\tilde{x}_t) := \|\tilde{x}_t \boxminus \mu_t\|_{\Sigma_t^{-1}}^2 + \sum_{(x_i, f_j) \in S_{z,1} \cup S_{z,2}} \|z_{i,j} \boxminus h(x_i, f_j)\|_{\Sigma_v^{-1}}^2,$$

where $z_{i,j} \in \mathbb{R}^{d_z}$ denotes the feature measurement of feature j observed from pose $x_i \in \mathcal{X}_p$. By using Gauss-Newton linearization, we leverage constraints posed by the measurement residuals to construct an updated mean for \tilde{x}_t , denoted $\bar{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, and an updated covariance for \tilde{x}_t , denoted $\bar{\Sigma}_t \in \mathbb{R}^{(d_{\text{IMU}} + nd_x) \times (d_{\text{IMU}} + nd_x)}$. As a result, our cost will be updated to $c_{\text{MSCKF},t,3} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathbb{R}$:

$$c_{\text{MSCKF},t,3}(\tilde{x}_t) := \|\tilde{x}_t \boxminus \bar{\mu}_t\|_{\bar{\Sigma}_t^{-1}}^2,$$

which assumes the form of $c_{\text{MSCKF},t,0}$.

The *state propagation step* propagates the full state by incorporating dynamics residuals into the running cost $c_{\text{MSCKF},t,0}$, creating a new cost $c_{\text{MSCKF},t,4} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathcal{X}_{\text{IMU}} \rightarrow \mathbb{R}$:

$$c_{\text{MSCKF},t,4}(\tilde{x}_t, x_{t+1,\text{IMU}}) := \|\tilde{x}_t \boxminus \bar{\mu}_t\|_{\bar{\Sigma}_t^{-1}}^2 + \|x_{t+1,\text{IMU}} \boxminus g_{\text{IMU}}(x_t, \text{IMU})\|_{\Sigma_t^{-1}}^2.$$

In effect, $c_{\text{MSCKF},t,4}$ appends the new IMU variable $x_{t+1,\text{IMU}} \in \mathcal{X}_{\text{IMU}}$ to the current full state $\tilde{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, and constrains this new full state via the dynamics residuals. A marginalization step, with $\tilde{x}_{t,K} := (x_{t+1,\text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and $\tilde{x}_{t,M} := x_{t,\text{IMU}} \in \mathcal{X}_{\text{IMU}}$, then removes the previous IMU state, $x_{t,\text{IMU}}$, from the running cost. This produces a mean $\mu_{t+1} \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and a covariance $\Sigma_{t+1} \in \mathbb{R}^{(d_{\text{IMU}} + nd_x) \times (d_{\text{IMU}} + nd_x)}$ for the new MSCKF full state, $\tilde{x}_{t+1} := \tilde{x}_{t,K} = (x_{t+1,\text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$. The running cost is updated to $c_{\text{MSCKF},t+1,0} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathbb{R}$:

$$c_{\text{MSCKF},t+1,0}(\tilde{x}_{t+1}) := \|\tilde{x}_{t+1} \boxminus \mu_{t+1}\|_{\Sigma_{t+1}^{-1}}^2,$$

which returns the running cost to the form of $c_{\text{MSCKF},t,0}$.

The theorems below establish that the feature augmentation, feature update, and state propagation steps of the MSCKF, presented above in our optimization framework, correspond precisely to those presented in the standard MSCKF (Alg. 9) [6]. (For proofs, see Appendix C5).

Algorithm 2: Multi-State Constrained Kalman Filter (MSCKF) on manifolds, as iterative optimization.

Data: Prior $\mathcal{N}(\mu_0, \Sigma_0)$ on $x_{\text{IMU},0} \in \mathcal{X}_{\text{IMU}}$, noise covariances Σ_w, Σ_v , dynamics g_{IMU} , measurement map h , time horizon T , Pose transform ψ (IMU \rightarrow global), $\epsilon > 0$.

Result: Estimates \hat{x}_t for all desired timesteps $t \in \{1, \dots, T\}$.

```

1 costt ← ||x0 ⊖ μ0||Σ02. (Initialize objective function).
2 Sz, Sx, Sz,1, Sz,2 ← φ
3 (n, p) ← (0, 0)
4 for t = 0, ⋯, T do
5   while new pose xn+1 ∈ Xp recorded, new IMU
     measurement not received do
6     costt ← costt + ε-1 ||xn+1 ⊖ ψ(x̃t, xn+1IMU)||22.
7     μt, Σt, costt ← 1 Gauss-Newton costt (Alg. 3),
       about (μt, ψ(μt, xn+1IMU)) with ε → 0.
8     {zn+1,j} ← Feature measurements at xn+1
9     Sz ← Sz ∪ {(xn+1, fj) | fj observed at n + 1}
10    n ← n + 1
11    if n ≥ Nmax - 1 then
12      Sx ← {xi | i mod 3 = 2, and 1 ≤ i ≤ n.}
13      Sz,1 ← {(xi, fj) ∈ Sz | xi ∈ Sx}
14      Sx, feature j observed at each pose in Sx
15    end
16    Sz,2 ← {(xi, fj) ∈ Sz | fj not observed at xn}.
17    costt ←
       costt + ∑(xi, fj) ∈ Sz,1 ∪ Sz,2 ||zi,j ⊖ h(xi, ft,j)||Σv-1
       about μt (Alg. 3)
18    x̃t ← μ̄t ∈ XIMU × (Xp)n.
19    Sz ← Sz \ (Sz,1 ∪ {(xi, fj) | xi ∈ Sx})
20    Reindex poses and features in ascending order.
21    (p, n) ← (p - |Sf|, n - |Sx|)
22  end
23  if t < T then
24    costt ← costt + ||xt+1,IMU ⊖ gIMU(xt,IMU)||Σw2.
25    μt+1, Σt+1, costt ← 1 Marginalization step on costt,
       about (μ̄t, g(μt,IMU)) (Alg. 4)
26  end
27 end
28 return x̂0, ⋯, x̂T ∈ XIMU × (Xp)n

```

Theorem 5.4: The pose augmentation step of the standard MSCKF (Alg. 10) is equivalent to applying a Gauss-Newton step to $c_{\text{MSCKF},t,1} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathcal{X}_{\text{IMU}} \rightarrow \mathbb{R}$, with:

$$c_{\text{MSCKF},t,1}(\tilde{x}_t, x_{n+1}) = \|\tilde{x}_t \ominus \mu_t\|_{\Sigma_t}^2 + \epsilon^{-1} \|x_{n+1} \ominus \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}})\|_2^2,$$

and taking $\epsilon \rightarrow 0$ in the resulting (augmented) mean μ_t and covariance Σ_t .

Theorem 5.5: The feature update step of the standard MSCKF (Alg. 11) is equivalent to applying a Marginalization step to $c_{\text{MSCKF},t,2} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathbb{R}^{|S_f|d_f} \rightarrow \mathbb{R}$, with:

$$c_{\text{MSCKF},t,2}(\tilde{x}_t, f_{S_f}) := \|\tilde{x}_t \ominus \mu_t\|_{\Sigma_t}^2 + \sum_{(x_i, f_j) \in S_{z,1} \cup S_{z,2}} \|z_{i,j} \ominus h(x_i, f_j)\|_{\Sigma_v}^2,$$

where $f_{S_f} \in \mathbb{R}^{|S_f|d_f}$ denotes the stacked vector of all feature positions in S_f (see Alg. 9).

Theorem 5.6: The state propagation step of the standard MSCKF (Alg. 12) is equivalent to applying a Marginalization

step to $c_{\text{MSCKF},t,4} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathcal{X}_{\text{IMU}} \rightarrow \mathbb{R}$, with:

$$c_{\text{MSCKF},t,4}(\tilde{x}_t, x_{t+1, \text{IMU}}) := \|\tilde{x}_t \ominus \mu_t\|_{\Sigma_t}^2 + \|x_{t+1, \text{IMU}} \ominus g_{\text{IMU}}(x_t, \text{IMU})\|_{\Sigma_t}^2.$$

with $\tilde{x}_{t,K} := (x_{t+1, \text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and $\tilde{x}_{t,M} = x_{t, \text{IMU}} \in \mathcal{X}_{\text{IMU}}$.

C. State-of-the-Art SLAM Algorithms

Our framework balances the need for computational efficiency, estimation accuracy, and map precision, tradeoffs observed in design choices of existing SLAM algorithms.

- **Extended Kalman Filter (EKF)** [5, 21, 22]—The EKF iteratively updates position estimates of the current pose and all observed features; all past poses are marginalized. This design favors computational speed over localization precision. A variant, the *iterated Extended Kalman Filter* (iEKF), takes multiple Gauss-Newton steps before marginalization to tune the linearization point. This improves mapping and localization accuracy but increases computation time.
- **Multi-State Constrained Kalman Filter** [6, 7, 23]—The MSCKF iteratively updates a full state, with the current IMU state and n past poses; here $n \leq N_{\text{max}}$, a specified upper bound that trades off accuracy and computational speed. Features are stored separately.
- **Sliding Window Smoother, Fixed-Lag Smoother** [10, 24, 25]—The fixed-lag smoother resembles the MSCKF, but performs multiple steps of Gauss-Newton descent before the marginalization step, to tune the linearization point. This improves mapping and localization accuracies at the cost of increasing computation time.
- **Open Keyframe Visual-Inertial SLAM (OKVIS)** [12]—OKVIS updates a sliding window of “keyframes”, poses deemed most informative, which may be arbitrarily spaced in time. Keyframe poses leaving the sliding window, and associated landmarks, are marginalized. This design aims to improve estimation accuracy by maximizing information encoded by the stored poses, without increasing computation time.
- **GraphSLAM and Bundle Adjustment** [21, 26]—These algorithms solve the full SLAM problem, with no marginalization. Their state estimation can be more accurate than the above algorithms, but also far slower.

VI. EXPERIMENTS

This section describes the empirical performance of different marginalization schemes on pose tracking of real-world data. We examine the MSCKF [6], a standard sliding window filter, and the keyframe-based OKVIS algorithm [12], each implemented as an incremental optimization algorithm of the form presented in this paper.

A. Simulation Settings

Experiments are performed on the EuRoC MAV dataset of stereo image sequences and IMU data [27]. We standardize the front-end across all experiments and implementations, using

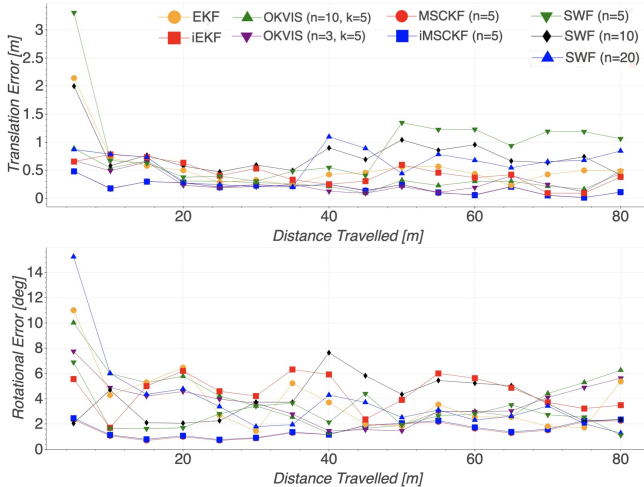


Fig. 1. Localization on Vicon Room 2 (medium). Drift from ground-truth is plotted against the distance traveled along the ground-truth trajectory, sampled at 5 meter intervals. (Note: MSCKF and iMSCKF curves almost overlap). We apply trajectory alignment as in [30].

BRISK keypoint features with brute-force matching. Outlier rejection between the two cameras in the stereo setup is performed using a simple epipolar constraint test, and outlier rejection between stereo frames taken at subsequent timesteps is performed with reprojection distance test using the latest estimate of the feature position and camera pose. We use GTSAM in C++ in the back-end to construct and update costs, compute Jacobians, and implement Gauss-Newton and marginalization steps [8, 28]. To construct dynamics and measurement maps, we collect on-board IMU odometry measurements, and apply the IMU pre-integration scheme in [29]. (Appendix D). We apply trajectory alignment as in [30].

B. Results and Discussion

Localization root-mean-squared error on Vicon Room and Machine Hall sequences from the Euroc MAV dataset are presented in Table I. Due to space constraints, only the estimator drift on the V2_02 sequence is plotted (Fig 1). First, we analyze standard sliding window filters of window size $n = 5, 10, 20$ frames. Features are marginalized when they are only visible in the oldest frame in the optimization window. EKF and iEKF are also included, and are implemented as sliding window filters with window size 1. For the former, only 1 Gauss-Newton step is taken, and in the latter, steps are taken until convergence. Next, we implement MSCKF as an incremental optimization algorithm (Section V-B), with window size $n = 5$, and with two optimization schemes: (1) the standard formulation, with one Gauss-Newton step after marginalization, and (2) a version that takes multiple Gauss-Newton steps until convergence (“Iterated MSCKF,” or iMSCKF). Finally, we implement OKVIS with IMU window size $n = 3, 10$, keyframe window size $k = 5$, and marginalization and keyframe selection schemes identical to those in Leutenegger et al. [12].

Our experiments show that, overall, OKVIS outperforms baseline sliding window filters, even when the latter has a

larger window size. Moreover, our MSCKF implementation outperforms sliding window filters and OKVIS, even under challenging camera motions, despite the latter maintaining nonlinear constraints between camera poses and landmarks, and taking multiple Gauss-Newton steps per iteration. This persists even for sliding window filters with larger window sizes. Taking multiple Gauss-Newton steps in the iMSCKF estimator did not noticeably improve performance over the standard MSCKF, as illustrated in Figure 1.

In contrast with sliding window filter and OKVIS implementations of comparable sizes, the MSCKF recovers better from localization errors, by employing a marginalization scheme that always maintains poses arbitrarily far in the past. This is because older poses represent higher baselines and thus supply better localization information [7]. For instance, the MSCKF maintains the first pose in the estimator for a long time, rendering subsequent estimates more consistent with the initial pose, and thus minimizing drift at the start of the trajectory. In contrast, although OKVIS allows keyframes to be maintained arbitrarily far in the past, keyframes are usually roughly evenly spaced and form a sliding temporal window in camera motions. Thus, earlier poses are quickly marginalized, causing estimates to drift more at the start of the trajectory. Furthermore, the MSCKF includes features in the optimization window only after they have matured, and thus maximally utilizes localization information with fewer updates. Finally, incorporating only matured features ensures that each feature is always initialized through multiple-view triangulation instead of merely stereo triangulation. This minimizes the linearization error when features are marginalized.

VII. CONCLUSION

This paper presents a framework formulating and analyzing optimization and filtering-based SLAM approaches as the iterative application of key algorithm submodules, and proves that it encompasses state-of-the-art filtering algorithms as special cases. Experimental analysis indicate our formulation is useful for analyzing various design choices inherent in these existing SLAM algorithms, and implementing them in a modular fashion for a wide range of robotics applications, which we are eager to test on hardware.

As future work, we wish to apply our analysis to the *dynamic SLAM* problem, which concerns highly mobile features [2, 31] in practical multi-agent interactions, e.g., real-life traffic scenarios [32], by designing marginalization strategies for estimators that jointly track moving and stationary landmarks. We expect good performance on the dynamic SLAM problem, since it enables flexible user-selected design choices.

REFERENCES

- [1] J.J. Leonard and H.F. Durrant-Whyte. “Simultaneous Map Building and Localization for an Autonomous Mobile Robot”. In: *IEEE IROS*. Vol. 3. 1991, pp. 1442–7.
- [2] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard. “Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age”. In: *IEEE T-RO* 32.6 (2016), pp. 1309–1332.
- [3] A. J. Davison. “FutureMapping: The Computational Structure of Spatial AI Systems”. In: *arXiv*. 2018.

Data	MSCKF(n=5)	iMSCKF(n=5)	OKVIS(n=10,k=5)	OKVIS(n=3,k=5)	SWF(n=10)	SWF(n=20)	SWF(n=5)	EKF	iEKF
V1_01	0.09m, 2.87°	0.09m, 2.87°	0.20m, 3.62°	0.23m, 3.68°	0.36m, 3.12°	0.36m, 4.72°	1.00m, 8.22°	1.10m, 8.79°	0.04m, 59.10°
V1_02	0.15m, 1.52°	0.16m, 1.55°	-	-	0.33m, 3.90°	1.16m, 4.69°	0.30m, 3.68°	0.75m, 8.13°	0.42m, 5.48°
V1_03	1.00m, 4.84°	1.12m, 4.90°	0.42m, 6.49°	0.46m, 6.99°	-	6.36m, 25.85°	0.79m, 6.47°	14.84m, 24.58°	1.83m, 8.69°
V2_01	0.14m, 0.83°	0.14m, 0.80°	0.45m, 2.88°	0.39m, 2.58°	0.75m, 5.23°	0.23m, 1.71°	0.26m, 1.96°	0.79m, 3.67°	0.48m, 3.58°
V2_02	0.24m, 1.67°	0.24m, 1.74°	0.44m, 4.69°	0.38m, 4.13°	0.96m, 3.71°	0.75m, 4.49°	1.42m, 2.99°	0.83m, 4.32°	0.47m, 4.14°
MH01	0.07m, 1.03°	0.07m, 1.03°	0.63m, 8.44°	0.72m, 11.61°	0.19m, 1.31°	0.11m, 1.14°	0.43m, 4.02°	0.46m, 3.09°	0.42m, 3.11°
MH02	0.14m, 1.05°	0.19m, 1.86°	0.76m, 8.79°	0.93m, 9.99°	0.23m, 1.89°	0.23m, 1.58°	0.28m, 2.80°	0.34m, 2.21°	0.50m, 2.68°
MH03	0.26m, 1.36°	0.25m, 1.35°	0.64m, 4.29°	0.86m, 5.79°	0.32m, 1.77°	0.54m, 1.75°	0.25m, 1.60°	1.19m, 4.62°	1.33m, 5.23°
MH04	1.11m, 1.62°	1.06m, 1.52°	-	-	0.79m, 1.75°	0.86m, 1.79°	0.59m, 1.28°	4.10m, 4.58°	5.21m, 6.32°

TABLE I

ROOT-MEAN-SQUARED ERROR IN TRANSLATION AND ROTATION ON VICON ROOM AND MACHINE HALL SEQUENCES FROM THE EUROC MAV DATASET. WE APPLY TRAJECTORY ALIGNMENT AS IN [30].

- [4] A. J. Davison and J. Ortiz. “FutureMapping 2: Gaussian Belief Propagation for Spatial AI”. In: *arXiv*. 2019.
- [5] J. Solà. “Simultaneous Localization and Mapping with the Extended Kalman Filter”. In: *arXiv* (2014).
- [6] A. I. Mourikis and S. I. Roumeliotis. “A Multi-State Constraint Kalman Filter for Vision-aided Inertial Navigation”. In: *IEEE ICRA* (2007), pp. 3565–3572.
- [7] M. Li and A. I. Mourikis. “Improving the Accuracy of EKF-based Visual-Inertial Odometry”. In: *IEEE ICRA* (2012), pp. 828–835.
- [8] F. Dellaert. “Factor Graphs and GTSAM: A Hands-on Introduction”. In: *Technical Report GT-RIM-CP&R-2012-002*. 2012.
- [9] M. Kaess, A. Ranganathan, and F. Dellaert. “iSAM: Incremental Smoothing and Mapping”. In: *IEEE T-RO* 24 (Jan. 2009), pp. 1365–1378.
- [10] F. Dellaert and M. Kaess. “Square Root SAM: Simultaneous Localization and Mapping via Square Root Information Smoothing”. In: *IJRR* 25.12 (2006), pp. 1181–1203.
- [11] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J.J. Leonard, and F. Dellaert. “iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree”. In: *IJRR* 31 (2 Feb. 2012), pp. 217–236.
- [12] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale. “Keyframe-based Visual-Inertial Odometry using Non-linear Optimization”. In: *IJRR* 34 (2015), pp. 314–334.
- [13] K. Eickenhoff, P. Geneva, and G. Huang. “Closed-Form Preintegration Methods for Graph-based Visual-Inertial Navigation”. In: *IJRR* 38 (2019), pp. 563–586.
- [14] D. Scaramuzza and F. Fraundorfer. “Visual Odometry, Part I: The First 30 Years and Fundamentals”. In: *IEEE RAM* 18.4 (2011), pp. 80–92.
- [15] F. Fraundorfer and D. Scaramuzza. “Visual Odometry, Part II: Matching, Robustness, Optimization, and Applications”. In: *IEEE RAM* 19.2 (2012), pp. 78–90.
- [16] U. Frese, R. Wagner, and T. Röfer. “A SLAM Overview from a User’s Perspective”. In: *Künstliche Intelligenz* 24 (2010), pp. 191–198.
- [17] S. Huang and G. Dissanayake. “A Critique of Current Developments in Simultaneous Localization and Mapping”. In: *IJARS* 13.5 (2016).
- [18] Kasra Khosoussi, Shoudong Huang, and Gamini Dissanayake. “A Sparse Separable SLAM Back-End”. In: *IEEE Transactions on Robotics* 32.6 (2016).
- [19] H. Strasdat, J.M.M. Montiel, and A. J. Davison. “Visual SLAM: Why Filter?” In: *IVC* 30.2 (2012), pp. 65–77.
- [20] B. Bell. “The Iterated Kalman Smoother as a Gauss-Newton Method”. In: *SIAM J. Optim.* 4 (1994), pp. 626–636.
- [21] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press, 2005.
- [22] Z. Zhang. “Parameter Estimation Techniques: a Tutorial with Application to Conic Fitting”. In: *IVC* 15.1 (1997), pp. 59–76.
- [23] M. Li and A. I. Mourikis. “Optimization-based Estimator Design for Vision-aided Inertial Navigation: Supplemental Materials”. In: *RSS* (2012).
- [24] P. S. Maybeck. *Stochastics Models, Estimation, and Control: Introduction*. Academic Press, Inc., 1979.
- [25] G. Sibley, L. H. Matthies, and G. S. Sukhatme. “Sliding window filter with application to planetary landing.” In: *JFR* 27.5 (2010), pp. 587–608.
- [26] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard. “A Tutorial on Graph-Based SLAM”. In: *IEEE ITSM 2.4* (2010), pp. 31–43.
- [27] M. et al. Burri. “The EuRoC Micro Aerial Vehicle Datasets”. In: *IJRR* (2016).
- [28] F. Dellaert. “Factor Graphs for Robot Perception”. In: *Foundations and Trends® in Robotics* 6.1-2 (2017), pp. 1–139.
- [29] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. “IMU Preintegration on Manifold for Efficient Visual-Inertial Maximum-a-Posteriori Estimation”. In: *RSS*. 2015.
- [30] Shinji Umeyama. “Least-Squares Estimation of Transformation Parameters Between Two Point Patterns”. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 13 (1991), pp. 376–380.
- [31] J. Zhang, M. Henein, R. Mahony, and V. Ila. “VDO-SLAM: A Visual Dynamic Object-aware SLAM System”. In: *arXiv*. 2020.
- [32] D. Fridovich-Keil, E. Ratner, A. Dragan, and C.J. Tomlin. “Efficient Iterative Linear-Quadratic Approximations for Nonlinear Multi-Player General-Sum Differential Games”. In: *arXiv* (2019).
- [33] J. Solà. “Quaternion kinematics for the error-state KF”. In: *arXiv* (2014).

APPENDIX

Please use the following link <https://arxiv.org/pdf/2112.05921.pdf> to access the ArXiv version of the paper with the appendix. The authors will ensure that this link stays active.

The following supplementary material includes the appendix, which contains proofs and figures omitted in the main paper due to space limitations.

A. Appendix for Section III

1) SLAM on Manifolds: Dynamics and Measurement Maps

To formulate SLAM on a manifold, we must alter our definitions of the state variables, features, image positions, dynamics map, and measurement map. Let \mathcal{X} be a smooth manifold of dimension d_x , on which the system state are defined. Similarly, let \mathcal{F} be a smooth manifold of dimension d_f , on which features are defined, and let \mathcal{Z} be the smooth manifold of dimension d_z , on which image measurements are defined (Often, $\mathcal{F} = \mathbb{R}^{d_f}$ and $\mathcal{Z} \in \mathbb{R}^{d_z}$, e.g., with $d_f = 3$ and $d_z = 2$). We then have:

$$\begin{aligned} x_{t+1} &= g(x_t) \boxplus w_t, & w_t &\sim \mathcal{N}(0, \Sigma_w), \\ z_{t,j} &= h(x_t, f_{t,j}) \boxplus v_{t,j}, & v_{t,j} &\sim \mathcal{N}(0, \Sigma_v). \end{aligned}$$

where $x_t \in \mathcal{X}$ denotes the state at time t , $g: \mathcal{X} \rightarrow \mathcal{X}$ denotes the discrete-time dynamics map, and $w_t \in \mathbb{R}^{d_x}$ denotes the dynamics noise, with covariance $\Sigma_w \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_w \succ 0$. Moreover, $f_{t,j} \in \mathcal{F}$ denotes feature position j estimated at the camera pose at time t , $z_{t,j} \in \mathcal{Z}$ denotes the image measurement of feature j measured from the camera pose at time t , $h: \mathcal{X} \times \mathcal{F} \rightarrow \mathcal{Z}$ denotes the measurement map, and $v_t \in \mathbb{R}^{d_z}$ denotes the measurement noise, with covariance $\Sigma_v \in \mathbb{R}^{d_z \times d_z}$, $\Sigma_v \succ 0$.

As before, SLAM concerns an optimization problem over a collection of poses and features, e.g., a sliding window of the most recent poses in the states $\{x_i \in \mathcal{X} | i = t-n+1, \dots, t\}$ and features $\{f_{t,j} \in \mathcal{F} | j = p-q+1, \dots, p\}$:

$$\begin{aligned} \bar{x}_t &:= (x_{t-n+1}, \dots, x_t, f_{t,p-q+1}, \dots, f_{t,p}) \\ &\in \mathcal{X}^n \times \mathcal{F}^q. \end{aligned}$$

We assume that \bar{x}_t is associated with a prior distribution with mean $\mu_0 \in \mathcal{X}^n \times \mathcal{F}^q$ and covariance $\Sigma_0 \in \mathbb{R}^{(nd_x+qd_f) \times (nd_x+qd_f)}$.

2) SLAM as an Optimization Problem on Manifolds

In this subsection, we interpret the SLAM problem on manifolds as the optimization of a cost function $c: \mathcal{X}^n \times \mathcal{F}^q \rightarrow \mathbb{R}$, constructed from residual terms of the same dimension of the minimal coordinates of \bar{x}_t , x_t and z_t . In particular, we must generalize (3) to the case where the states, dynamics and measurement maps are defined on and between manifolds. This involves replacing + and - operators with \boxplus and \boxminus operators, when necessary. For example, the sliding-window filter window presented in Section II, would be associated with the cost $c: \mathcal{X}^n \times \mathcal{F}^q \rightarrow \mathbb{R}$, given by:

$$\begin{aligned} c(\bar{x}_t) &= \|\bar{x}_t \boxminus \mu_0\|_{\Sigma_0^{-1}}^2 \\ &+ \sum_{i=t-n+1}^{t-1} \|x_{i+1} \boxminus g(x_i)\|_{\Sigma_w^{-1}}^2 \\ &+ \sum_{j=p-q+1}^p \sum_{i=t-n+1}^t \|z_{i,j} \boxminus h(x_i, f_{t,j})\|_{\Sigma_v^{-1}}^2 \end{aligned}$$

Similar to Section II, we stack all residual terms into a single residual vector $C(\bar{x}_t)$. For example, for the sliding-window filter above, we have:

$$\begin{aligned} C(\bar{x}_t) &:= \left[(\Sigma_0^{-1/2}(\bar{x}_t \boxminus \mu_0))^\top \right. \\ &(\Sigma_w^{-1/2}(x_{t-n+1} \boxminus g(x_{t-n})))^\top \cdots (\Sigma_w^{-1/2}(x_t \boxminus g(x_{t-1})))^\top \\ &(\Sigma_v^{-1/2}(z_{t-n+1,p-q+1} \boxminus h(x_{t-n+1}, f_{t,p-q+1})))^\top \cdots \\ &(\Sigma_v^{-1/2}(z_{t-n+1,p} \boxminus h(x_{t-n+1}, f_{t,p})))^\top \cdots \\ &(\Sigma_v^{-1/2}(z_{t,p-q+1} \boxminus h(x_t, f_{t,p-q+1})))^\top \cdots \\ &\left. (\Sigma_v^{-1/2}(z_{t,p} \boxminus h(x_t, f_{t,p})))^\top \right]^\top \\ &\in \mathbb{R}^{(2n-1)d_x + pd_f + nqd_z}. \end{aligned}$$

As a result, $c(\bar{x}_t) = C(\bar{x}_t)^\top C(\bar{x}_t)$, and the SLAM problem is now reduced to the following nonlinear least squares problem:

$$\min_{\bar{x}_t} c(\bar{x}_t) = \min_{\bar{x}_t} C(\bar{x}_t)^\top C(\bar{x}_t) \quad (11)$$

Section B4 introduces the main algorithmic submodules used to find an approximate solution to (11).

3) Jacobians Under Box Operators

In subsequent proofs, we require the following results regarding the behavior of Jacobian matrices under the \boxplus and \boxminus operators. In particular, we will focus on the case where the box operator acts on elements of $SO(3)$, since $SO(3)$ is the non-Euclidean-space Lie group that appears most often in this paper.

Definition A.1 (Jacobian on $SO(3)$): We say that $f: SO(3) \rightarrow SO(3)$ is differentiable at $R \in SO(3)$, with Jacobian denoted by:

$$\frac{\partial}{\partial \theta} f(R) \in \mathbb{R}^{3 \times 3},$$

if the following equality holds:

$$\lim_{\delta \theta \rightarrow 0} \frac{\|f(R \boxplus \delta \theta) \boxminus f(R) - \frac{\partial f}{\partial \theta} \delta \theta\|_2}{\|\delta \theta\|_2} = 0.$$

Theorem A.2: Suppose $f: SO(3) \rightarrow SO(3)$ is differentiable. Then, for any fixed $y \in SO(3)$:

$$\frac{\partial}{\partial \theta} (y \boxminus f(R)) = -\frac{\partial f}{\partial \theta} + O(y \boxminus f(R)).$$

Proof: Using the definition of the \boxminus operator for $SO(3)$, we have, for any $\delta \theta \in \mathbb{R}^3$:

$$\begin{aligned} &y \boxminus f(R \boxplus \delta \theta) \\ &= y \boxminus \left(f(R) \boxplus \frac{\partial f}{\partial \theta} \delta \theta \right) + o(\delta \theta) \\ &= y \boxminus \left(f(R) \text{Exp} \left(\frac{\partial f}{\partial \theta} \delta \theta \right) \right) + o(\delta \theta) \\ &= \text{Log} \left(\text{Exp} \left(-\frac{\partial f}{\partial \theta} \delta \theta \right) f(R)^\top y \right) + o(\delta \theta) \\ &= \text{Log} \left(\text{Exp} \left(-\frac{\partial f}{\partial \theta} \delta \theta \right) \text{Exp}(y \boxminus f(R)) \right) + o(\delta \theta) \\ &= y \boxminus f(R) + J_\ell^{-1}(y \boxminus f(R)) \cdot \left(-\frac{\partial f}{\partial \theta} \delta \theta \right) + o(\delta \theta). \end{aligned}$$

Here, we have defined $J_\ell^{-1} : \mathbb{R}^3 \setminus \{0\} \rightarrow SO(3)$ by:

$$J_\ell^{-1}(\theta) = I - \frac{1}{2}\theta^\wedge + \left(\frac{1}{\|\theta\|^2} - \frac{1 + \cos\|\theta\|}{2\|\theta\| \cdot \sin\|\theta\|} \right) (\theta^\wedge)^2,$$

in accordance with [33]. Here, $\wedge : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 3}$ denotes the wedge operator. Thus, we have:

$$\frac{\partial}{\partial \theta} (y \boxplus f(R)) = -J_\ell^{-1}(y \boxplus f(R)) \cdot \frac{\partial f}{\partial \theta}.$$

The theorem statement now follows. \blacksquare

Theorem A.3: Suppose $f : SO(3) \rightarrow SO(3)$ is differentiable. Then for any fixed $y \in SO(3)$:

$$\frac{\partial}{\partial \theta} (f(R) \boxplus y) = \frac{\partial f}{\partial \theta} + O(f(R) \boxplus y).$$

Proof: Using the definition of the \boxplus operator for $SO(3)$, we have, for any $\delta\theta \in \mathbb{R}^3$:

$$\begin{aligned} & f(R \boxplus \delta\theta) \boxplus y \\ &= \left(f(R) \boxplus \frac{\partial f}{\partial \theta} \delta\theta \right) \boxplus y + o(\delta\theta) \\ &= \left(f(R) \text{Exp} \left(\frac{\partial f}{\partial \theta} \delta\theta \right) \right) \boxplus y + o(\delta\theta) \\ &= \text{Log} \left(y^\top f(R) \text{Exp} \left(\frac{\partial f}{\partial \theta} \delta\theta \right) \right) + o(\delta\theta) \\ &= \text{Log} \left(\text{Exp}(f(R) \boxplus y) \text{Exp} \left(\frac{\partial f}{\partial \theta} \delta\theta \right) \right) + o(\delta\theta) \\ &= f(R) \boxplus y + J_r^{-1}(f(R) \boxplus y) \cdot \left(\frac{\partial f}{\partial \theta} \delta\theta \right) + o(\delta\theta), \end{aligned}$$

which implies:

$$\frac{\partial}{\partial \theta} (f(R) \boxplus y) = J_r^{-1}(f(R) \boxplus y) \cdot \frac{\partial f}{\partial \theta}.$$

Here, we have defined $J_r^{-1} : \mathbb{R}^3 \setminus \{0\} \rightarrow SO(3)$ by:

$$J_r^{-1}(\theta) = I + \frac{1}{2}\theta^\wedge + \left(\frac{1}{\|\theta\|^2} - \frac{1 + \cos\|\theta\|}{2\|\theta\| \cdot \sin\|\theta\|} \right) (\theta^\wedge)^2,$$

in accordance with [33]. The theorem statement now follows. \blacksquare

In subsequent discussions, when we consider dynamics and measurement models of the form $y = f(R) \boxplus n$, where $n \in \mathbb{R}^3$ denotes small-magnitude, zero-sum noise, terms of order $O(y \boxplus f(R))$ and $O(f(R) \boxplus y)$ are often ignored.

B. Appendix for Section IV

1) Algorithm Sub-blocks

Algorithm 3: Gauss-Newton Step.

Data: Objective $C^\top C$, linearization point \bar{x}_t^* .

Result: Mean μ , covariance Σ after a Gauss-Newton step.

- 1 $J \leftarrow \frac{\partial C}{\partial \bar{x}_t} \Big|_{\bar{x}_t^*}$
 - 2 $\Sigma_t \leftarrow J^\top J$
 - 3 $\mu_t \leftarrow \bar{x}_t^* - (J^\top J)^{-1} J^\top C(\bar{x}_t^*)$
 - 4 **return** μ_t, Σ_t
-

Algorithm 4: Marginalization

Data: Objective $f = C^\top C$, vector of variables to marginalize $x_{t,M}$, linearization point x_t^* .

Result: Mean $\mu_{t,K}$ and covariance $\Sigma_{t,K}$ of non-marginalized variables $x_{t,K}^*$.

- 1 $C \leftarrow$ subvector of C containing entries dependent on x_M .
 - 2 $J := [J_K \quad J_M] \leftarrow \left[\frac{\partial C}{\partial \bar{x}_{t,K}} \Big|_{x_t^*} \quad \frac{\partial C}{\partial \bar{x}_{t,M}} \Big|_{x_t^*} \right]$.
 - 3 $\Sigma_{t,K} \leftarrow (J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K) \leftarrow^{-1}$
 - 4 $\mu_{t,K} \leftarrow x_{t,K}^* - \Sigma_{t,K} J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C(x^*)$
 - 5 **return** $\mu_{t,K}, \Sigma_{t,K}$
-

2) Proof of Theorem 4.1 (Gauss-Newton Steps)

Here, we present the proof of Theorem 4.1, reproduced below.

Theorem A.4: (Gauss-Newton Step) Let $\bar{x}_t^* \in \mathbb{R}^d$ denote a given linearization point, and suppose $J := \frac{\partial C}{\partial \bar{x}_t} \in \mathbb{R}^{d_C \times d}$ has full column rank. Then applying a Gauss-Newton step to the cost $c(\bar{x}_t)$, about $\bar{x}_t^* \in \mathbb{R}^d$ yields the new cost:

$$c(\bar{x}_t) = \|\bar{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + o(\bar{x}_t - \bar{x}_t^*),$$

where $\mu_t \in \mathbb{R}^d$ and $\Sigma_t \in \mathbb{R}^{d \times d}$ are given by:

$$\begin{aligned} \Sigma_t &\leftarrow (J^\top J)^{-1}, \\ \mu_t &\leftarrow \bar{x}_t^* - (J^\top J)^{-1} J^\top C(\bar{x}_t^*). \end{aligned}$$

Proof: We have:

$$\begin{aligned} c(\bar{x}_t) &= C(\bar{x}_t)^\top C(\bar{x}_t) \\ &= [C(\bar{x}_t^*) + J(\bar{x}_t - \bar{x}_t^*)]^\top [C(\bar{x}_t^*) + J(\bar{x}_t - \bar{x}_t^*)] \\ &\quad + o(\bar{x}_t - \bar{x}_t^*) \\ &= (\bar{x}_t - \mu_t)^\top \Sigma_t^{-1} (\bar{x}_t - \mu_t) + c_0(\bar{x}_t^*) + o(\bar{x}_t - \bar{x}_t^*), \end{aligned}$$

where $c_0(\bar{x}_t^*) \in \mathbb{R}$ denotes a scalar-valued function of \bar{x}_t^* that is independent of the variable \bar{x}_t . This concludes the proof. \blacksquare

3) Proof of Theorem 4.2 (Marginalization Steps)

Theorem A.5 (Marginalization Step): Let $\bar{x}_t^* \in \mathbb{R}^d$ denote a given linearization point, and suppose $J := \frac{\partial C}{\partial \bar{x}_t} \in \mathbb{R}^{d_C \times d}$ has full column rank. Define $J_K := \frac{\partial C}{\partial \bar{x}_{t,K}} \in \mathbb{R}^{d_C \times d_K}$ and $J_M := \frac{\partial C}{\partial \bar{x}_{t,M}} \in \mathbb{R}^{d_C \times d_M}$. If $C_2(\bar{x}_{t,M}, \bar{x}_{t,K})$ were a linear function of $\bar{x}_t = (\bar{x}_{t,M}, \bar{x}_{t,K})$, then applying a Marginalization step to the cost $c(\bar{x}_t)$, about the linearization point $\bar{x}_t^* = (\bar{x}_{t,K}^*, \bar{x}_{t,M}^*) \in \mathbb{R}^d$ yields:

$$\min_{\bar{x}_{t,M}} c_2(\bar{x}_{t,K}, \bar{x}_{t,M}) = \|\bar{x}_{t,K} - \mu_{t,K}\|_{\Sigma_{t,K}^{-1}}^2,$$

where $\Sigma_{t,K} \in \mathbb{R}^{d_K \times d_K}$ and $\mu_{t,K} \in \mathbb{R}^{d_K}$ are given by:

$$\begin{aligned} \Sigma_{t,K} &:= (J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K) \leftarrow^{-1}, \\ \mu_{t,K} &:= \bar{x}_{t,K}^* - \Sigma_{t,K} J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C(\bar{x}_t^*). \end{aligned}$$

Proof: It suffices to show that:

$$\begin{aligned} & \min_{\bar{x}_{t,M}} c_2(\bar{x}_{t,K}, \bar{x}_{t,M}) \\ &= (\bar{x}_{t,K} - \mu_{t,K})^\top \Sigma_{t,K}^{-1} (\bar{x}_{t,K} - \mu_{t,K}) + c'(\bar{x}_t^*). \end{aligned}$$

To do so, we first note that since $C_2(\bar{x}_t)$ is linear in \bar{x}_t :

$$c_2(\bar{x}_t) = \|C_2(\bar{x}_t)\|_2^2 = \|C_2(\bar{x}_t^*) + J_2 \Delta \bar{x}_t\|_2^2$$

$$= \|C_2(\bar{x}_t^*) + J_K \Delta \bar{x}_{t,K} + J_M \Delta \bar{x}_{t,M}\|_2^2.$$

By the method of least-squares, the optimal $\Delta \bar{x}_{t,M}$ is given by the normal equation:

$$\Delta \bar{x}_{t,M} = -(J_M^\top J_M)^{-1} J_M^\top (C_2(\bar{x}_t^*) + J_K \Delta \bar{x}_{t,K})$$

Substituting back into our expression for $c(\bar{x}_t)$, we have:

$$\begin{aligned} & \min_{\bar{x}_t, M} . c_2(\bar{x}_t) \\ &= \|(I - J_M (J_M^\top J_M)^{-1} J_M^\top) (C_2(\bar{x}_t^*) + J_K \Delta \bar{x}_{t,K})\|_2^2 \\ &= (C_2(\bar{x}_t^*) + J_K \Delta \bar{x}_{t,K})^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] \\ & \quad (C_2(\bar{x}_t^*) + J_K \Delta \bar{x}_{t,K}) \\ &= (\bar{x}_{t,K} - \bar{x}_{t,K}^*)^\top \underbrace{J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K}_{:= \Sigma_K^{-1}} \\ & \quad (\bar{x}_{t,K} - \bar{x}_{t,K}^*) \\ & \quad + 2(\bar{x}_{t,K} - \bar{x}_{t,K}^*)^\top J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}_t^*) \\ & \quad + C_2(\bar{x}_t^*)^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}_t^*) \\ &= (\bar{x}_{t,K} - \bar{x}_{t,K}^* + \Sigma_K J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}_t^*))^\top \Sigma_K^{-1} \\ & \quad \underbrace{:= -\mu_K} \\ & \quad (\bar{x}_{t,K} - \bar{x}_{t,K}^* + \Sigma_K J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}_t^*)) \\ & \quad \underbrace{:= -\mu_K} \\ & \quad + C_2(\bar{x}_t^*) (I - J_M (J_M^\top J_M)^{-1} J_M^\top) \\ & \quad + c'(\bar{x}_t^*) \\ &= (\bar{x}_{t,K} - \mu_K)^\top \Sigma_K^{-1} (\bar{x}_{t,K} - \mu_K) + c'(\bar{x}_t^*). \end{aligned}$$

with Σ_K and μ_K as defined in the theorem statement, and $c'(\bar{x}_t^*) \in \mathbb{R}$ as a term independent of \bar{x}_t . ■

4) Main Algorithm on Manifolds

Recall the states, poses, features, dynamics and measurement maps, and costs defined in Appendix A2. In particular, the cost $c(\bar{x})$ is given by:

$$c(\bar{x}) = C(\bar{x})^\top C(\bar{x}) \quad (12)$$

Now, let \bar{x}^* be a chosen linearization point. Let $\widehat{C}_{\bar{x}^*} := C \circ \pi_{\bar{x}^*}^{-1}$ be the coordinate representation of the function C near \bar{x}^* . Recall that $\widehat{C}_{\bar{x}^*}$ is simply a function from one Euclidean space to another. We can now Taylor expand to write:

$$\begin{aligned} C(\bar{x}) &= (C \circ \pi_{\bar{x}^*}^{-1})(\pi_{\bar{x}^*}(\bar{x})) = \widehat{C}_{\bar{x}^*}(\Delta \bar{\chi}) \\ &= \widehat{C}_{\bar{x}^*}(0) + J \Delta \bar{\chi} + o(\Delta \bar{\chi}), \end{aligned} \quad (13)$$

where $\Delta \bar{\chi} = \bar{x} \boxminus \bar{x}^*$ and J is the Jacobian of $\widehat{C}_{\bar{x}^*}$ with respect to $\Delta \bar{\chi}$ evaluated at zero. We then apply a modified version of the algorithms from Section IV-A:

- 1) *Gauss-Newton Descent*: Used to update the current linearization point $\bar{x}^{(k)}$ to a new linearization point $\bar{x}^{(k+1)}$.

$$\bar{x}^{(k+1)} \leftarrow \bar{x}^{(k)} \boxplus (- (J^\top J)^{-1} J^\top C(\bar{x}^{(k)})) \quad (14)$$

After Gauss-Newton steps have been taken, the linearization point \bar{x}^* is fixed, and all or part of the original optimization problem, reproduced below:

$$\min_{\bar{x}} c(\bar{x}) = \min_{\bar{x}} C(\bar{x})^\top C(\bar{x})$$

is replaced with the following linear least squares optimization problem:

$$\min_{\bar{x}} . (\bar{x} \boxminus \mu)^\top \Sigma^{-1} (\bar{x} \boxminus \mu) \quad (15)$$

where the algorithm assigns:

$$\begin{aligned} \mu &\leftarrow \bar{x}^* \boxplus (J^\top J)^{-1} J^\top C(\bar{x}^*) \\ \Sigma &\leftarrow (J^\top J)^{-1}. \end{aligned}$$

- 2) *Marginalization*: Used to remove variables \bar{x}_M from the optimization problem by applying linear approximation to C —in particular, the optimization problem:

$$\min_{\bar{x}} . c(\bar{x})$$

is approximated by:

$$\min_{\bar{x}_K} . (\bar{x}_K \boxminus \mu_K)^\top \Sigma_K^{-1} (\bar{x}_K \boxminus \mu_K),$$

where the algorithm assigns:

$$\begin{aligned} \mu_K &\leftarrow \bar{x}_K^* \boxplus \\ & \quad (- \Sigma_K J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_2(\bar{x}^*)) \\ \Sigma_K &\leftarrow (J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K)^{-1}, \end{aligned}$$

C. Appendix for Section V

1) Continuous to Discrete Time Formulation

The robotic systems analyzed by most mainstream SLAM algorithms are described using continuous-time dynamics models. The implementation of these algorithms thus involves discrete-time propagation. We sketch this process below, in a manner which is both mathematically aligned with the existing literature and convenient for illustrating our optimization algorithm, which operates using a discrete-time dynamics model.

To begin, let \mathbb{H} and \mathbb{H}_u denote the space of quaternions and the space of unit quaternions, respectively. Recall that the time derivative of a smooth curve on \mathbb{H}_u , e.g., $q \in \mathbb{H}_u$, can be expressed as an element of \mathbb{H} ; in particular, there exists some $\omega \in \mathbb{R}^3$ such that:

$$\dot{q} = q \star \begin{bmatrix} 0 \\ \omega \end{bmatrix} \in \mathbb{H}. \quad (16)$$

Let \mathcal{X} be the d_x -dimensional Lie group inhabited by the robot state $x \in \mathcal{X}$, given by the finite Cartesian product of Euclidean spaces and unit quaternions. In other words, \mathcal{X} can be written as:

$$\mathcal{X} := \mathcal{X}_1 \times \cdots \times \mathcal{X}_N,$$

where \mathcal{X}_k is either \mathbb{H}_u or an Euclidean space, for each $k \in \{1, \dots, N\}$. Inspired by (16), we define:

$$\mathcal{X}' := \mathcal{X}'_1 \times \cdots \times \mathcal{X}'_N,$$

where, for each $k \in \{1, \dots, N\}$, we have $\mathcal{X}'_k = \mathbb{H}$ if $\mathcal{X}_k = \mathbb{H}_u$, and $\mathcal{X}'_k = \mathcal{X}_k$ otherwise.

The fundamental question underlying the time discretization process is as follows. Suppose the continuous-time dynamics model for the robot state is given by:

$$\dot{x}(s) = g_{ct}(x(s), w_c(s)), \quad \forall s \in \mathbb{R}, \quad (17)$$

with $g_{ct} : \mathcal{X} \times \mathbb{R}^{d_x} \rightarrow \mathcal{X}'$ smooth, and with $w_c(s) \in \mathbb{R}^{d_x}$ as zero-mean white noise with autocorrelation $\mathbb{E}[w_c(s)w_c(s')] = K \cdot \delta(s - s')$ for each $s \in \mathbb{R}$, where $K \in \mathbb{R}^{d_x \times d_x}$, $K \succ 0$, and $\delta(\cdot)$ denotes the Dirac delta function. Suppose $w_c(s)$ and deviations in $x(s)$ are treated as first-order perturbations. Can we straightforwardly construct an approximate discrete-time additive noise model, of the form:

$$x_{t+1} = g(x_t) \boxplus w(t), \quad (18)$$

where $x_t := x(t_0) \in \mathcal{X}$, $g : \mathcal{X} \rightarrow \mathcal{X}$ smooth, and $x_{t+1} := x(t_1) \in \mathcal{X}$ and $w(t) \in \mathbb{R}^{d_x}$, whose evolution agrees with that of (17) up to first-order perturbations?

We answer the above question in the affirmative, by detailing below the time discretization process used throughout the remainder of this section. With a slight abuse of notation, we write the discrete-time states as $x_t := x(t_0) \in \mathcal{X}$ and $x_{t+1} := x(t_1) \in \mathcal{X}$ for some $t_0, t_1 \in \mathbb{R}$, with $t_0 < t_1$. The main idea is to first perform non-linear integration on a set of nominal dynamics, which ignores error in the initial state and the dynamics noise. We then correct for these perturbations by assuming that they evolve as a linear dynamical system.

- Nominal state dynamics:

We define the *nominal state dynamics* as:

$$\dot{\hat{x}}(s) = g_{ct}(\hat{x}(s), 0) \quad \forall s \in \mathbb{R}.$$

This is the set of dynamics that would be obeyed in the absence of the white noise $w_c(s)$. If the solution $\hat{x}(s)$ uniquely exists, we can construct a smooth function $\hat{g} : \mathcal{X} \rightarrow \mathcal{X}$ that maps $\hat{x}(t_0)$ to $\hat{x}(t_1)$, i.e.:

$$\hat{x}(t_1) = \hat{g}(\hat{x}(t_0)), \quad (19)$$

which describes the discrete-time propagation of the nominal state $x(t)$.

- Error state dynamics:

To characterize the drift of the true state $x(s)$ away from the nominal state $\hat{x}(s)$, we define the error state $\delta x(s) := x(s) \boxminus \hat{x}(s) \in \mathbb{R}^{d_x}$, and characterize the corresponding *error-state dynamics* as:

$$\begin{aligned} \dot{\delta x}(s) &= g_{ct}(x(s), w_c(s)) \boxminus g_{ct}(\hat{x}(s), 0) \\ &= \frac{\partial f}{\partial x}(\hat{x}(s), 0) \cdot \delta x(s) \\ &\quad + \frac{\partial f}{\partial w_c}(\hat{x}(s), 0) \cdot w_c(s) + o(\delta x(s) \cdot w_c(s)), \end{aligned} \quad (20)$$

By approximating the dynamics of the error state $\delta x(s) \in \mathbb{R}^{d_x}$ as a linear, time-varying system, we can write, for each $s \in (t_0, t_1)$:

$$\begin{aligned} \delta x(s) &= \Phi(s, t_0) \delta x(t_0) \\ &\quad + \int_{t_0}^s \Phi(s, \tau) \cdot \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot w_c(\tau) d\tau, \end{aligned} \quad (21)$$

where $\Phi(s, t_0) \in \mathbb{R}^{d_x \times d_x}$ is the state transition matrix satisfying:

$$\frac{d}{ds} \Phi(s, t_0) = \frac{\partial f}{\partial x}(\hat{x}(s), 0) \cdot \Phi(s, t_0).$$

- True State dynamics—Discrete-time propagation of mean and covariance:

From the discrete-time propagation of the nominal and error dynamics, i.e., (19) and (21), we have:

$$\begin{aligned} x(t_1) &= \hat{x}(t_1) \boxplus \delta x(t_1) \\ &= \hat{g}(\hat{x}(t_0)) \boxplus \left(\Phi(t_1, t_0) \delta x(t_0) \right. \\ &\quad \left. + \int_{t_0}^{t_1} \Phi(t_1, \tau) \cdot \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot w_c(\tau) d\tau \right) \\ &= \left(\hat{g}(\hat{x}(t_0)) \boxplus \Phi(t_1, t_0) \delta x(t_0) \right) \\ &\quad \boxplus \left(\int_{t_0}^{t_1} \Phi(t_1, \tau) \cdot \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot w_c(\tau) d\tau \right) \\ &\quad + o(\delta x(t_0), w_c(\cdot)), \end{aligned}$$

where we have used the fact that the \boxplus operator, which operates through:

$$\text{Exp}(\delta\theta + \delta\phi) = \text{Exp}(\delta\theta) \cdot \text{Exp}(\delta\phi) + o(\delta\theta \cdot \delta\phi)$$

for any $x \in \mathbb{H}$ and $\delta\theta, \delta\phi \in \mathbb{R}^3$.

Finally, by identifying:

$$\begin{aligned} x_{t+1} &\leftarrow x(t_1) \in \mathcal{X}, \\ g(x_t) &\leftarrow \hat{x}(t_1) \boxplus \Phi(t_1, t_0) \delta x(t_0) \in \mathcal{X}, \\ w_t &\leftarrow \int_{t_0}^{t_1} \Phi(t_1, \tau) \cdot \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot w_c(\tau) d\tau \in \mathbb{R}^{d_x}, \end{aligned}$$

we obtain our discrete-time dynamics model:

$$x_{t+1} = g(x_t) \boxplus w_t.$$

Notice that $w_t \in \mathbb{R}^{d_x \times d_x}$ is a zero-mean random variable, with covariance $\Sigma_w \in \mathbb{R}^{d_x \times d_x}$ given by:

$$\begin{aligned} \Sigma_w &:= \mathbb{E}[w_t w_t^\top] \\ &= \mathbb{E} \left[\int_{t_0}^{t_1} \int_{t_0}^{t_1} \Phi(t_1, \tau) \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot w_c(\tau) w_c(\tau')^\top \right. \\ &\quad \left. \cdot \left(\frac{\partial f}{\partial w_c}(\hat{x}(\tau'), 0) \right)^\top \Phi(t_1, \tau')^\top d\tau d\tau' \right] \\ &= \int_{t_0}^{t_1} \Phi(t_1, \tau) \frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \cdot K \cdot \left(\frac{\partial f}{\partial w_c}(\hat{x}(\tau), 0) \right)^\top \\ &\quad \Phi(t_1, \tau)^\top d\tau \end{aligned}$$

where we have used the fact that $\mathbb{E}[w_c(\tau)w_c(\tau')^\top] = K \cdot \delta(\tau - \tau')$, with $\delta(\cdot)$ denoting the Dirac delta function.

2) EKF, Setup

The Extended Kalman Filter (EKF), whose standard formulation is presented in Algorithm 5, is an iterative algorithm for updating estimates of the current pose x_t (i.e. $n = 1$) and positions of all observed features at the current time, $f_t := (f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{pd_f}$. This corresponds to the sliding

window filter in our formulation, with $n = 1$ and $q = p$. Below, as an application of our optimization-based SLAM framework, we present the dynamics and measurement maps of the EKF algorithm in \mathbb{R}^2 , as well as the associated cost functions. Dimension-wise, in its standard formulation, the 2D EKF is an instantiation of Algorithm 5 with $d_x = 3$, $d_f = 2$, and $d_z = 2$. To unify our notation, we will suppose that d_x, d_f, d_z assume these values throughout the rest of this section.

Let $x_t := (x_t^1, x_t^2, \theta_t) \in \mathbb{R}^{d_x}$ denote the *robot pose*, comprising its position and angle in \mathbb{R}^{d_f} , let $f_{t,k} := (f_{t,k}^1, f_{t,k}^2) \in \mathbb{R}^{d_f}$ denote the position of each *feature* $f_k \in \{f_1, \dots, f_p\}$ visible at time t , and let $z_{t,k} := (z_{t,k}^1, z_{t,k}^2) \in \mathbb{R}^{d_z}$ denote the measurement of feature f_k at time t . The dynamics map $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$, with $\dot{x}_t = g(x_t)$ is obtained by performing numerical integration on the continuous-time dynamics:

$$\begin{aligned}\dot{x}_t^1 &= v \cos \theta + w_t^1, \\ \dot{x}_t^2 &= v \sin \theta + w_t^2, \\ \dot{\theta}_t &= \omega + w_t^3,\end{aligned}$$

where $w_t := (w_t^1, w_t^2, w_t^3) \in \mathbb{R}^{d_x}$ denotes additive zero-mean Gaussian noise on the (x, y, θ) coordinates of the state variable, respectively, with joint covariance $w_t \sim \mathcal{N}(0, \Sigma_w)$ for some covariance matrix $\Sigma_w \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_w \succ 0$. For more details regarding the numerical integration process, see Section C1 in the Appendix.

The measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$ is given by:

$$\begin{aligned}z_{t,k}^1 &= f_{t,k}^1 - x_t^1 + v_t^1, \\ z_{t,k}^2 &= f_{t,k}^2 - x_t^2 + v_t^2,\end{aligned}$$

where $v_t := (v_t^1, v_t^2) \in \mathbb{R}^{d_z}$ denotes additive zero-mean Gaussian noise on the measurements $z_{t,j}^1, z_{t,j}^2 \in \mathbb{R}$, respectively, with joint covariance $v_t \sim \mathcal{N}(0, \Sigma_v)$ for some covariance matrix $\Sigma_v \in \mathbb{R}^{d_z \times d_z}$, $\Sigma_v \succ 0$. The measurement vector $z_t \in \mathbb{R}^{p d_z}$ is then given by concatenating each of the q residual measurements obtained at time t , i.e. $z_t := (z_{t,1}, \dots, z_{t,p}) \in \mathbb{R}^{p d_z}$.

3) Proofs from Section V-A

Theorem A.6: The feature augmentation step of the standard EKF SLAM algorithm (Alg. 6) is equivalent to applying a Gauss-Newton step to $c_{EKF,t,1} : \mathbb{R}^{d_x + p d_f} \rightarrow \mathbb{R}$, given by:

$$\begin{aligned}c_{EKF,t,1}(\tilde{x}_t, f_{t,p+1}, \dots, f_{t,p+p'}) \\ = \|\tilde{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + \sum_{k=p+1}^{p+p'} \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v^{-1}}^2.\end{aligned}$$

Proof:

To simplify the analysis below, we assume all degrees of freedom of new features are observed. More specifically, we assume the existence of an *inverse observation map* $\ell : \mathbb{R}^{d_x} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_f}$, satisfying $h(x_t, \ell(x_t, z_t)) = z_t$ for each $x_t \in \mathbb{R}^{d_x}, z_t \in \mathbb{R}^{d_z}$, which directly generates position estimates of new features from their feature measurements and the current pose, by effectively ‘‘inverting’’ the measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$ [5]. When full observations are unattainable, the missing degrees of freedom are introduced as a prior to the system [5]; in this case, similar results follow.

Algorithm 5: Extended Kalman Filter SLAM, Standard Formulation.

Data: Prior distribution on $x_0 \in \mathbb{R}^{d_x} : \mathcal{N}(\mu_0, \Sigma_0)$, dynamics and measurement noise covariances $\Sigma_w \in \mathbb{R}^{d_x \times d_x}, \Sigma_v \in \mathbb{R}^{d_z \times d_z}$, (discrete-time) dynamics map $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$, measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{p d_f} \rightarrow \mathbb{R}^{d_z}$, time horizon $T \in \mathbb{N}$.
Result: Estimates \hat{x}_t for all desired timesteps $t \leq T$.

```

1 for  $t = 0, \dots, T$  do
2   if detect new feature measurements
3      $z_{t,p+1:p+p'} := (z_{t,p+1}, \dots, z_{t,p+p'}) \in \mathbb{R}^{p' d_z}$  then
4        $\mu_t, \Sigma_t, p \leftarrow$  Alg. 6, EKF feature augmentation
5        $(\mu_t, \Sigma_t, p, z_{t,p+1:p+p'}, h(\cdot))$ 
6     end
7      $z_{t,1:p} := (z_{t,1}, \dots, z_{t,p}) \in \mathbb{R}^{p d_z} \leftarrow$  New measurements
8     of existing features.
9      $\bar{\mu}_t, \bar{\Sigma}_t \leftarrow$  Alg. 7, EKF feature update
10     $(\bar{\mu}_t, \bar{\Sigma}_t, z_{t,1:p}, h(\cdot))$ .
11  if  $t < T$  then
12     $\mu_{t+1}, \Sigma_{t+1} \leftarrow$  Alg. 8, EKF state propagation
13     $(\mu_t, \Sigma_t, g(\cdot))$ 
14  end
15 end
16 return  $\hat{x}_0, \dots, \hat{x}_T \in \mathbb{R}^{d_x}$ .
```

First, to simplify notation, define:

$$\begin{aligned}z_{t,p+1:p+p'} &= (z_{t,p+1}, \dots, z_{t,p+p'}) \in \mathbb{R}^{p' d_z}, \\ f_{t,p+1:p+p'} &= (f_{t,p+1}, \dots, f_{t,p+p'}) \in \mathbb{R}^{p' d_f}, \\ \tilde{h}(x_t, f_{t,p+1:p+p'}) &:= (h(x_t, f_{t,p+1}), \dots, h(x_t, f_{t,p+p'})) \\ &\in \mathbb{R}^{p' d_z}, \\ \tilde{\Sigma}_v &= \text{diag}\{\Sigma_v, \dots, \Sigma_v\} \in \mathbb{R}^{p' d_z \times p' d_z}.\end{aligned}$$

We can now rewrite the cost $c_{EKF,t,1}$ as:

$$\begin{aligned}c_{EKF,t,1}(\tilde{x}_t, f_{t,p+1:p+p'}) \\ = \|\tilde{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + \|z_{t,p+1:p+p'} - \tilde{h}(x_t, f_{t,p+1:p+p'})\|_{\tilde{\Sigma}_v^{-1}}^2.\end{aligned}$$

To apply a Gauss-Newton step, our first task is to find a vector $C_1(\tilde{x}_t, f_{t,p+1:p+p'})$ of an appropriate dimension such that $c_{EKF,t,1}(\tilde{x}_t, f_{t,p+1:p+p'}) = C_1(\tilde{x}_t, f_{t,p+1:p+p'})^\top C_1(\tilde{x}_t, f_{t,p+1:p+p'})$. A natural choice is furnished by $C_1(\tilde{x}_t, f_{t,p+1:p+p'}) \in \mathbb{R}^{d_x + p' d_z}$, as defined below:

$$\begin{aligned}C_1(\tilde{x}_t, f_{t,p+1:p+p'}) \\ := \begin{bmatrix} \Sigma_t^{-1/2}(\tilde{x}_t - \mu_t) \\ \Sigma_v^{-1/2}(z_{t,p+1:p+p'} - \tilde{h}(x_t, f_{t,p+1:p+p'})) \end{bmatrix}.\end{aligned}$$

Thus, our parameters for the Gauss-Newton algorithm submodule are:

$$\begin{aligned}\tilde{x}_t^* &:= (x_t^*, f_{t,1:p}^*, f_{t,p+1:p+p'}^*) \\ &= (\bar{\mu}_t, \ell(x_t^*, z_{t,p+1}), \dots, \ell(x_t^*, z_{t,p+p'})) \in \mathbb{R}^{d_x + (p+p')d_f}, \\ &\text{where } x_t^* \in \mathbb{R}^{d_x}, f_{t,1:p}^* \in \mathbb{R}^{p d_f}, f_{t,p+1:p+p'}^* \in \mathbb{R}^{p' d_f}, \\ C_1(\tilde{x}_t^*) &= \begin{bmatrix} \Sigma_t^{-1/2}(\tilde{x}_t^* - \mu_t) \\ \tilde{\Sigma}_v^{-1/2}(z_{t,p+1:p+p'} - \tilde{h}(x_t^*, f_{t,p+1:p+p'}^*)) \end{bmatrix}\end{aligned}$$

Algorithm 6: Extended Kalman Filter, Feature Augmentation Sub-block.

Data: Current EKF state $\tilde{x}_t \in \mathbb{R}^{d_x+pd_f}$, with mean $\mu_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$, current number of features p , observations of new features at current pose $z_{t,p+1:p+p'} := (z_{t,p+1}, \dots, z_{t,p+p'}) \in \mathbb{R}^{p'd_z}$, measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$, inverse measurement map $\ell : \mathbb{R}^{d_x} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_f}$.

Result: Updated number of features p , updated EKF state mean $\mu_t \in \mathbb{R}^{d_x+pd_f}$, covariance $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$ (with p already updated)

- 1 $(\mu_{t,x}, \mu_{t,f,1:p}) \leftarrow \mu_t \in \mathbb{R}^{d_x+pd_f}$, with $\mu_{t,x} \in \mathbb{R}^{d_x}$, $\mu_{t,f,1:p} \in \mathbb{R}^{pd_f}$.
- 2 $\ell : \mathbb{R}^{d_x} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_f} \leftarrow$ Inverse measurement map, satisfying $z_{t,k} = h(x_t, \ell(x_t, z_{t,k}))$ for each $x_t \in \mathbb{R}^{d_x}$, $z_{t,k} \in \mathbb{R}^{d_z}$, $\forall k = p+1, \dots, p+p'$.
- 3 $\tilde{\ell}(\mu_{t,x}, z_{t,p+1}, \dots, z_{t,p+p'}) \leftarrow (\ell(\mu_{t,x}, z_{t,p+1}), \dots, \ell(\mu_{t,x}, z_{t,p+p'})) \in \mathbb{R}^{p'd_f \times (d_x+p'd_z)}$
- 4 $\mu_t \leftarrow (\mu_t, \tilde{\ell}(\mu_{t,x}, z_{t,p+1}, \dots, z_{t,p+p'})) \in \mathbb{R}^{d_x+(p+p')d_f}$
- 5 $\begin{bmatrix} \Sigma_{t,xx} & \Sigma_{t,xf} \\ \Sigma_{t,fx} & \Sigma_{t,ff} \end{bmatrix} \leftarrow \Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$, with $\Sigma_{t,xx} \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_{t,xf} = \Sigma_{t,fx}^\top \in \mathbb{R}^{d_x \times pd_f}$, $\Sigma_{t,ff} \in \mathbb{R}^{pd_f \times pd_f}$.
- 6 $L_x \leftarrow \frac{\partial \tilde{\ell}}{\partial x} \Big|_{(\mu_t, z_t)} \in \mathbb{R}^{p'd_f \times d_x}$
- 7 $L_z \leftarrow \frac{\partial \tilde{\ell}}{\partial z} \Big|_{(\mu_t, z_t)} \in \mathbb{R}^{p'd_f \times p'd_z}$
- 8 $\tilde{\Sigma}_v \leftarrow \text{diag}\{\Sigma_v, \dots, \Sigma_v\} \in \mathbb{R}^{p'd_z \times p'd_z}$
- 9 $\Sigma_t \leftarrow \begin{bmatrix} \Sigma_{t,xx} & \Sigma_{t,xf} & \Sigma_{t,xx}L_x^\top \\ \Sigma_{t,fx} & \Sigma_{t,ff} & \Sigma_{t,fx}L_x^\top \\ L_x\Sigma_{t,xx} & L_x\Sigma_{t,xf} & L_x\Sigma_{t,xx}L_x^\top + L_z\tilde{\Sigma}_vL_z^\top \end{bmatrix} \in \mathbb{R}^{(d_x+(p+p')d_f) \times (d_x+(p+p')d_f)}$
- 10 $p \leftarrow p+p'$
- 11 **return** $\mu_t \in \mathbb{R}^{d_x+pd_f}$, $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$, $p \geq 0$

$$= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \mathbb{R}^{d_x+pd_f+p'd_z},$$

$$J = \begin{bmatrix} \Sigma_t^{-1/2} & O \\ -\tilde{\Sigma}_v^{-1/2}\tilde{H}_{t,x} [I \ O] & -\tilde{\Sigma}_v^{-1/2}\tilde{H}_{t,f} \end{bmatrix}$$

$$\in \mathbb{R}^{(d_x+pd_f+p'd_z) \times (d_x+(p+p')d_f)},$$

where $\tilde{H}_t := [\tilde{H}_{t,x} \ \tilde{H}_{t,f}] \in \mathbb{R}^{p'd_z \times (d_x+p'd_f)}$ is defined as the Jacobian of $\tilde{h} : \mathbb{R}^{d_x} \times \mathbb{R}^{p'd_f} \rightarrow \mathbb{R}^{p'd_z}$ at $(x_t^*, f_{t,p+1:p+p'}^*) \in \mathbb{R}^{d_x+p'd_f}$, with $\tilde{H}_{t,x} \in \mathbb{R}^{p'd_z \times d_x}$ and $\tilde{H}_{t,f} \in \mathbb{R}^{p'd_z \times pd_f}$. By Algorithm 3, the Gauss-Newton update is thus given by:

$$\begin{aligned} & \Sigma_t \\ & \leftarrow (J^\top J)^{-1} \\ & = \left(\begin{bmatrix} \Sigma_t^{-1/2} & -[O] \\ O & -\tilde{\Sigma}_v^{-1/2}\tilde{H}_{t,f} \end{bmatrix} \right)^{-1} \end{aligned} \quad (23)$$

Algorithm 7: Extended Kalman Filter, Feature Update Sub-block.

Data: Current EKF state $\tilde{x}_t \in \mathbb{R}^{d_x+pd_f}$, with mean $\mu_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$, new measurements of existing features $z_{t,1:p} := (z_{t,1}, \dots, z_{t,p}) \in \mathbb{R}^{pd_z}$, measurement map $h : \mathbb{R}^{d_x} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$

Result: Updated EKF state mean $\mu_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$

- 1 $f_{t,1:p} \leftarrow (f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{pd_f}$.
- 2 $\tilde{h}(x_t, f_{t,1:p}) \leftarrow (h(x_t, f_{t,1}), \dots, h(x_t, f_{t,p})) \in \mathbb{R}^{pd_z}$
- 3 $H_t \leftarrow \frac{\partial \tilde{h}}{\partial (x_t, f_{t,1:p})} \Big|_{\mu_t}$ Jacobian of $\tilde{h} : \mathbb{R}^{d_x} \times \mathbb{R}^{pd_f} \rightarrow \mathbb{R}^{pd_z}$ evaluated at $\mu_t \in \mathbb{R}^{d_x+pd_f}$.
- 4 $\tilde{\Sigma}_v \leftarrow \text{diag}\{\Sigma_v, \dots, \Sigma_v\} \in \mathbb{R}^{pd_z \times pd_z}$.
- 5 $\bar{\mu}_t \leftarrow \mu_t + \Sigma_t H_t^\top (H_t \Sigma_t H_t^\top + \tilde{\Sigma}_v)^{-1} (z_{t,1:p} - \tilde{h}(\mu_t, f_{t,1:p})) \in \mathbb{R}^{d_x+pd_f}$.
- 6 $\bar{\Sigma}_t \leftarrow \Sigma_t - \Sigma_t H_t^\top (H_t \Sigma_t H_t^\top + \tilde{\Sigma}_v)^{-1} H_t \Sigma_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$.
- 7 **return** $\bar{\mu}_t \in \mathbb{R}^{d_x+pd_f}$, $\bar{\Sigma}_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$.

Algorithm 8: Extended Kalman Filter, State Propagation Sub-block.

Data: Current EKF state $\tilde{x}_t \in \mathbb{R}^{d_x+pd_f}$, with mean $\bar{\mu}_t \in \mathbb{R}^{d_x+pd_f}$ and covariance $\bar{\Sigma}_t \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$, (discrete-time) dynamics map $g : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$

Result: Propagated EKF state mean $\mu_{t+1} \in \mathbb{R}^{d_x+pd_f}$ and covariance $\Sigma_{t+1} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$

- 1 $(\bar{\mu}_{t,x}, \bar{\mu}_{t,f,1:p}) \leftarrow \bar{\mu}_t$, with $\bar{\mu}_{t,x} \in \mathbb{R}^{d_x}$, $\bar{\mu}_{t,f,1:p} \in \mathbb{R}^{pd_f}$.
- 2 $\begin{bmatrix} \bar{\Sigma}_{t,xx} & \bar{\Sigma}_{t,xf} \\ \bar{\Sigma}_{t,fx} & \bar{\Sigma}_{t,ff} \end{bmatrix} \leftarrow \bar{\Sigma}_t \in \mathbb{R}^{d_x \times d_x}$, with $\bar{\Sigma}_{t,xx} \in \mathbb{R}^{d_x \times d_x}$, $\bar{\Sigma}_{t,xf} = \bar{\Sigma}_{t,fx}^\top \in \mathbb{R}^{d_x \times pd_f}$, $\bar{\Sigma}_{t,ff} \in \mathbb{R}^{pd_f \times pd_f}$.
- 3 $G_t \leftarrow \frac{\partial g}{\partial x} \Big|_{\bar{\mu}_{t,x}}$.
- 4 $\mu_{t+1} \leftarrow (g(\bar{\mu}_t), \bar{\mu}_{t,f,1:p}) \in \mathbb{R}^{d_x+pd_f}$.
- 5 $\Sigma_{t+1} \leftarrow \begin{bmatrix} G_t \bar{\Sigma}_{t,xx} G_t^\top + \Sigma_w & G_t \bar{\Sigma}_{t,xf} \\ \bar{\Sigma}_{t,fx} G_t^\top & \bar{\Sigma}_{t,ff} \end{bmatrix} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$.
- 6 **return** $\mu_{t+1} \in \mathbb{R}^{d_x+pd_f}$, $\Sigma_{t+1} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$.

$$\left[\begin{array}{cc} \Sigma_t^{-1/2} & O \\ -\tilde{\Sigma}_v^{-1/2}\tilde{H}_{t,x} [I \ O] & -\tilde{\Sigma}_v^{-1/2}\tilde{H}_{t,f} \end{array} \right]^{-1} \quad (24)$$

$$= \begin{bmatrix} \Sigma_t^{-1} + [O] & \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} [I \ O] & [O] \\ \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} [I \ O] & \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} & [O] \\ \Omega_{t,xx} + \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} & \Omega_{t,xf} & \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} \\ \Omega_{t,fx} & \Omega_{t,ff} & O \\ \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} & O & \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} \end{bmatrix}^{-1}, \quad (25)$$

$$\begin{aligned} \bar{\mu}_t & \leftarrow \tilde{x}_t^* - (J^\top J)^{-1} J^\top C_1(\tilde{x}_t^*) \\ & = (\bar{\mu}_t, \ell(x_t^*, z_{t,p+1}), \dots, \ell(x_t^*, z_{t,p+p'})). \end{aligned}$$

Here, we have defined $\Omega_{t,xx} \in \mathbb{R}^{d_x \times d_x}$, $\Omega_{t,xf} = \Omega_{t,fx}^\top \in$

$\mathbb{R}^{d_x \times pd_f}$ and $\Omega_{t,ff} \in \mathbb{R}^{pd_f \times pd_f}$ by:

$$\begin{bmatrix} \Omega_{t,xx} & \Omega_{t,xf} \\ \Omega_{t,fx} & \Omega_{t,ff} \end{bmatrix} := \begin{bmatrix} \Sigma_{t,xx} & \Sigma_{t,xf} \\ \Sigma_{t,fx} & \Sigma_{t,ff} \end{bmatrix}^{-1} \quad (26)$$

To conclude the proof, we must show that (25) is identical to the update equations for covariance matrix in the standard formulation of the Extended Kalman Filter algorithm, i.e., we must show that:

$$\begin{bmatrix} \Sigma_{t,xx} & \Sigma_{t,xf} & \Sigma_{t,xx} L_x^\top \\ \Sigma_{t,fx} & \Sigma_{t,ff} & \Sigma_{t,fx} L_x^\top \\ L_x \Sigma_{t,xx} & L_x \Sigma_{t,xf} & L_x \Sigma_{t,xx} L_x^\top + L_z \Sigma_v L_z^\top \end{bmatrix} \cdot \begin{bmatrix} \Omega_{t,xx} + \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} & \Omega_{t,xf} & \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} \\ \Omega_{t,fx} & \Omega_{t,ff} & O \\ \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} & O & \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} \end{bmatrix}$$

equals the $(d_x + (p + p')d_f) \times (d_x + (p + p')d_f)$ identity matrix. This follows by applying (26), as well as the matrix equalities resulting from taking the derivative of the equation $z_t := h(x_t, \ell(x_t, z_t))$ with respect to $x_t \in \mathbb{R}^{d_x}$ and $z_t \in \mathbb{R}^{d_z}$, respectively:

$$\begin{aligned} I &= \tilde{H}_{t,f} L_z, \\ O &= \tilde{H}_{t,x} + H_{t,f} L_x. \end{aligned}$$

Theorem A.7: The feature update step of the standard EKF SLAM algorithm (Alg. 7) is equivalent to applying a Gauss-Newton step on $c_{EKF,t,2} : \mathbb{R}^{d_x + pd_f} \rightarrow \mathbb{R}$, given by:

$$\begin{aligned} c_{EKF,t,2}(\tilde{x}_t) \\ := \|\tilde{x}_t - \mu_t\|_{\Sigma_t^{-1}}^2 + \sum_{k=1}^p \|z_{t,k} - h(x_t, f_{t,k})\|_{\Sigma_v^{-1}}^2. \end{aligned}$$

Proof: First, to simplify notation, define:

$$\begin{aligned} z_{t,1:p} &:= (z_{t,1}, \dots, z_{t,p}) \in \mathbb{R}^{pd_z}, \\ f_{t,1:p} &:= (f_{t,1}, \dots, f_{t,p}) \in \mathbb{R}^{pd_f}, \\ \tilde{h}(x_t, f_{t,1:p}) &:= (h(x_t, f_{t,1}), \dots, h(x_t, f_{t,p})) \in \mathbb{R}^{pd_z}, \\ \tilde{\Sigma}_v &:= \text{diag}\{\Sigma_v, \dots, \Sigma_v\} \in \mathbb{R}^{pd_z \times pd_z}. \end{aligned}$$

We can then rewrite the cost as:

$$c_{EKF,t,2}(\tilde{x}_t) = \|\tilde{x}_t^* - \mu_t\|_{\Sigma_t^{-1}}^2 + \|z_{t,1:p} - \tilde{h}(\tilde{x}_t^*)\|_{\tilde{\Sigma}_v^{-1}}^2.$$

To apply a Gauss-Newton step, our first task is to find a vector $C_2(\tilde{x}_t)$ of an appropriate dimension such that $c_{EKF,t,2}(\tilde{x}_t) = C_2(\tilde{x}_t)^\top C_2(\tilde{x}_t)$. A natural choice is furnished by $C_2(\tilde{x}_t) \in \mathbb{R}^{d_x + pd_f + pd_z}$, as defined below:

$$C_2(\tilde{x}_t) := \begin{bmatrix} \Sigma_t^{-1/2}(\tilde{x}_t - \mu_t) \\ \tilde{\Sigma}_v^{-1/2}(z_{t,1:p} - \tilde{h}(\tilde{x}_t)) \end{bmatrix}.$$

Thus, our parameters for the Gauss-Newton algorithm submodule are:

$$\begin{aligned} \tilde{x}_t^* &= \mu_t \in \mathbb{R}^{d_x + pd_f}, \\ C_2(\tilde{x}_t^*) &= \begin{bmatrix} \Sigma_t^{-1/2}(\tilde{x}_t^* - \mu_t) \\ \tilde{\Sigma}_v^{-1/2}(z_{t,1:p} - \tilde{h}(\tilde{x}_t^*)) \end{bmatrix} = \begin{bmatrix} 0 \\ \tilde{\Sigma}_v^{-1/2}(z_{t,1:p} - \tilde{h}(\mu_t)) \end{bmatrix} \\ &\in \mathbb{R}^{d_x + pd_f + pd_z}, \end{aligned}$$

$$J = \begin{bmatrix} \Sigma_t^{-1/2} \\ -\tilde{\Sigma}_v^{-1/2} H_t \end{bmatrix} \in \mathbb{R}^{(d_x + pd_f + pd_z) \times (d_x + pd_f)},$$

where $\tilde{H}_t \in \mathbb{R}^{pd_z} \times \mathbb{R}^{d_x + pd_f}$ is defined as the Jacobian of $\tilde{h} : \mathbb{R}^{d_x} \times \mathbb{R}^{pd_f} \rightarrow \mathbb{R}^{pd_z}$ at $\tilde{x}_t^* \in \mathbb{R}^{d_x + pd_f}$. By Algorithm 3, the Gauss-Newton update is thus given by:

$$\begin{aligned} \bar{\Sigma}_t &\leftarrow (J^\top J)^{-1} \\ &= (\Sigma_t^{-1} + H_t^\top \Sigma_v H_t)^{-1} \\ &= \Sigma_t - \Sigma_t H_t^\top (\Sigma_v^{-1} + H_t \Sigma_t^{-1} H_t^\top)^{-1} H_t \Sigma_t, \\ \bar{\mu}_t &\leftarrow \mu_t - (J^\top J)^{-1} J^\top C_2(\tilde{x}_t^*) \\ &= \mu_t - (\Sigma_t^{-1} + H_t^\top \Sigma_v^{-1} H_t)^{-1} \begin{bmatrix} \Sigma_t^{-1/2} & -H_t^\top \Sigma_v^{-1/2} \\ 0 & \Sigma_v^{-1/2}(z_{t,1:p} - \tilde{h}(\mu_t)) \end{bmatrix} \\ &= \mu_t + (\Sigma_t^{-1} + H_t^\top \Sigma_v^{-1} H_t)^{-1} H_t^\top \Sigma_v^{-1} (z_{t,1:p} - \tilde{h}(\mu_t)), \\ &= \mu_t + \Sigma_v^{-1} H_t^\top (\Sigma_t^{-1} + H_t^\top \Sigma_v^{-1} H_t)^{-1} (z_{t,1:p} - \tilde{h}(\mu_t)), \end{aligned}$$

which are identical to the feature update equations for the mean and covariance matrix in the Extended Kalman Filter algorithm, i.e. (4) and (5) respectively. Note that, in the final step, we have used a variant of the Woodbury Matrix Identity. ■

Theorem A.8: The state propagation step of the standard EKF SLAM algorithm (Alg. 8) is equivalent to applying a Marginalization step to $c_{EKF,t,4} : \mathbb{R}^{2d_x + pd_f} \rightarrow \mathbb{R}$, given by:

$$\begin{aligned} c_{EKF,t,4}(\tilde{x}_t, x_{t+1}) \\ := \|\tilde{x}_t - \bar{\mu}_t\|_{\bar{\Sigma}_t^{-1}}^2 + \|x_{t+1} - g(x_t)\|_{\Sigma_w^{-1}}^2. \end{aligned}$$

Proof: Intuitively, the state propagation step marginalizes out $\tilde{x}_t \in \mathbb{R}^{d_x}$ and retain $x_{t+1} \in \mathbb{R}^{d_x}$. In other words, in the notation of our Marginalization algorithm submodule, we have:

$$\begin{aligned} \tilde{x}_{t,K} &= x_{t+1} \in \mathbb{R}^{d_x + pd_f}, \\ \tilde{x}_{t,M} &= \tilde{x}_t \in \mathbb{R}^{d_x + pd_f}. \end{aligned}$$

To apply a marginalization step, our first task is to find vectors $C_K(x_K) = C_K(\tilde{x}_t)$ and $C_M(x_K, x_M) = C_M(\tilde{x}_t, x_{t+1})$ of appropriate dimensions such that $c_{EKF,t,4}(\tilde{x}_t, x_{t+1}) = C_K(x_{t+1})^\top C_K(x_{t+1}) + C_M(\tilde{x}_t, x_{t+1})^\top C_M(\tilde{x}_t, x_{t+1})$. A natural choice is furnished by $C_K(x_{t+1}) \in \mathbb{R}$ and $C_M(\tilde{x}_t, x_{t+1}) \in \mathbb{R}^{d_x}$, as defined below:

$$\begin{aligned} c_K(x_{t+1}) &= 0 \\ c_M(\tilde{x}_t, x_{t+1}) &= \|\tilde{x}_t - \bar{\mu}_t\|_{\bar{\Sigma}_t^{-1}}^2 + \|x_{t+1} - g(x_t)\|_{\Sigma_w^{-1}}^2. \end{aligned}$$

where we have identified the following parameters, in the language of a Marginalization step (Section II):

$$\begin{aligned} C_K(\tilde{x}_{t,K}) &= 0 \in \mathbb{R} \\ C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}) &= \begin{bmatrix} \bar{\Sigma}_t^{-1/2}(\tilde{x}_t - \bar{\mu}_t) \\ \Sigma_w^{-1/2}(x_{t+1} - g(x_t)) \end{bmatrix} \in \mathbb{R}^{2d_x + pd_f}. \end{aligned}$$

For convenience, we will define the pose and feature track components of the mean $\mu_t \in \mathbb{R}^{d_x + pd_f}$ by $\mu_t := (\mu_{t,x}, \mu_{t,f}) \in \mathbb{R}^{d_x + pd_f}$, with $\mu_{t,x} \in \mathbb{R}^{d_x}$ and $\mu_{t,f} \in \mathbb{R}^{pd_f}$, respectively. This mirrors our definition of $x_t \in \mathbb{R}^{d_x}$ and $f_{t,1:p} \in \mathbb{R}^{pd_f}$ as the

components of the full state $\tilde{x}_t := (x_t, f_{t,1:p}) \in \mathbb{R}^{d_x+pd_f}$. In addition, we will define the components of $\bar{\Sigma}_t^{-1/2} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$ and $\bar{\Sigma}_t^{-1} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}$ by:

$$\begin{aligned} \begin{bmatrix} \Omega_{t,xx} & \Omega_{t,xf} \\ \Omega_{t,fx} & \Omega_{t,ff} \end{bmatrix} &:= \bar{\Sigma}_t^{-1} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}, \\ \begin{bmatrix} \Lambda_{t,xx} & \Lambda_{t,xf} \\ \Lambda_{t,fx} & \Lambda_{t,ff} \end{bmatrix} &:= \bar{\Sigma}_t^{-1/2} \in \mathbb{R}^{(d_x+pd_f) \times (d_x+pd_f)}, \end{aligned}$$

where $\Sigma_{t,xx}, \Lambda_{t,xx} \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_{t,xf}, \Lambda_{t,xf} \in \mathbb{R}^{d_x \times pd_f}$, $\Sigma_{t,fx}, \Lambda_{t,fx} \in \mathbb{R}^{pd_f \times d_x}$, and $\Sigma_{t,ff}, \Lambda_{t,ff} \in \mathbb{R}^{pd_f \times pd_f}$. Using the above definitions, we can reorder the residuals in $C_K \in \mathbb{R}$ and $C_M \in \mathbb{R}^{2d_x+pd_f}$, and thus redefine them by:

$$\begin{aligned} C_K(\tilde{x}_{t,K}) &= 0 \in \mathbb{R} \\ C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}) &= \begin{bmatrix} \Lambda_{t,xx}(x_t - \mu_{t,x}) + \Lambda_{t,xf}(f_{t,1:p} - \mu_{t,f}) \\ \Sigma_w^{-1/2}(x_{t+1} - g(x_t)) \\ \Lambda_{t,fx}(x_t - \mu_{t,x}) + \Lambda_{t,ff}(f_{t,1:p} - \mu_{t,f}) \end{bmatrix} \\ &\in \mathbb{R}^{2d_x+pd_f}. \end{aligned}$$

Our state variables and cost functions for the Gauss-Newton algorithm submodule are:

$$\begin{aligned} \overline{x_M^*} &= \tilde{x}_t^* = \overline{\mu_t} \in \mathbb{R}^{d_x+pd_f}, \\ \overline{x_K^*} &= g(\tilde{x}_t^*) = g(\overline{\mu_t}) \in \mathbb{R}^{d_x+pd_f}, \\ C_K(\tilde{x}_{t,K}^*) &= 0 \in \mathbb{R}, \\ C_M(\tilde{x}_{t,K}^*, \tilde{x}_{t,M}^*) &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \mathbb{R}^{2d_x+pd_f}, \\ J_M &= \begin{bmatrix} O & \Lambda_{xf} \\ \Sigma_w^{-1/2} & O \\ O & \Lambda_{ff} \end{bmatrix} \in \mathbb{R}^{(2d_x+pd_f) \times (d_x+pd_f)} \\ J_K &= \begin{bmatrix} \Lambda_{xx} \\ -\Sigma_w^{-1/2} G_t \\ \Lambda_{xf} \end{bmatrix} \in \mathbb{R}^{(2d_x+pd_f) \times d_x}, \end{aligned}$$

where we have defined G_t to be the Jacobian of $g: \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ at $\overline{\mu_{t,x}} \in \mathbb{R}^{d_x}$, i.e.:

$$G_t := \left. \frac{\partial g}{\partial x_t} \right|_{x_t = \overline{\mu_{t,x}}}$$

Applying the Marginalization equations, we thus have:

$$\begin{aligned} \mu_{t+1} &\leftarrow \tilde{x}_{t,K} - \Sigma_{t+1} J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] \\ &\quad C_M(\overline{x_K^*}, \overline{x_M^*}) \\ &= g(\overline{\mu_t}), \\ \Sigma_{t+1} &\leftarrow (J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K)^{-1}, \\ &= (J_K^\top J_K - J_K^\top J_M (J_M^\top J_M)^{-1} J_M^\top J_K)^{-1}, \\ &= \left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Lambda_{fx} \Lambda_{xf} + \Lambda_{ff}^2 \end{bmatrix} - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Lambda_{fx} \Lambda_{xx} + \Lambda_{ff} \Lambda_{fx} \end{bmatrix} \right. \\ &\quad \left. (\Lambda_{xx} + \Lambda_{xf} \Lambda_{fx} + G_t^\top \Sigma_w^{-1} G_t)^{-1} \right. \\ &\quad \left. \cdot \begin{bmatrix} -G_t^\top \Sigma_w^{-1} & \Lambda_{xx} \Lambda_{xf} + \Lambda_{fx} \Lambda_{ff} \end{bmatrix} \right)^{-1} \end{aligned}$$

$$\begin{aligned} &= \left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Omega_{ff} \end{bmatrix} - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Omega_{fx} \end{bmatrix} \right. \\ &\quad \left. (\Omega_{xx} + G_t^\top \Sigma_w^{-1} G_t)^{-1} \begin{bmatrix} -G_t^\top \Sigma_w^{-1} & \Omega_{xf} \end{bmatrix} \right)^{-1} \end{aligned}$$

To show that this is indeed identical to the propagation equation for the covariance matrix in the Extended Kalman Filter algorithm, i.e. Algorithm 5, Line 5, we must show that:

$$\begin{aligned} &\left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Omega_{ff} \end{bmatrix} - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Omega_{fx} \end{bmatrix} (\Omega_{xx} + G_t^\top \Sigma_w^{-1} G_t)^{-1} \right. \\ &\quad \left. \begin{bmatrix} -G_t^\top \Sigma_w^{-1} & \Omega_{xf} \end{bmatrix} \right)^{-1} \\ &= \begin{bmatrix} G_t \bar{\Sigma}_{t,xx} G_t^\top + \Sigma_w & G_t \bar{\Sigma}_{t,xf} \\ \bar{\Sigma}_{t,xf} G_t^\top & \bar{\Sigma}_{t,ff} \end{bmatrix} \end{aligned}$$

This follows by brute-force expanding the above block matrix components, and applying Woodbury's Matrix Identity, along with the definitions of $\Sigma_{t,xx}$, $\Lambda_{t,xx}$, $\Sigma_{t,xf}$, $\Lambda_{t,xf}$, $\Sigma_{t,fx}$, $\Lambda_{t,fx}$, $\Sigma_{t,ff}$, and $\Lambda_{t,ff}$. ■

4) MSCKF, Setup

The Multi-State Constrained Kalman Filter (MSCKF) algorithm iteratively refines the mean and covariance of a MSCKF full state, consisting of the most recent IMU state and a sliding window of n poses. (Algorithm 9) In particular, when a set of new IMU measurements is obtained, the MSCKF full state is propagated forward in time. When a new image measurement arrives, the current pose is appended to the MSCKF full state vector. Features not observed in the current pose are marginalized. If the number of poses maintained in the MSCKF full state, denoted n , exceeds a pre-specified upper bound N_{\max} , then features common to every third currently maintained pose, starting from the second oldest pose, are marginalized. Below, we discuss the key components of the MSCKF algorithm—the IMU state, poses, MSCKF full states, features, image measurements, dynamics, pose augmentation and measurement maps—in more detail.

The IMU state $x_{t,\text{IMU}}$ takes the form:

$$x_{t,\text{IMU}} := (q_{WS}, v_S, b_g, b_a, r_{WS})(t) \in \mathbb{R}^3 \times \mathbb{H}_u \times \mathbb{R}^9, \quad (27)$$

where we use S and W to represent the sensor frame and world frame, respectively. Here, $r_{WS} \in \mathbb{R}^3$ denotes the position of the IMU sensor frame represented in the world frame, $q_{WS} \in \mathbb{H}_u$ denotes the unit quaternion of axis rotation from world frame to IMU sensor frame, and $R(q_{WS}) \in SO(3)$ denotes the rotation matrix associated with q_{WS} . Moreover, $v_S \in \mathbb{R}^3$ denotes the linear velocity of the IMU sensor frame relative to the world frame, as represented in the world frame, while $b_g \in \mathbb{R}^3$ and $b_a \in \mathbb{R}^3$ denote the sensor biases of the gyroscope and accelerometer, respectively. Finally, $\tilde{\omega}_S \in \mathbb{R}^3$ and $\tilde{a}_S \in \mathbb{R}^3$ denote gyroscope and accelerometer measurements, respectively.

For convenience, define $\mathcal{X}_{\text{IMU}} := \mathbb{R}^3 \times \mathbb{H}_u \times \mathbb{R}^9$ and $\mathcal{X}'_{\text{IMU}} := \mathbb{R}^3 \times \mathbb{H} \times \mathbb{R}^9$. The continuous-time IMU dynamics map $g_{\text{IMU},ct}: \mathcal{X}_{\text{IMU}} \rightarrow \mathcal{X}'_{\text{IMU}}$ is given by:

$$\dot{q}_{WS} = q_{WS} \star \frac{1}{2} \begin{bmatrix} 0 \\ \tilde{\omega}_S - b_g - w_g \end{bmatrix},$$

$$\begin{aligned}
\dot{b}_g &= w_{b_g}, \\
\dot{v}_S &= R(q_{WS})^\top (\tilde{a}_S - b_a + w_a) + g_W, \\
\dot{b}_a &= w_{b_a}, \\
\dot{r}_{WS} &= v_S.
\end{aligned}$$

Algorithm 9: Multi-State Constrained Kalman Filter, Standard Formulation.

Data: Prior distribution on $x_{\text{IMU},0} \in \mathcal{X}_p: \mathcal{N}(\mu_0, \Sigma_0)$, dynamics and measurement noise covariances $\Sigma_w \in \mathbb{R}^{d_x \times d_x}$, $\Sigma_v \in \mathbb{R}^{d_x \times d_x}$, discrete-time dynamics map $g_{\text{IMU}}: \mathbb{R}^{d_{\text{IMU}}} \times \mathbb{R}^{d_{\text{IMU}}}$, measurement map $h: \mathcal{X}_p \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$, time horizon T , pose transformation $\psi: \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathcal{X}_p \rightarrow \mathcal{X}_p$ (IMU \rightarrow global).

Result: Estimates $\hat{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ for all desired timesteps $t \leq T$, where $n :=$ number of poses in \hat{x}_t at time t .

```

1  $S_z, S_x, S_{z,1}, S_{z,2} \leftarrow \phi$ 
2  $(n, p) \leftarrow (0, 0)$ 
3 for  $t = 0, \dots, T$  do
4   while new image  $\mathcal{I}$  with new pose  $x_{n+1} \in \mathcal{X}_p$  recorded,
   next IMU measurement not yet received do
5      $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n, \Sigma_t \in$ 
      $\mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)} \leftarrow$  Alg. 10 ( $\tilde{x}_t, \mu_t, \Sigma_t,$ 
      $x_{n+1}, x_{n+1}^{\text{IMU}}, \psi(\cdot)$ )
6      $\{z_{n+1,j}\}$  feature  $j$  is observed at  $x_{n+1}$   $\leftarrow$  Feature
     measurements at  $x_{n+1}$ 
7      $\{f_j^*\}$  Feature  $j$  is observed at  $x_{n+1}$   $\leftarrow$  Feature
     position estimates at  $x_{n+1}$ .
8     Record new estimates of existing features and first
     estimate of new features at  $x_{n+1} \in \mathcal{X}_p$ .
9      $S_z \leftarrow$ 
      $S_z \cup \{(x_{n+1}, f_j) \mid \text{Feature } j \text{ observed at } n+1\}$ 
10     $n \leftarrow n+1$ 
11    if  $n \geq N_{\text{max}} - 1$  then
12       $S_x \leftarrow \{x_i \mid i \bmod 3 = 2, \text{ and } 1 \leq i \leq n\}$ 
13       $S_{z,1} \leftarrow \{(x_i, f_j) \in S_z \mid x_i \in$ 
       $S_x, \text{ feature } j \text{ observed at each pose in } S_x\}$ 
14    end
15     $S_{z,2} \leftarrow \{(x_i, f_j) \in S_z \mid x_i \in$ 
     $x_{1:n}, \text{ feature } j \text{ observed at } x_i \text{ but not at } x_n\}$ .
16     $S_f \leftarrow \{f_j \mid \exists x_i \in x_{1:n} \text{ s.t. } (x_i, f_j) \in S_{z,1} \cup S_{z,2}\}$ 
17    if  $S_f \neq \phi$  then
18       $\bar{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n, \bar{\Sigma}_t \in$ 
       $\mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)} \leftarrow$  Alg. 11 ( $\tilde{x}_t, \mu_t,$ 
       $\Sigma_t, x_{n+1}, S_{z,1} \cup S_{z,2}, S_f, h(\cdot)$ )
19       $\hat{x}_t \leftarrow \bar{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ .
20    end
21     $S_z \leftarrow S_z \setminus (S_{z,1} \cup \{(x_i, f_j) \mid x_i \in S_x\})$ 
22    Reindex poses and features, in ascending order of
    index, i.e.,  $\{x_1, \dots, x_{n-|S_x|}\}$  and
     $\{f_1, \dots, f_{p-|S_f|}\}$ .
23     $(p, n) \leftarrow (p - |S_f|, n - |S_x|)$ 
24  end
25  if  $t < T$  then
26     $\mu_{t+1} \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n, \Sigma_{t+1} \in$ 
     $\mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)} \leftarrow$  Alg. 12, MSCKF State
    Propagation ( $\tilde{x}_t, \bar{\mu}_t, \bar{\Sigma}_t$ )
27  end
28 end
29 return  $\hat{x}_0, \dots, \hat{x}_T \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ 

```

where \star denotes quaternion multiplication, and $w_g, w_a, w_{b_g}, w_{b_a} \in \mathbb{R}^3$ denote zero-mean standard Gaussian noise.

Each pose $x_k \in \{x_1, \dots, x_n\}$ currently maintained in the sliding window of poses takes the form $x_k := (q_{WC_k}, r_{WC_k}) \in \mathbb{H}_u \times \mathbb{R}^3$, where $r_{WC_k} \in \mathbb{R}^3$ denotes the position of the camera at pose x_k in the world frame, while $q_{WC_k} \in \mathbb{H}_u$ denotes the quaternion associated with the axis rotation from the world frame to the camera frame at pose $x_k \in \mathbb{H} \times \mathbb{R}^3$. For convenience, we define $\mathcal{X}_p := \mathbb{H}_u \times \mathbb{R}^3$.

The MSCKF full state maintained throughout the operation of the MSCKF algorithm contains the IMU state at the current time, as well as a collection of n poses, where n is constrained to remain below a pre-specified, fixed upper bound N_{max} :

$$\tilde{x}_t := (x_{t,\text{IMU}}, x_1, \dots, x_n) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n. \quad (28)$$

The state space is thus $\mathcal{X} := \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$.

When a new image measurement arrives, the estimate of the current camera pose in the IMU frame, denoted $x_{n+1}^{\text{IMU}} \in \mathcal{X}_{\text{IMU}}$, is transformed to the global frame, and appended to the MSCKF full state \tilde{x}_t . This coordinate transformation is realized by the map $\psi: \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathcal{X}_p$, defined by:

$$\begin{aligned}
&\psi(q_{WS}, v_S, b_g, b_a, r_{WS}, q_{WC_1}, r_{WC_1}, \dots, \\
&\quad q_{WC_n}, r_{WC_n}, q_{WI_{n+1}}, r_{WI_{n+1}}) \\
&= (q_{IC} \star q_{WI_{n+1}}, r_{WI_{n+1}} + C(q_{WI_{n+1}})r_{IC}) \\
&:= (q_{WC_{n+1}}, r_{WC_{n+1}}),
\end{aligned}$$

where q_{IC} denotes the quaternion encoding the (fixed) transformation from the IMU frame to the camera frame. In summary, the MSCKF algorithm defines the new pose x_{n+1} and updates the MSCKF full state \tilde{x}_t as follows:

$$\begin{aligned}
x_{n+1} &\leftarrow \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}}) \in \mathcal{X}_p, \\
\tilde{x}_t &\leftarrow (\tilde{x}_t, x_{n+1}) = (\tilde{x}_t, \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}})) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^{n+1},
\end{aligned}$$

with the map ψ as defined above.

When new feature position estimates are detected from a new image measurement, the new camera pose corresponding to this image measurement is appended to \tilde{x}_t , and n is incremented by 1. If $n = N_{\text{max}}$ the upper limit N_{max} , a third of all old poses in \tilde{x}_t is discarded, starting from the second oldest pose. Then, feature measurements, corresponding to features unobserved at the current pose, are marginalized and used to update the mean and covariance of the new MSCKF full state \tilde{x}_t .

As is the case with the EKF algorithm, we assume that the image measurement space and feature space are given by \mathbb{R}^{d_z} and \mathbb{R}^{d_f} , respectively, with $d_z = 2$ and $d_f = 3$. Throughout the duration of the MSCKF algorithm, poses and features are added into, dropped from, and marginalized from the MSCKF full state. Suppose at a given time, the MSCKF maintains n poses in the *MSCKF full state* \tilde{x}_t , and retains

measurements of p features. For each pose $i \in \{1, \dots, n\}$ and feature $j \in \{1, \dots, p\}$ currently maintained in the SLAM algorithm, if feature j were detected at pose i , let $z_{i,j} \in \mathbb{R}^{d_z}$ denote the associated feature measurement. For the MSCKF, the measurement map $h : \mathcal{X}_{\text{IMU}} \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$ is given by:

$$z_{i,j} = h(x_i, f_j) \\ := \frac{1}{(R(q_{WC_k})f_j - r_{WC_k})_z} \left[\begin{array}{c} (R(q_{WC_k})f_j - r_{WC_k})_x \\ (R(q_{WC_k})f_j - r_{WC_k})_y \end{array} \right] (t) \\ + v_{i,j}.$$

where $R(q_{WC_k}) \in SO(3)$ denotes the rotation matrix associated with the quaternion q_{WC_k} , $f_j \in \mathbb{R}^3$ denotes the position of feature j in the world frame, while the subscript indices “ x, y, z ” refer to the respective coordinates of the vector $R(q_{WS})f_j - r_{WS} \in \mathcal{X}_p$. Meanwhile, $v_{i,j} \in \mathbb{R}^{d_z}$ denotes zero-mean standard Gaussian noise in the measurement at time t , with covariance matrix $\Sigma_v \in \mathbb{R}^{d_z \times d_z}$, $\Sigma_v \succ 0$.

When a new image measurement is received, the MSCKF algorithm performs marginalization, described in Section IV-C, using two sets of feature measurements—the set of all feature measurements common to old poses x_i to be dropped, denoted $S_{z,1}$, as well as the set of all feature measurements of features f_j not seen in the current pose, denoted $S_{z,2}$. These are more precisely defined in Section V-B. The measurement vector used for marginalization, denoted $\tilde{z} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}| d_z}$, is then given by concatenating the q residual measurements obtained at times $t - n + 1, \dots, t$, i.e.:

$$\tilde{z} := \{z_{i,j} | (x_i, f_j) \in S_{z,1} \cup S_{z,2}\} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}| d_z}.$$

5) Proofs from Section V-B

Theorem A.9: The pose augmentation step of the standard MSCKF SLAM algorithm (Alg. 10) is equivalent to applying a Gauss-Newton step to $c_{\text{MSCKF},t,1} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathbb{R}$, given by:

$$c_{\text{MSCKF},t,1}(\tilde{x}_t, x_{n+1}) \\ = \|\tilde{x}_t \boxminus \mu_t\|_{\Sigma_t^{-1}}^2 + \epsilon^{-1} \|x_{n+1} \boxminus \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}})\|_2^2,$$

and taking $\epsilon \rightarrow 0$ in the resulting (augmented) mean μ_t and covariance Σ_t .

Proof: We claim that from an optimization perspective, the state augmentation step is equivalent to applying one Gauss-Newton step to the cost function $c_{\text{MSCKF},t,1}(\tilde{x}_t, x_{n+1})$, specified above, and then taking the limit $\epsilon \rightarrow 0$ in the resulting augmented mean $\mu_t(\epsilon) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^{(n+1)}$ and augmented covariance $\mu_t(\epsilon) \in \mathbb{R}^{(d_{\text{IMU}}+(n+1)d_x) \times (d_{\text{IMU}}+(n+1)d_x)}$.

To apply a Gauss-Newton step, our first task is to find a vector $C(\tilde{x}_t, x_{n+1})$ of an appropriate dimension such that $c_{\text{MSCKF},t,1}(\tilde{x}_t, x_{n+1}) = C_1(\tilde{x}_t, x_{n+1})^\top C_1(\tilde{x}_t, x_{n+1})$. A natural choice is furnished by $C_1(\tilde{x}_t, x_{n+1}) \in \mathbb{R}^{d_{\text{IMU}}+(n+1)d_x}$, as defined below:

$$C_1(\tilde{x}_t, x_{n+1}) := \left[\begin{array}{c} \Sigma_t^{-1/2}(\tilde{x}_t \boxminus \mu_t) \\ \epsilon^{-1/2}(x_{n+1} \boxminus \psi(\tilde{x}_t, x_{n+1}^{\text{IMU}})) \end{array} \right].$$

Thus, our parameters for the Gauss-Newton algorithm sub-module are:

$$(\tilde{x}_t^*, x_{n+1}^*) := (\mu_t, \psi(\mu_t, x_{n+1}^{\text{IMU}})) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n,$$

Algorithm 10: Multi-State Constrained Kalman Filter, Pose Augmentation Sub-block.

Data: MSCKF state $\tilde{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, with mean $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and covariance $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$, New pose $x_{n+1} \in \mathcal{X}_p$, measurement of new pose in IMU frame $x_{n+1}^{\text{IMU}} \in \mathcal{X}_p$, Transformation of poses from IMU frame to global frame $\psi : \mathbb{R}^{(d_{\text{IMU}}+nd_x)} \times \mathcal{X}_p \rightarrow \mathcal{X}_p$

Result: Updated MSCKF state mean

$\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and covariance $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$, updated number of poses n .

- 1 $\tilde{x}_t \leftarrow (\tilde{x}_t, x_{n+1}) \in \mathbb{R}^{d_{\text{IMU}}+(n+1)d_x}$, where $x_{n+1} \in \mathcal{X}_p$ is the new pose vector.
 - 2 $\{z_{n+1,j} | \text{Feature } j \text{ is observed at pose } n+1\} \leftarrow$ Feature measurements at pose x_{n+1}
 - 3 $\{f_j^* | \text{Feature } j \text{ is observed at pose } x_{n+1}\} \leftarrow$ Feature position estimates at pose x_{n+1} .
 - 4 $\mu_t \leftarrow (\mu_t, \psi(\mu_t, x_{n+1}^{\text{IMU}})) \in \mathbb{R}^{d_{\text{IMU}}+(n+1)d_x}$, where $\mu_{t,\text{IMU}} \in \mathbb{R}^{d_{\text{IMU}}} :=$ IMU component of μ_t , $x_{n+1}^{\text{IMU}} \in \mathcal{X}_p :=$ pose estimate of x_{n+1} from the IMU frame.
 - 5 $\Sigma_t \leftarrow \left[\begin{array}{c} I_{d_{\text{IMU}}+(n+1)d_x} \\ \frac{\partial \psi}{\partial (\tilde{x}_t, x_{n+1}^{\text{IMU}})} \end{array} \right] \Sigma_t \left[\begin{array}{c} I_{d_{\text{IMU}}+(n+1)d_x} \\ \frac{\partial \psi}{\partial (\tilde{x}_t, x_{n+1}^{\text{IMU}})} \end{array} \right]^\top$
 - 6 **return** $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$, $n \geq 0$
-

$$C_1(\tilde{x}_t^*, x_{n+1}^*) = \left[\begin{array}{c} \Sigma_t^{-1/2}(\tilde{x}_t^* - \mu_t) \\ \epsilon^{-1/2}(x_{n+1}^* - \psi(\tilde{x}_t^*, x_{n+1}^{\text{IMU}})) \end{array} \right] = \left[\begin{array}{c} 0 \\ 0 \end{array} \right] \\ \in \mathbb{R}^{d_{\text{IMU}}+(n+1)d_x}, \\ J = \left[\begin{array}{cc} \Sigma_t^{-1/2} & O \\ -\epsilon^{-1/2}\Psi & \epsilon^{-1/2}I_{d_x} \end{array} \right] \\ \in \mathbb{R}^{(d_{\text{IMU}}+(n+1)d_x) \times (d_{\text{IMU}}+(n+1)d_x)},$$

where $\Psi \in \mathbb{R}^{d_x \times (d_{\text{IMU}}+nd_x)}$ is defined as the Jacobian of $\psi : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \rightarrow \mathcal{X}_p$ with respect to \tilde{x}_t at $(\tilde{x}_t^*, x_{n+1}^{\text{IMU}}) \in \mathbb{R}^{d_{\text{IMU}}+(n+1)d_x}$. By Algorithm 3, the Gauss-Newton update is thus given by:

$$\Sigma_t(\epsilon) \leftarrow (J^\top J)^{-1} = \left[\begin{array}{cc} \Sigma_t^{1/2} & O \\ \Psi \Sigma_t^{1/2} & \epsilon^{1/2} I_{d_x} \end{array} \right] \left[\begin{array}{cc} \Sigma_t^{1/2} & \Sigma_t^{1/2} \Psi^\top \\ O & \epsilon^{1/2} I_{d_x} \end{array} \right] \\ = \left[\begin{array}{cc} \Sigma_t & \Sigma_t \Phi^\top \\ \Psi \Sigma_t & \Psi \Sigma_t \Psi^\top + \epsilon I_{d_x} \end{array} \right], \\ \mu_t(\epsilon) \leftarrow \tilde{x}_t^* - (J^\top J)^{-1} J^\top C_1(\tilde{x}_t^*, x_{n+1}^*) \\ = 0.$$

Taking $\epsilon \rightarrow 0$ concludes the proof. ■

Theorem A.10: The feature update step of the standard MSCKF algorithm (Alg. 11) is equivalent to applying a marginalization step to $c_{\text{MSCKF},t,2} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathbb{R}^{|S_f|d_f} \rightarrow \mathbb{R}$, given by:

$$c_{\text{MSCKF},t,2}(\tilde{x}_t, f_{S_f})$$

Algorithm 11: Multi-State Constrained Kalman Filter, Feature Update Sub-block.

Data: MSCKF state $\tilde{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, with mean $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and covariance $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$, Set of image measurements for marginalization $S_{z,1} \cup S_{z,2}$, Set of features to marginalize S_f , measurement map $h : \mathcal{X}_p \times \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_z}$.

Result: Updated MSCKF state mean

$$\begin{aligned} \bar{\mu}_t &\in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \text{ and covariance} \\ \bar{\Sigma}_t &\in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}. \end{aligned}$$

- 1 $f_{S_f} \in \mathbb{R}^{|S_f|d_f} \leftarrow$ Concatenation of all features in S_f
 - 2 $f_{S_f}^* \in \mathbb{R}^{|S_f|d_f} \leftarrow$ Concatenation of position estimate of all features in S_f
 - 3 $\tilde{h}(\tilde{x}_t, f_{S_f}) \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z} \leftarrow$ Concatenation of measurement map outputs $\{h(x_i, f_j) | (x_i, f_j) \in S_{z,1} \cup S_{z,2}\}$.
 - 4 $\tilde{z} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z} \leftarrow$ Concatenation of feature measurements $\{z_{ij} | (x_i, f_j) \in S_{z,1} \cup S_{z,2}\}$.
 - 5 $\tilde{H}_{t,x} \leftarrow \frac{\partial \tilde{h}}{\partial \tilde{x}_t}(\mu_t, f_{S_f}^*) \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times (d_{\text{IMU}}+nd_x)}$.
 - 6 $\tilde{H}_{t,f} \leftarrow \frac{\partial \tilde{h}}{\partial f_{S_f}}(\mu_t, f_{S_f}^*) \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times |S_f|d_f}$.
 - 7 $\{a_1, \dots, a_{|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f}\} \subset \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z} \leftarrow$ Orthonormal basis for $N(\tilde{H}_{t,f}^\top)$.
 - 8 $A \leftarrow \begin{bmatrix} a_1 & \dots & a_{|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f} \end{bmatrix} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times (|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f)}$.
 - 9 $QT \leftarrow$ QR Decomposition of $A^\top \tilde{H}_{t,x}$, with $Q \in \mathbb{R}^{(|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f) \times (|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f)}$, $T \in \mathbb{R}^{(|S_{z,1} \cup S_{z,2}|d_z - |S_f|d_f) \times (d_{\text{IMU}}+nd_x)}$.
 - 10 $\bar{\Sigma}_t^{-1} \leftarrow \Sigma_t^{-1} + T^\top (Q^\top A^\top R A Q)^{-1} T \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$.
 - 11 $\bar{\mu}_t \leftarrow \mu_t \boxplus (\Sigma_t^{-1} + T^\top (Q^\top A^\top R A Q)^{-1} T)^{-1} T^\top (Q^\top A^\top R A Q)^{-1} (\tilde{z} \boxplus \tilde{h}(\tilde{x}_t)) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$.
 - 12 $\hat{x}_t \leftarrow \bar{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$.
 - 13 **return** $\bar{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, $\bar{\Sigma}_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$
-

$$:= \|\tilde{x}_t \boxplus \mu_t\|_{\Sigma_t^{-1}}^2 + \sum_{(x_i, f_j) \in S_{z,1} \cup S_{z,2}} \|z_{i,j} \boxplus h(x_i, f_j)\|_{\Sigma_v^{-1}}^2,$$

where $f_{S_f} \in \mathbb{R}^{|S_f|d_f}$ denotes the stacked vector of all feature positions in S_f (see Algorithm 9).

Proof: First, we rewrite $c_{\text{MSCKF},t,2}$ as:

$$\begin{aligned} c_{\text{MSCKF},t,2}(\tilde{x}_t, f_{S_f}) \\ := \|\tilde{x}_t \boxplus \mu_t\|_{\Sigma_t^{-1}}^2 + \|\tilde{z} \boxplus \tilde{h}(\tilde{x}_t, f_{S_f})\|_{\Sigma_v^{-1}}^2, \end{aligned}$$

where $\tilde{z} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z}$, $\tilde{h} : \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \times \mathbb{R}^{|S_f|d_f} \rightarrow \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z}$: are defined as follows— \tilde{z} denotes the stacked measurement vectors in $\{z_{i,j} | (x_i, f_j) \in S_{z,1} \cup S_{z,2}\} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z}$, $\tilde{h}(\tilde{x}_t, f_{S_f})$ denotes the stacked outputs of the measurement map in $\{h(x_i, f_j) | (x_i, f_j) \in S_{z,1} \cup S_{z,2}\} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z}$, and $\Sigma_v := \text{diag}\{\Sigma_v, \dots, \Sigma_v\} \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times |S_{z,1} \cup S_{z,2}|d_z}$.

Essentially, by marginalizing the feature position estimates, this step utilizes information from feature measurements to

Algorithm 12: Multi-State Constrained Kalman Filter, State Propagation Sub-block.

Data: MSCKF state $\tilde{x}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, with mean $\mu_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$ and covariance $\Sigma_t \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$, (discrete-time) dynamics map $g : \mathbb{R}^{d_{\text{IMU}}} \rightarrow \mathbb{R}^{d_{\text{IMU}}}$.

Result: Updated MSCKF state mean

$$\begin{aligned} \mu_{t+1} &\in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n \text{ and covariance} \\ \Sigma_{t+1} &\in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}. \end{aligned}$$

- 1 $(\bar{\mu}_{t,\text{IMU}}, \bar{\mu}_{t,x,1:n}) \leftarrow \bar{\mu}_t$, with $\bar{\mu}_{t,\text{IMU}} \in \mathbb{R}^{d_{\text{IMU}}}$, $\bar{\mu}_{t,x,1:n} \in \mathbb{R}^{nd_x}$.
 - 2 $G_t \leftarrow$ Jacobian of $g_{\text{IMU}} : \mathbb{R}^{d_{\text{IMU}}} \rightarrow \mathbb{R}^{d_{\text{IMU}}}$ evaluated at $\bar{\mu}_{t,\text{IMU}} \in \mathbb{R}^{d_{\text{IMU}}}$.
 - 3 $\mu_{t+1} \leftarrow (g_{\text{IMU}}(\bar{\mu}_{t,\text{IMU}}), \bar{\mu}_{t,x,1:n}) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$.
 - 4 $\Sigma_{t+1} \leftarrow \begin{bmatrix} G_t & O \\ O & I_{nd_x} \end{bmatrix} \bar{\Sigma}_t \begin{bmatrix} G_t^\top & O \\ O & I_{nd_x} \end{bmatrix} + \begin{bmatrix} \Sigma_w & O \\ O & O \end{bmatrix} \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$.
 - 5 **return** $\mu_{t+1} \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n$, $\Sigma_{t+1} \in \mathbb{R}^{(d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)}$.
-

constrain our state estimates. To accomplish this, we choose our algorithm variables as follows:

$$\begin{aligned} \tilde{x}_{t,K} &:= \tilde{x}_t = (x_{t,\text{IMU}}, x_1, \dots, x_n) \\ &\in \mathbb{R}^{d_{\text{IMU}}+nd_x+|S_f|d_f}, \end{aligned}$$

$$\tilde{x}_{t,M} := f_{S_f} \in \mathbb{R}^{|S_f|d_f},$$

$$\bar{x} := (\tilde{x}_{t,K}, \tilde{x}_{t,M})$$

$$\in \mathbb{R}^{d_{\text{IMU}}+nd_x+|S_f|d_f},$$

$$\begin{aligned} C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}) &:= \begin{bmatrix} \Sigma_t^{-1/2}(\tilde{x}_t \boxplus \mu_t) \\ \tilde{\Sigma}_v^{-1/2}(\tilde{z} \boxplus \tilde{h}(\tilde{x}_t, f_{S_f})) \end{bmatrix} \\ &\in \mathbb{R}^{d_{\text{IMU}}+nd_x+|S_{z,1} \cup S_{z,2}|d_z}. \end{aligned}$$

The Marginalization algorithm block then implies that:

$$\begin{aligned} J_K &:= \frac{\partial C_M}{\partial \tilde{x}_t}(\bar{\mu}_t, f_{S_f}^*) = \begin{bmatrix} \Sigma_t^{-1/2} \\ -\tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,x} \end{bmatrix} \\ &\in \mathbb{R}^{(d_{\text{IMU}}+nd_x+|S_{z,1} \cup S_{z,2}|d_z) \times (d_{\text{IMU}}+nd_x)}, \end{aligned}$$

$$\begin{aligned} J_M &:= \frac{\partial C_M}{\partial f_{S_f}}(\bar{\mu}_t, f_{S_f}^*) = \begin{bmatrix} O \\ -\tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} \end{bmatrix} \\ &\in \mathbb{R}^{(d_{\text{IMU}}+nd_x+|S_{z,1} \cup S_{z,2}|d_z) \times |S_f|d_f}, \end{aligned}$$

where we have defined:

$f_{S_f}^* \in \mathbb{R}^{|S_f|d_f} \leftarrow$ Stacked position estimates of features in S_f ,

$$\tilde{H}_{t,x} := \frac{\partial \tilde{h}}{\partial \tilde{x}_t}(\bar{\mu}_t, f_{S_f}^*) \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times (d_{\text{IMU}}+nd_x)},$$

$$\tilde{H}_{t,f} := \frac{\partial \tilde{h}}{\partial f_{S_f}}(\bar{\mu}_t, f_{S_f}^*) \in \mathbb{R}^{|S_{z,1} \cup S_{z,2}|d_z \times |S_f|d_f}.$$

Recall that the marginalization equations (10) and (9) in our formulation read:

$$\begin{aligned} \mu_K &\leftarrow \mu_K - \Sigma_K J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}), \\ \Sigma_K &\leftarrow (J_K^\top (I - J_M (J_M^\top J_M)^{-1} J_M^\top) J_K)^{-1}. \end{aligned}$$

Substituting in the above expressions for J_K, J_M , and $C_M(\bar{\mu}_t, f_{S_f}^*)$, we have:

$$\begin{aligned} \bar{\Sigma}_t &\leftarrow (J_K^\top (I - J_M (J_M^\top J_M)^{-1} J_M^\top) J_K)^{-1}, \\ &= \left(\begin{bmatrix} \Sigma_t^{-1/2} & -\tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} \\ O & I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2} \end{bmatrix} \right. \\ &\quad \left. \begin{bmatrix} I & O \\ O & I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2} \end{bmatrix} \right)^{-1} \\ &= (\Sigma_t^{-1} + \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} [I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \\ &\quad \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2}] \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,x})^{-1} \\ \bar{\mu}_t &\leftarrow \mu_K - \Sigma_K J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] C_M(\bar{\mu}_t, f_{S_f}^*) \\ &= \mu_t + (\Sigma_t^{-1} + \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} \\ &\quad [I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2}] \\ &\quad \cdot \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,x})^{-1} \\ &\quad \cdot \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} [I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \\ &\quad \cdot \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2}] \tilde{\Sigma}_v^{-1/2} (\tilde{z} - \tilde{h}(\tilde{x}_t, f_{S_f})). \end{aligned}$$

Comparing with the update step in the MSCKF algorithm, i.e., (10) and (9), reproduced below:

$$\begin{aligned} \bar{\Sigma}_t^{-1} &\leftarrow \Sigma_t^{-1} + T^\top (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} T, \\ \bar{\mu}_t &\leftarrow \mu_t + (\Sigma_t^{-1} + T^\top (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} T)^{-1} \\ &\quad T^\top (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} (\tilde{z} - \tilde{h}(\tilde{x}_t, f_{S_f})) \end{aligned}$$

we find that it suffices to show:

$$\begin{aligned} &T^\top (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} \\ &= \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} [I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2}] \\ &\quad \cdot \tilde{\Sigma}_v^{-1/2} \\ &= \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} - \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1}. \end{aligned}$$

To see this, recall that A is defined as a full-rank matrix whose columns span $N(\tilde{H}_{t,f}^\top)$. Thus:

$$(\tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f})^\top \cdot \tilde{\Sigma}_v^{1/2} A Q = \tilde{H}_{t,f}^\top A Q = O.$$

In other words, the columns of $\tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f}$ and of $\tilde{\Sigma}_v^{1/2} A Q$ form bases of orthogonal subspaces whose direct sum equals $\mathbb{R}^{nq d_z}$. We thus have:

$$\begin{aligned} &\tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2} \\ &\quad + \tilde{\Sigma}_v^{1/2} A Q (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} Q^\top A^\top \tilde{\Sigma}_v^{1/2} = I, \end{aligned}$$

which in turn implies that:

$$\begin{aligned} &T^\top (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} \\ &= \tilde{H}_{t,x}^\top A Q (Q^\top A^\top \tilde{\Sigma}_v A Q)^{-1} Q^\top A^\top \\ &= \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} (\tilde{\Sigma}_v^{1/2} A Q) \\ &\quad (Q^\top A^\top \tilde{\Sigma}_v^{1/2} \cdot \tilde{\Sigma}_v^{1/2} A Q)^{-1} (Q^\top A^\top \tilde{\Sigma}_v^{1/2}) \tilde{\Sigma}_v^{-1/2} \\ &= \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1/2} (I - \tilde{\Sigma}_v^{-1/2} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1/2}) \\ &\quad \tilde{\Sigma}_v^{-1/2} \end{aligned}$$

$$= \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,x} - \tilde{H}_{t,x}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f} (\tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1} \tilde{H}_{t,f})^{-1} \tilde{H}_{t,f}^\top \tilde{\Sigma}_v^{-1},$$

as claimed. \blacksquare

Theorem A.11: The state propagation step of the standard MSCKF SLAM algorithm (Alg. 12) is equivalent to applying a Marginalization step once to $c_{MSCKF,t,4} : \mathbb{R}^{2d_{IMU} + nd_x} \rightarrow \mathbb{R}$, given by:

$$\begin{aligned} &c_{MSCKF,t,4}(\tilde{x}_t, x_{t+1,IMU}) \\ &:= \|\tilde{x}_t \boxminus \bar{\mu}_t\|_{\Sigma_t^{-1}}^2 + \|x_{t+1,IMU} \boxminus g_{IMU}(x_{t,IMU})\|_{\Sigma_t^{-1}}^2. \end{aligned}$$

Proof: We claim that from an optimization perspective, the update step is equivalent to applying one marginalization step to the cost function $c_{MSCKF,t,4}(\tilde{x}_t, x_{t+1,IMU})$ specified above. In particular, we wish to marginalize out $x_{t,IMU} \in \mathcal{X}_{IMU}$ and retain $x_{t+1,IMU} \in \mathcal{X}_{IMU}$; in other words, in the notation of our Marginalization algorithm submodule, we have:

$$\begin{aligned} \tilde{x}_{t,K} &:= (x_{t+1,IMU}, x_1, \dots, x_n) \in \mathcal{X}_{IMU} \times (\mathcal{X}_p)^n, \\ \tilde{x}_{t,M} &:= x_{t,IMU} \in \mathcal{X}_{IMU}. \end{aligned}$$

To apply a marginalization step, our first task is to find vectors $C_K(x_K) = C_K(\tilde{x}_t)$ and $C_M(x_K, x_M) = C_M(\tilde{x}_t, x_{t+1,IMU})$ of appropriate dimensions such that $c_{MSCKF,t,4}(\tilde{x}_t, x_{t+1,IMU}) = C_K(x_{t+1,IMU})^\top C_K(x_{t+1,IMU}) + C_M(\tilde{x}_t, x_{t+1,IMU})^\top C_M(\tilde{x}_t, x_{t+1,IMU})$. A natural choice is furnished by $C_K(x_{t+1,IMU}) \in \mathbb{R}$ and $C_M(\tilde{x}_t, x_{t+1,IMU}) \in \mathcal{X}_p$, as defined below:

$$C_K(\tilde{x}_{t,K}) = 0 \in \mathbb{R}$$

$$C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}) = \begin{bmatrix} \tilde{\Sigma}_t^{-1/2} (\tilde{x}_t - \bar{\mu}_t) \\ \Sigma_w^{-1/2} (x_{t+1,IMU} - g_{IMU}(x_{t,IMU})) \end{bmatrix} \in \mathbb{R}^{2d_{IMU} + nd_x}.$$

For convenience, we will define the IMU state and pose components of the mean $\mu_t \in \mathcal{X}_{IMU} \times (\mathcal{X}_p)^n$ by $\mu_t := (\mu_{t,IMU}, \mu_{t,IMU}) \in \mathcal{X}_{IMU} \times (\mathcal{X}_p)^n$, with $\mu_{t,IMU} \in \mathcal{X}_p$ and $\mu_{t,x} \in (\mathcal{X}_p)^n$, respectively. This mirrors our definition of $x_t \in \mathcal{X}_p$ and $x_{n+1} \in (\mathcal{X}_p)^n$ as the components of the full state $\tilde{x}_t := (x_t, x_{n+1}) \in \mathcal{X}_{IMU} \times (\mathcal{X}_p)^n$. In addition, we will define the components of $\tilde{\Sigma}_t^{-1/2} \in \mathbb{R}^{(d_{IMU} + nd_x) \times (d_{IMU} + nd_x)}$ and $\tilde{\Sigma}_t^{-1} \in \mathbb{R}^{(d_{IMU} + nd_x) \times (d_{IMU} + nd_x)}$ by:

$$\begin{aligned} &\begin{bmatrix} \Omega_{t,IMU,IMU} & \Omega_{t,IMU,x} \\ \Omega_{t,x,IMU} & \Omega_{t,x,x} \end{bmatrix} := \tilde{\Sigma}_t^{-1} \\ &\quad \in \mathbb{R}^{(d_{IMU} + nd_x) \times (d_{IMU} + nd_x)}, \\ &\begin{bmatrix} \Lambda_{t,IMU,IMU} & \Lambda_{t,IMU,x} \\ \Lambda_{t,x,IMU} & \Lambda_{t,x,x} \end{bmatrix} := \tilde{\Sigma}_t^{-1/2} \\ &\quad \in \mathbb{R}^{(d_{IMU} + nd_x) \times (d_{IMU} + nd_x)}, \end{aligned}$$

with the dimensions of the above block matrices given by $\Sigma_{t,IMU,IMU}, \Lambda_{t,IMU,IMU} \in \mathbb{R}^{d_{IMU} \times d_{IMU}}$, $\Sigma_{t,IMU,x}, \Lambda_{t,IMU,x} \in \mathbb{R}^{d_{IMU} \times nd_x}$, $\Sigma_{t,x,IMU}, \Lambda_{t,x,IMU} \in \mathbb{R}^{nd_x \times d_{IMU}}$, and $\Sigma_{t,x,x}, \Lambda_{t,x,x} \in \mathbb{R}^{nd_x \times nd_x}$. Using the above definitions, we can reorder the residuals in $C_K \in \mathbb{R}$ and $C_M \in \mathbb{R}^{2d_{IMU} + nd_x}$, and thus redefine them by:

$$C_K(\tilde{x}_{t,K}) = 0 \in \mathbb{R}$$

$$\begin{aligned}
& C_M(\tilde{x}_{t,K}, \tilde{x}_{t,M}) \\
&= \begin{bmatrix} \Lambda_{t,\text{IMU},\text{IMU}}(x_{t,\text{IMU}} - \mu_{t,\text{IMU}}) + \Lambda_{t,\text{IMU},x}(x_{1:n} - \mu_{t,x}) \\ \Sigma_w^{-1/2}(x_{t+1,\text{IMU}} - g_{\text{IMU}}(x_{t,\text{IMU}})) \\ \Lambda_{t,x,\text{IMU}}(x_{t,\text{IMU}} - \mu_{t,\text{IMU}}) + \Lambda_{t,x,x}(x_{1:n} - \mu_{t,x}) \end{bmatrix} \\
&\in \mathbb{R}^{2d_{\text{IMU}}+nd_x},
\end{aligned}$$

where $x_{1:n} := (x_1, \dots, x_n) \in (\mathcal{X}_p)^n$.

Our state variables and cost functions for the Gauss-Newton algorithm submodule are:

$$\begin{aligned}
\overline{x_M^*} &= \tilde{x}_t^* = \overline{\mu}_t \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n, \\
\overline{x_K^*} &= g(\tilde{x}_t^*) = g(\overline{\mu}_t) \in \mathcal{X}_{\text{IMU}} \times (\mathcal{X}_p)^n, \\
C_K(\tilde{x}_{t,K}^*) &= 0 \in \mathbb{R}, \\
C_M(\tilde{x}_{t,K}^*, \tilde{x}_{t,M}^*) &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \mathbb{R}^{2d_{\text{IMU}}+nd_x}, \\
J_K &= \begin{bmatrix} O & \Lambda_{\text{IMU},x} \\ \Sigma_w^{-1/2} & O \\ O & \Lambda_{xx} \end{bmatrix} \\
&\in \mathbb{R}^{(2d_{\text{IMU}}+nd_x) \times (d_{\text{IMU}}+nd_x)} \\
J_M &= \begin{bmatrix} \Lambda_{\text{IMU},\text{IMU}} \\ -\Sigma_w^{-1/2} G_t \\ \Lambda_{x,\text{IMU}} \end{bmatrix} \in \mathbb{R}^{(2d_{\text{IMU}}+nd_x) \times d_x},
\end{aligned}$$

where we have defined G_t to be the Jacobian of $g_{\text{IMU}} : \mathcal{X}_{\text{IMU}} \rightarrow \mathcal{X}_{\text{IMU}}$ at $\overline{\mu}_{t,\text{IMU}} \in \mathcal{X}_{\text{IMU}}$, i.e.:

$$G_t := \left. \frac{\partial g}{\partial x_{t,\text{IMU}}} \right|_{x_{t,\text{IMU}} = \overline{\mu}_{t,\text{IMU}}}$$

Applying the Marginalization update equations, we thus have:

$$\begin{aligned}
\mu_{t+1} &\leftarrow \tilde{x}_{t,K} - \Sigma_{t+1} J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] \\
&C_M(\overline{x_K^*}, \overline{x_M^*}) \\
&= g(\overline{\mu}_t), \\
\Sigma_{t+1} &\leftarrow (J_K^\top [I - J_M (J_M^\top J_M)^{-1} J_M^\top] J_K)^{-1}, \\
&= (J_K^\top J_K - J_K^\top J_M (J_M^\top J_M)^{-1} J_M^\top J_K)^{-1}, \\
&= \left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Lambda_{x,\text{IMU}} \Lambda_{\text{IMU},x} + \Lambda_{xx}^2 \end{bmatrix} \right. \\
&\quad \left. - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Lambda_{x,\text{IMU}} \Lambda_{\text{IMU},\text{IMU}} + \Lambda_{xx} \Lambda_{x,\text{IMU}} \end{bmatrix} \right. \\
&\quad \cdot (\Lambda_{\text{IMU},\text{IMU}}^2 + \Lambda_{\text{IMU},x} \Lambda_{x,\text{IMU}} + G_t^\top \Sigma_w^{-1} G_t)^{-1} \\
&\quad \left. \cdot [-G_t^\top \Sigma_w^{-1} \quad \Lambda_{\text{IMU},\text{IMU}} \Lambda_{\text{IMU},x} + \Lambda_{x,\text{IMU}} \Lambda_{xx}] \right)^{-1} \\
&= \left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Omega_{xx} \end{bmatrix} - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Omega_{x,\text{IMU}} \end{bmatrix} \right. \\
&\quad \left. (\Omega_{\text{IMU},\text{IMU}} + G_t^\top \Sigma_w^{-1} G_t)^{-1} \right. \\
&\quad \left. [-G_t^\top \Sigma_w^{-1} \quad \Omega_{\text{IMU},x}] \right)^{-1}
\end{aligned}$$

To show that this is indeed identical to the propagation equation for the covariance matrix in the Extended Kalman

Filter algorithm, i.e. Algorithm 5, Line 5, we must show that:

$$\begin{aligned}
& \left(\begin{bmatrix} \Sigma_w^{-1} & O \\ O & \Omega_{xx} \end{bmatrix} - \begin{bmatrix} -\Sigma_w^{-1} G_t \\ \Omega_{x,\text{IMU}} \end{bmatrix} \right. \\
&\quad \left. (\Omega_{\text{IMU},\text{IMU}} + G_t^\top \Sigma_w^{-1} G_t)^{-1} [-G_t^\top \Sigma_w^{-1} \quad \Omega_{\text{IMU},x}] \right)^{-1} \\
&= \begin{bmatrix} G_t \overline{\Sigma}_{t,\text{IMU},\text{IMU}} G_t^\top + \Sigma_w & G_t \overline{\Sigma}_{t,\text{IMU},x} \\ \overline{\Sigma}_{t,\text{IMU},x} G_t^\top & \overline{\Sigma}_{t,x,x} \end{bmatrix}
\end{aligned}$$

This follows by brute-force expanding the above block matrix components, and applying Woodbury's Matrix Identity, along with the definitions of $\Sigma_{t,\text{IMU},\text{IMU}}$, $\Lambda_{t,\text{IMU},\text{IMU}}$, $\Sigma_{t,\text{IMU},x}$, $\Lambda_{t,\text{IMU},x}$, $\Sigma_{t,x,\text{IMU}}$, $\Lambda_{t,x,\text{IMU}}$, $\Sigma_{t,x,x}$, and $\Lambda_{t,x,x}$. ■

D. Simulation Settings

1) Dataset

All experiments are performed on the EuRoC MAV dataset [27], a popular public SLAM dataset of stereo image sequences and inertial measurement unit (IMU) measurements. The stereo images and IMU measurements arrive at rates of 20Hz and 200Hz, respectively. The sensor suite is mounted on board a micro aerial vehicle. Ground-truth poses, recorded using a Vicon motion capture system and IMU biases, are also available for each sequence. Our experiments are carried out on the Vicon Room 2 01 and 02 sequences, which contain about 2300 stereo images each and span about 2 minutes of real-time operation. The first of these sequences is easier to analyze via our SLAM algorithms, as it corresponds to simple and slow evolution of the camera, while the latter contains some jerky and quick motions that prove challenging to some algorithms.

2) Front-end and Back-end

Since the focus of this work is the SLAM back-end, we standardize the front-end across all experiments, altering only the back-end used to process the abstracted data produced by the front-end. We use keypoint features as environment landmarks, as is standard in visual SLAM. First, BRISK features are extracted from both images of the input stereo pair. Then, feature matching is carried out between the left and right image frames. We then carry out brute force matching using Hamming distance on the binary BRISK descriptors, and filter outliers via an epipolar constraint check, using the known relative pose between the two cameras in the stereo set-up. Only keypoints for which a stereo match was found are kept. Next, a four-way consistency check is carried out. i.e. a match between two stereo frames S_1, S_2 is accepted if and only if both observations of a given feature in S_1 are matched to the respective observations of the same feature in S_2 . Finally, outlier matches are rejected by projecting the best estimate of the matched feature onto the best estimate of the current camera pose, and rejecting matches that have a high reprojection error. Any stereo matches in the current frame that were not matched with a previously seen is recorded as a newly detected landmark, and initialized using stereo triangulation from the best estimate of the current camera

pose. The front-end maintains data structures allowing two-way access between features and camera poses: for each feature index, it is possible to look up all camera poses from which that feature is visible, and likewise for each camera pose it is possible to query which features are visible in that frame.

For book-keeping the cost function in the back-end, computing Jacobians, and implementing Gauss-Newton optimization, we use GTSAM in C++ [8, 28].

3) Dynamics and Image Measurement Models

We use an on-board IMU to collect odometry measurements, i.e., body-frame angular velocity and linear acceleration, and apply the IMU pre-integration scheme detailed in [29], as summarized below. The objective of IMU preintegration is to establish a discrete-time dynamics map $x_{t+1} = g(x_t)$ that allows us to predict the pose of the robot at time x_{t+1} given the pose at time t and the IMU measurement at time t . Here, each timestep corresponds to a new *image* measurement. Since IMU measurements arrive at a faster rate than image measurements, to compute this map, we must stack several IMU measurements into a relative state measurement, which can then be concatenated with the state at time x_t to get the predicted state at time x_{t+1} .

The robot state x_t consists of the orientation, position, and velocity of the body frame relative to the world frame, and the IMU biases, so that $x_t = (R_t, p_t, v_t, b_t)$, where $(R_t, p_t) \in SE(3)$, $v_t \in \mathbb{R}^3$ and $b_t = (b_t^g, b_t^a) \in \mathbb{R}^3 \times \mathbb{R}^3 \simeq \mathbb{R}^6$ are the IMU biases in the gyroscope and accelerometer respectively. The IMU biases are slowly varying and generally unknown, so they are included in the robot state and are also jointly estimated.

The IMU measures angular velocity and accelerations. The measurements are denoted ${}_B\tilde{a}$ and ${}_B\tilde{\omega}_{WB}$. Here, B refers to the robot’s body frame and W the world frame. The prefix B means that the quantity is expressed in the B frame, and the suffix WB denotes that the quantity represents the motion of the B frame relative to W . So the pose of the robot is $(R_{WB}, Wp) \in SE(3)$. The measurements are affected by additive white noise η and the slowly varying IMU biases:

$$\begin{aligned} {}_B\tilde{\omega}_{WB}(t) &= {}_B\omega_{WB}(t) + b^\omega(t) + \eta^\omega(t) \\ {}_B\tilde{a}(t) &= R_{WB}^\top(t)({}_W a(t) - {}_W g) + b^a(t) + \eta^a(t) \end{aligned}$$

where ${}_W\omega_{WB}$ is the true angular velocity, ${}_W a$ the true acceleration, and ${}_W g$ the gravity acceleration vector in the world frame. Assume that between timestep t and $t + 1$, we received m IMU measurements, at constant time increments Δt . [29] provides expressions for $(x_{t+1} \boxminus g(x_t))$ directly in terms of the IMU measurements. Additionally, expressions for noise propagation are also provided, which allows us to compute the covariance Σ_v over the error $(x_{t+1} \boxminus g(x_t))$ in terms of the measurement noise η , which is what we need to compute the required cost function $\|x_{t+1} \boxminus g(x_t)\|_{\Sigma_v}^2$ and Jacobians.

Given a camera pose $x_t = (R_t, p_t)$ and a 3D feature location $f_j = (f_j^x, f_j^y, f_j^z)$, the camera measurement model $h(x_t, f_j)$ predicts the projected pixel location of the point in both stereo images. We assume the stereo camera pair is calibrated and

rectified, so that both cameras have the same pinhole camera matrix K , and the epilines are horizontal. As such, the two measurements of a point in the two cameras will share the same v -coordinate in (u, v) image space. Therefore, an image measurement will be stored as a 3-vector (u_L, u_R, v) , where the coordinates of the measurement in the left and right image are (u_L, v) and (u_R, v) respectively. The measurement map h predicts the image location (u_L, u_R, v) by projecting the point f_j onto both image frames (with the standard pinhole projection), using the known poses of the two cameras in the robot’s body-frame. Due to noise, the actual measurement will not have the exact same v coordinate, so the image measurement is collected by averaging the two v -coordinates of the two keypoints. This measurement z_{tj} is then compared to the predicted coordinates by a Mahalanobis distance in \mathbb{R}^3 space to get the cost function $\|z_{tj} - h(x_t, f_j)\|_{\Sigma_w}^2$. The covariance Σ_w is the expected noise in image space, which is a design parameter. In our experiments we choose $\Sigma_w = \sigma I_3$ where I_3 is the 3×3 identity matrix and $\sigma = 0.05$ pixels, which was found to work well in practice.