Robust Preintegrated Wheel Odometry for Off-road Autonomous Ground Vehicles

Easton R. Potokar, Daniel McGann, and Michael Kaess

***Abstract*—Wheel odometry is not often used in state estimation for off-road vehicles due to frequent wheel slippage, varying wheel radii, and the 3D motion of the vehicle not fitting with the 2D nature of integrated wheel odometry. This paper attempts to overcome these issues by proposing a novel 3D preintegration of wheel encoder measurements on manifold. Our method additionally estimates wheel slip, radii, and baseline online to improve accuracy and robustness. Further, due to the preintegration, many measurements can be summarized into a single motion constraint using first-order updates for wheel slippage and intrinsics, allowing for efficient usage in an optimization-based state estimation framework. While our method can be used with any sensors in a factor graph framework, we validate its effectiveness and observability of parameters in a vision-wheel-odometry system (VWO) in a Monte Carlo simulation. Additionally, we illustrate its accuracy and demonstrate it can be used to overcome other sensor failures in real-world off-road scenarios in both a VWO and visual-inertial-wheel odometry (VIWO) system.**

***Index Terms*—Localization, Wheeled Robots, Field Robots**

I. Introduction

OFF-ROAD autonomous ground vehicles (AGVs) have a variety of applications that can have significant impacts in industries such as autonomous driving, agriculture, and space exploration. The use of AGVs in these industries has the potential to greatly increase quality of life, safety, and productivity in their various fields. However, effective usage of AGVs requires accurate 3D localization. For some of these applications GPS is often not available, leading to state estimation using sensors such as cameras, LiDARs, inertial measurement units (IMUs), and wheel encoders.

While fusion of IMUs in odometry frameworks has been well studied [1], [2] in off-road vehicles, fusion with wheel encoders has mainly been limited to planar environments such as urban driving or indoors. Wheel encoders have seen more limited use for off-road vehicle state estimation due to a number of inherit limitations. One such obstacle is that integration of wheel encoder measurements is inherently 2D. Because of this, previous methods have provided either planar or no constraints on the additional roll, pitch, and z pose dimensions.

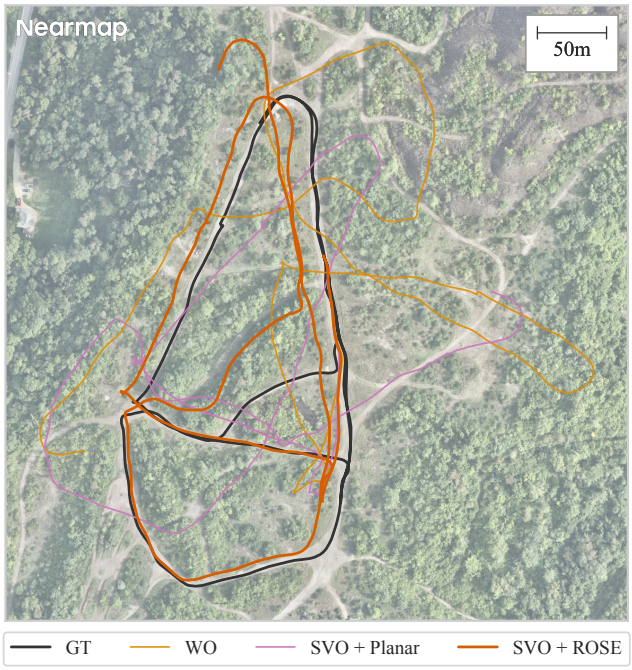

Fig. 1: Trajectory of results from a real world trial in Penn Hills, Pennsylvania. Where other baseline methods break down due to non-planar surfaces and wheel slippage, our method is able to minimize drift and output more accurate pose estimates. Shown above are the ground truth (GT), wheel odometry (WO), stereo visual odometry with planar wheel odometry (SVO + Planar), and stereo visual odometry with our method (SVO + ROSE).

Another common pitfall is outlier wheel slippage, which occurs more often in an off-road scenario. While some wheel slippage is expected due to the dynamics of AGVs [3], a single outlier encoder measurement from wheel slippage can cause significant degradation of vehicle tracking and must be detected in real-time. Additionally, many parameters such as wheel radii can vary in real time due to dropping tire pressure or rugged terrain. While the effect of incorrect wheel radii parameters is more subtle than that of wheel slippage, it can still cause meaningful increase in drift over time.

Further, wheel encoder measurements generally come in at a much higher rate than sensors such as cameras. When estimating other parameters such as wheel intrinsics or slip, this can lead to repeated integration of wheel encoder measurements during optimization, a significant computational burden. Some work has been done on preintegration theory to avoid the repeated integration of wheel encoder measurements, although it is generally on SE(2) and not suitable for off-road vehicles.

Manuscript received: July, 26, 2024; Revised October, 7, 2024; Accepted October, 31, 2024.

This paper was recommended for publication by Editor Giuseppe Loianno upon evaluation of the Associate Editor and Reviewers' comments.

This material is based upon work supported by the U.S. Army Research Office and the U.S. Army Futures Command under Contract No. W911NF20-D-0002. The content of the information does not reflect the position or the policy of the government and no official endorsement should be inferred.

Easton R. Potokar, Daniel McGann, and Michael Kaess are with the Robotics Institute at Carnegie Mellon University, Pittsburgh, PA USA {potokar, danmcgann, kaess}@cmu.edu

Digital Object Identifier (DOI): see top of this page.

We propose a Robust Offroad wheel odometry with Slip Estimation (ROSE) to solve these problems. ROSE is a novel way to incorporate wheel encoder measurements into optimization-based state estimation under the factor graph framework. Our approach seeks to address the above issues when using wheel encoders for off-road state estimation. In particular, our contributions are as follows:

1) By leveraging nonholonomic constraints and constraining the x and y angular velocity of the vehicle, we show the on-manifold SE(3) integration of wheel encoder measurements using local coordinates and how it can be used in a preintegrated fashion to form a piecewise-planar constraint.
2) We show how wheel slip, wheel radii, and the vehicle wheel baseline can be estimated online without repeated integration for more efficient optimization-based estimation.
3) We evaluate our proposed method via a Monte Carlo simulation in a VWO framework, showing convergence of wheel parameters and slippage as well as robustness to both significant slippage and incorrect initial parameters. Moreover, we validate on real-world data in both a VWO and VIWO framework.

Finally, we release our codebase as open-source[1] for widespread usage in the state estimation library gtsam [4].

II. Related Works

Most previous work done on integrating wheel encoders into a sensor fusion system doesn't consider the specific limitations of off-road vehicles. Previous frameworks are often underconstrained if not fused with an IMU, are focused on planar environments such as indoors, or disregard wheel slippage and intrinsics.

While IMUs are generally used to provide an odometry "backbone" for state estimation methods, there are scenarios when IMU failures may occur due to poor bias estimation, time-sync, or even hardware malfunction amongst others issues. In these scenarios, it would be ideal if the state estimation method could lean upon wheel odometry as an alternative odometry backbone. Some work isn't usable in this way due to tightly-fusing the IMU gyroscope with wheel encoder measurements [5], [6], [7]. Others use solely the IMU for angular measurements [8], [9], leaving no orientation measurements in cases of IMU failure. The majority of other work on VIWO methods are missing constraints in the pitch, roll, and z positions [10], [11] when the IMU isn't functioning, potentially leading to underconstrained systems with erroneous estimates.

Other works do constrain these additional directions, but do so by constraining some or all of them to zero [12], [13], [14]. While this is a valid assumption for planar environments such as indoors or some urban environments, the same does not hold true for off-road vehicles. Some work has been done on fitting complex surface manifolds to provide integration in these additional directions either implicitly [15], [16] or explicitly using cameras [17], but is unlikely to function off-road due to potentially rocky and non-smooth terrain.

While much work has been done on modeling wheel slippage for dynamics and controls [18], [19], [20], estimating outlier wheel slippage events and wheel intrinsics for state estimation hasn't been as widely studied. Some work has been done using factor graphs similar to our own [11], but is formulated for planar environments and removes the wheel odometry in cases of slippage. Others have used Kalman Filters to estimate slip parameters, but rely upon the IMU for heading measurements [21] or don't consider the impact slip may have on vehicle heading [8].

There's been significant work done in calibrating wheel intrinsics and extrinsics offline [22], [23], [24]. Online estimation of wheel intrinsics has been done when fused with an IMU in a multi-state constrained Kalman Filter (MSCKF) [10] or fused with LiDAR in a factor graph [25]. While these function for their corresponding applications, they don't adhere to the requirements of off-road driving mentioned above.

To fill the gap left by prior works in off-road driving, we propose our 6D wheel preintegrated method with online wheel slip and intrinsics estimation.

III. Mathematical Background

We briefly introduce some mathematical preliminaries that we will use throughout the paper. For a ground truth measurement z, we will use $\hat{z}$ to represent an estimate of z, and $\tilde{z}$ to represent z with noise throughout.

A. Lie Groups

We denote a Lie group by $\mathcal{G}$ and its corresponding Lie algebra by $\mathfrak{g}$. Examples include the group of rotations $\mathrm{SO}(3)$ along with $\mathfrak{so}(3)$, or, used throughout this manuscript, the group of 3D rigid body transformations $\mathrm{SE}(3)$ along with $\mathfrak{se}(3)$.

The group itself is generally highly nonconvex, making tasks such as optimization, estimation, and integration difficult. The algebra is a vector space and is isomorphic to $\mathbb{R}^{\dim \mathfrak{g}}$ using the linear "wedge" operator,

$$^{\wedge} : \mathbb{R}^{\dim \mathfrak{g}} \to \mathfrak{g}. \tag{1}$$

An inverse map also exists, given by the "vee" operator,

$$^{\vee} : \mathfrak{g} \to \mathbb{R}^{\dim \mathfrak{g}}. \tag{2}$$

Since the algebra is a vector space, it is often used in tasks such as optimization and integration. Finally, for elements near the identity of $\mathcal{G}$, the exponential map $\exp : \mathfrak{g} \to \mathcal{G}$ can be used to map the algebra to the group. It is often paired with the wedge operator,

$$\exp(\cdot^{\wedge}) : \mathbb{R}^{\dim \mathfrak{g}} \to \mathcal{G}. \tag{3}$$

Finally, we define the exponential map Jacobian [26] as

$$H(\theta) \triangleq \frac{\partial}{\partial \theta} \exp(\theta^{\wedge}). \tag{4}$$

It can be used as a first order approximation when separating matrix exponential components,

$$\exp(\theta^{\wedge} + \delta^{\wedge}) \approx \exp(\theta^{\wedge}) \exp((H(\theta)\delta)^{\wedge}). \tag{5}$$

We point the reader to other literature [26] for more details.

[1] Available at https://github.com/rpl-cmu/rose/

B. Local Coordinates

Later on, we will perform preintegration in local manifold coordinates [27], [28] that we introduce here. Consider a differential equation given by $X \in \mathrm{SE}(3)$ and $\dot{X} \in \mathfrak{se}(3)$,

$$\dot{X}(t) = F(X,t), \quad X(0) = X_0. \tag{6}$$

Ideally, we would transfer this differential equation to $\mathbb{R}^6$ to provide simple preintegration in a vector space. We assume a solution of the form

$$X(t) = X_0 \exp(\theta(t)^\wedge) \tag{7}$$

where $\theta(t) \in \mathbb{R}^6$. We can differentiate X with respect to time by using the first-order approximations $\theta(t+h) \approx \theta(t) + \dot{\theta}(t)h$ and $\exp(\theta^\wedge) \approx I + \theta^\wedge$ which hold under $h \to 0$,

$$\begin{aligned}
\dot{X}(t) &= \lim_{h\to 0} \frac{X_0 \exp(\theta(t+h)^\wedge) - X_0 \exp(\theta(t)^\wedge)}{h} \\
&\approx \lim_{h\to 0} \frac{X_0 \exp\Big((\theta(t) + \dot{\theta}(t)h)^\wedge\Big) - X_0 \exp(\theta(t)^\wedge)}{h} \\
&\approx \lim_{h\to 0} X_0 \exp(\theta(t)^\wedge) \frac{I + (H(\theta)\dot{\theta}(t)h)^\wedge - I}{h} \\
&= X_0 \exp(\theta(t)^\wedge)(H(\theta)\dot{\theta}(t))^\wedge \\
&= X(t)(H(\theta)\dot{\theta}(t))^\wedge.
\end{aligned} \tag{8}$$

Equating this result with Eq. 6 results in,

$$\begin{aligned}
&F(X,t) = \dot{X}(t) = X(t)(H(\theta)\dot{\theta}(t))^\wedge \\
\implies &\dot{\theta}(t) = H(\theta)^{-1}(X(t)^{-1}F(X,t))^\vee, \quad \theta(0) = 0.
\end{aligned} \tag{9}$$

This result gives a straightforward way to perform manifold integration in $\mathbb{R}^6$ as long as our assumption in Eq. 7 is valid, which should hold true over small timesteps. Later it will be seen how this result will continue to be simplified for wheel odometry.

C. Factor Graphs

In state estimation, the objective is often to optimize a maximum a posteriori (MAP) problem over variables Θ given measurements Z, given by

$$\Theta^* = \mathrm{argmax}_\Theta P(\Theta|Z) \propto \mathrm{argmax}_\Theta P(Z|\Theta)P(\Theta). \tag{10}$$

When the measurements are assumed to be Gaussian distributed, this simplifies to nonlinear least squares optimization,

$$\Theta^* = \mathrm{argmin}_\Theta \sum_{z\in Z} ||h_z(\Theta) - z||^2_{\Sigma_z} \tag{11}$$

where h_z is the prediction function of z given the estimated state Θ. The difference $h_z(\Theta) - z$ is often called the residual and written as

$$\Theta^* = \mathrm{argmin}_\Theta \sum_{z\in Z} ||r_z(\Theta)||^2_{\Sigma_z}. \tag{12}$$

This can be represented by a factor graph [29], which is a bipartite graph with a single factor node for each residual. The factors nodes are connected to variable nodes given their dependence on each variable. Factor graphs are a useful visualization and computational tool for state estimation.

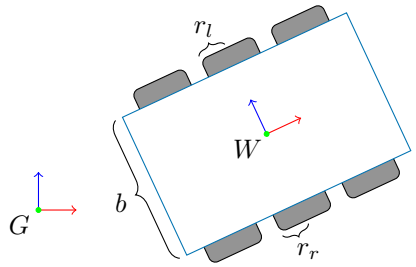

Fig. 2: Small example of vehicle where r_l and r_r are the left and right wheel radii, b the baseline between wheels, G the global frame, and W the wheel encoder frame.

IV. Preintegrated Wheel Odometry

In this section we will introduce the novel preintegrated wheel odometry by first walking through the 6D preintegration in local coordinates, then the first order updates to both wheel slip and intrinsics. We assume a single pair of wheels with encoders as seen in Fig. 2, but our method can be generalized to multiple pairs of wheels. Additionally, our use cases involve a skid-steer vehicle, but again our method generalizes to other vehicle types as well [16].

A. Measurement Transformation

Generally, wheel encoders measure the true angular speed $\omega = \begin{bmatrix}\omega_l & \omega_r\end{bmatrix}^\top$ of each wheel, or a value proportionate to it. Often, wheels will skid or slip on the ground, causing measured wheel speeds to deviate from the expected wheel speeds given the vehicle velocity. We introduce these slip events, along with Gaussian noise, as

$$\tilde{\omega} = \omega + s + \mathrm{w}_\omega, \qquad \mathrm{w}_\omega \sim \mathcal{N}(0, \Sigma_\omega) \tag{13}$$

where the wheel slip is represented by $s = \begin{bmatrix}s_l & s_r\end{bmatrix}^\top$. Note this wheel slip model is significantly simpler than others used for dynamics and control [3], [19], but is sufficiently accurate to capture outlier wheel slippage for state estimation [11]. Using the wheel baseline b and the left and right wheel radii r_l, r_r as defined Fig. 2, these measurements can be transformed to a vehicle forward linear velocity v_x and z angular velocity ω_z as follows,

$$\begin{bmatrix}\tilde{\omega_z} \\ \tilde{v_x}\end{bmatrix} = \begin{bmatrix}\frac{r_r\tilde{\omega}_r - r_l\tilde{\omega}_l}{b} \\ \frac{r_r\tilde{\omega}_r + r_l\tilde{\omega}_l}{2}\end{bmatrix} = \begin{bmatrix}\frac{-r_l}{b} & \frac{r_r}{b} \\ \frac{r_l}{2} & \frac{r_r}{2}\end{bmatrix}\tilde{\omega} \triangleq M(n)\tilde{\omega} \tag{14}$$

where we have defined the wheel intrinsics $n = \begin{bmatrix}b & r_l & r_r\end{bmatrix}^\top$ and $M(n)$ as the linear transformation that maps measurements to velocities. As this transformation is linear, we can exactly propagate the Gaussian noise through $M(n)$ to arrive at

$$\begin{aligned}
\begin{bmatrix}\tilde{\omega_z} \\ \tilde{v_x}\end{bmatrix} &= M(n)(\omega + s + \mathrm{w}_\omega) \\
&= M(n)(\omega + s) + \mathrm{w}_z \triangleq \begin{bmatrix}\omega_z \\ v_x\end{bmatrix} + \mathrm{w}_z, \\
\mathrm{w}_z &\sim \mathcal{N}(0, \Sigma_z), \quad \Sigma_z = M(n)\Sigma_\omega M(n)^\top.
\end{aligned} \tag{15}$$

To perform full 6 degrees of freedom (DoF) integration, additional constraints are needed for the y, z linear velocities

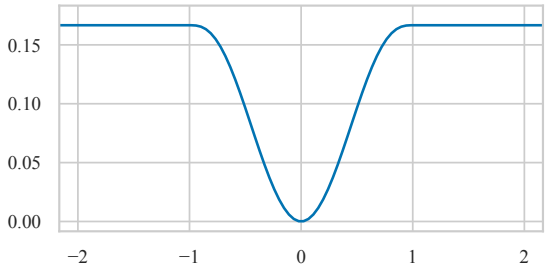

Fig. 3: Tukey-biweight robust kernel function, with parameter $c = 1$. Notice that outside of the basin, the function has gradient 0.

and x, y angular velocities. By leveraging the nonholonomic constraint, we assume that the y and z linear velocities can be represented by mean-zero Gaussians. Similarly over long time periods, we assume ground vehicles will have mean-zero x, y angular velocities and model them using mean-zero Gaussians as well. This results in full velocity measurement of the form,

$$\begin{aligned} \tilde{v} &= \Pi \begin{bmatrix} \omega_z \\ v_x \end{bmatrix} + \mathrm{w}, \qquad \mathrm{w} \sim \mathcal{N}(0_6, \Sigma_v) \\ \Sigma_v &= \text{block\_diag}(\Sigma_a, \Sigma_z, \Sigma_l) \\ \Pi &= \begin{bmatrix} 0_{2\times 2} \\ I_2 \\ 0_{2\times 2} \end{bmatrix}. \end{aligned} \tag{16}$$

As shown later, these assumptions will result in a piecewise planar constraint, with these modeled Gaussians providing a straightforward covariance propagation method for these extra dimensions.

B. Preintegration

Using this measurement, we can perform full 6 DoF integration using local coordinates. By using Eq. 9, the deterministic continuous dynamics will be

$$\begin{aligned} &\dot{X}(t) = X(t)v(t)^\wedge \\ \implies &\dot{\theta}(t) = H(\theta)^{-1}(X(t)^{-1}X(t)v(t)^\wedge)^\vee = H(\theta)^{-1}v. \end{aligned} \tag{17}$$

Using a simple Euler integration scheme results in our preintegration equation. Here we give the equations over a single timestep and over the entire preintegration window of size N as well,

$$\begin{aligned} \theta_{i+1} &= \theta_i + H(\theta_i)^{-1}v_i\Delta t_i \\ \implies \theta_N &= \sum_{k=1}^{N} H(\theta_k)^{-1}v_k\Delta t_k. \end{aligned} \tag{18}$$

Note the iteration is naturally started with $\theta_1 = 0$. This result is somewhat expected as wheel odometry can be integrated locally trivially. However this formulation provides a simple method for covariance propagation and first order updates for wheel slip and intrinsic estimation without the need to repeat the integration.

C. Covariance Propagation

For use in MAP estimation, we need a covariance estimate for θ_N. We obtain this by reintroducing Gaussian noise into Eq. 18 and expanding using a first order Taylor series about $\theta_i = 0$,

$$\begin{aligned} \theta_{i+1} &= \theta_i + \Delta t_i H(\theta_i)^{-1}(v_i + \mathrm{w}_i) \\ &\approx \frac{\partial \theta_{i+1}}{\partial \theta_i}\theta_i + \frac{\partial \theta_{i+1}}{\partial \mathrm{w}_i}\mathrm{w}_i \\ &\triangleq A_i\theta_i + B_i\mathrm{w}_i. \end{aligned} \tag{19}$$

Assuming a small θ_i which holds over small preintegration windows, the following approximation [26] is valid,

$$H(\theta)^{-1} \approx I_6 + \frac{1}{2}\begin{bmatrix} (\theta^\xi)^\wedge & 0 \\ (\theta^\rho)^\wedge & (\theta^\xi)^\wedge \end{bmatrix} \tag{20}$$

where θ^ξ corresponds to the rotational component, and θ^ρ the translational component. Using this approximation, A_i and B_i are given by

$$\begin{aligned} A_i &= I_6 + \Delta t_i \frac{\partial H(\theta_i)^{-1}v_i}{\partial \theta_i} \approx I_6 - \frac{\Delta t_i}{2}\begin{bmatrix} (v_i^a)^\wedge & 0 \\ (v_i^l)^\wedge & (v_i^a)^\wedge \end{bmatrix} \\ B_i &= \Delta t_i H(\theta_i)^{-1}. \end{aligned} \tag{21}$$

where v^a corresponds to the angular velocities, and v^l the linear. Using these approximations results in an iterative method for covariance propagation, letting Σ_i^θ be the covariance of θ_i,

$$\Sigma_{i+1}^\theta = A_i\Sigma_i^\theta A_i^\top + B_i\Sigma_v B_i^\top. \tag{22}$$

D. Slip & Intrinsics Estimation

Further, to remove the need to repeat calculations of Eq. 18, we compute first order updates for both the wheel intrinsics n and slip estimate s, which should be reasonably accurate over small preintegration windows. To do so, we make the assumption of constant wheel intrinsics across the preintegration window given by some initial estimate of $\bar{n}$. Additionally, we replace the slip in a single timestep s_i with its average across the preintegration window,

$$\bar{s} = \frac{1}{N}\sum_{k=1}^{N} s_i. \tag{23}$$

Introducing these assumptions into Eq. 18 results in

$$\begin{aligned} \theta_{i+1} &= \theta_i + H(\theta_i)^{-1}v_i\Delta t_i \\ &= \theta_i + H(\theta_i)^{-1}\Pi M(\bar{n})(\omega_i + \bar{s})\Delta t_i. \end{aligned} \tag{24}$$

As large slip outliers should occur infrequently, we begin by linearizing about a wheel slip of 0. Then, taking the first order expansion for wheel slip and using the chain rule results in,

$$\begin{aligned} G_{i+1} \triangleq \frac{\partial \theta_{i+1}}{\partial \bar{s}} &= \frac{\partial \theta_{i+1}}{\partial \theta_i}\frac{\partial \theta_i}{\partial \bar{s}} + \frac{\partial \theta_{i+1}}{\partial \bar{s}} \\ &= AG_i + H(\theta_i)^{-1}\Pi M(\bar{n})\Delta t_i. \end{aligned} \tag{25}$$

This gives an iterative equation for updating this Jacobian online. We can use G_N to approximately remove any effect slip had on θ_N by subtracting $G_N\hat{s}$ from θ_N. A similar setup can be done for the wheel intrinsics linearized about some estimate of the intrinsics $\bar{n}$,

$$\begin{aligned} F_{i+1} \triangleq \frac{\partial \theta_{i+1}}{\partial \bar{n}} &= \frac{\partial \theta_{i+1}}{\partial \theta_i}\frac{\partial \theta_i}{\partial \bar{n}} + \frac{\partial \theta_{i+1}}{\partial n} \\ &= AF_i + H(\theta_i)^{-1}\Pi N(\omega, n)\Delta t_i \\ N(\omega, n) &= \begin{bmatrix} -\frac{\omega_r r_r - \omega_l l_l}{b^2} & -\frac{\omega_l}{b} & \frac{\omega_r}{b} \\ 0 & \frac{\omega_l}{2} & \frac{\omega_r}{2} \end{bmatrix}. \end{aligned} \tag{26}$$

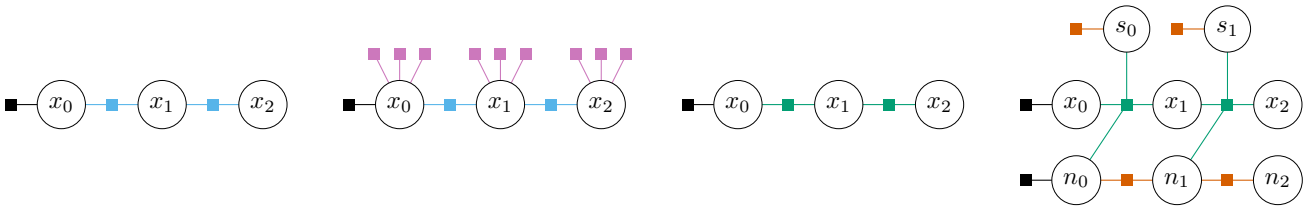

(a) Underconstrained (b) Planar (c) ROSE-- (d) ROSE

Fig. 4: Two steps of various wheel factors used in the manuscript. Light blue represents the standard 2D integration that only constrains x, y, and yaw. Pink are the planar roll, pitch, and z priors. Green is our on manifold 6D integration constraining all of SE(3). Finally, orange represents the slip prior and intrinsics constraint used in ROSE. More succinctly, (a) and (b) represent prior works, with (c) and (d) representing our contributions.

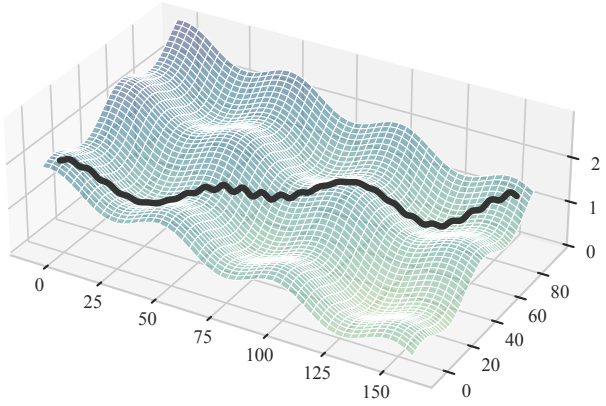

Fig. 5: 3D surface used for simulating off-road driving. The path taken by the simulated vehicle is shown by the black line. To not completely disqualify planar methods, we limited the manifold to have a z position range of $1.5\,\mathrm{m}$ and the pitch and roll a maximum value of $2.5°$.

In this case, we want to capture any effect $\hat{n} - \bar{n}$ that was missed in the original preintegration, so we add $F_N(\hat{n} - \bar{n})$ to θ_N. Putting Eqs. 18, 22, 25, and 26 all together results in our preintegration equation as

$$\begin{aligned}\tilde{\theta}_N(\hat{s}, \hat{n}) = &\left(\sum_{k=1}^{N} H(\theta_k)^{-1} \tilde{v} \Delta t_i\right) \\ &- G_N \hat{s} + F_N(\hat{n} - \bar{n}) + \mathrm{w}_N \\ \mathrm{w}_N &\sim \mathcal{N}(0, \Sigma_N^{\theta}).\end{aligned} \tag{27}$$

Each preintegration window will be initialized with $\theta_1 = 0, \Sigma_1^{\theta} = 0$.

E. Residuals

Putting the final preintegration equation in Eq. 27 into the original differential equation solution in Eq. 7 results in a prediction step for $X_i, X_j \in \mathrm{SE}(3)$ with N wheel encoder measurements between the poses,

$$\hat{X}_j = \hat{X}_i \exp\left((\theta_N(\hat{s}, \hat{n}) + \mathrm{w}_N)^{\wedge}\right). \tag{28}$$

Using a first order approximation to separate the noise, this can be rearranged to find the residual equation

$$r_{ij}^{\theta} = -\mathrm{w}_N = \log\left(\exp(-\theta_N(\hat{s}, \hat{n})^{\wedge}) \hat{X}_i^{-1} \hat{X}_j\right) \tag{29}$$

with r_{ij}^{θ} having the same distribution as w_N. This residual, along with our integration, results in a piecewise planar constraint on vehicle motion. Essentially, the x, y, and yaw directions are constrained using vehicle movement and pitch, roll, and z directions are constrained to be constant across the window. Also note, a transformation to other sensors frames can also trivially be done in this residual as well.

Additionally, to allow for time varying wheel intrinsics, in each preintegration window a wheel intrinsics $\hat{n}_i$ is estimated, with a simple residual disallowing significant changes over time,

$$r_{ij}^{n} = \hat{n}_i - \hat{n}_j \tag{30}$$

where r_{ij}^{n} is Gaussian distributed with covariance Σ^{n}.

Finally, since wheel slippage should only generally be mean-zero, we add a zero prior on the slip,

$$r_i^{s} = \hat{s}_i. \tag{31}$$

with corresponding variance for each slip prior given by σ^{s2}. We additionally wrap this residual in a Tukey robust kernel given by,

$$\rho(x) = \begin{cases} |x| \leq c & \frac{c^2}{6}\left(1 - \left(1 - (x/c)^2\right)^3\right) \\ |x| > c & \frac{c^2}{6}. \end{cases} \tag{32}$$

This prevents the optimizer from utilizing the slip estimate as an arbitrary slack variable, but if there is strong enough of a pull out of the robust basin, as shown in Fig. 3, it will enter a zone of zero gradient and won't impact the slip estimate.

These residuals can be summarized by the orange and green factors as shown in Fig. 4(d) and are combined into a single cost as,

$$r_{ij}^2 = r_{ij}^{\theta\top} (\Sigma_N^{\theta})^{-1} r_{ij}^{\theta} + r_{ij}^{n\top} (\Sigma^{n})^{-1} r_{ij}^{n} + \rho\left(\frac{r_i^s}{\sigma^s}\right) \tag{33}$$

This cost can then be summed over a fixed-lag window, and summed with other sensor costs such as cameras, IMUs, or LiDARs.

In summary, we have developed an on-manifold SE(3) preintegration of wheel odometry with additional wheel slip and wheel intrinsics estimation for usage in an optimization-based state estimation. It relies on the assumptions of mean-zero lateral and vertical linear velocities, mean-zero pitch and roll rates, and a handful of approximations that require small preintegration windows, all of which hold for our domain of off-road driving.

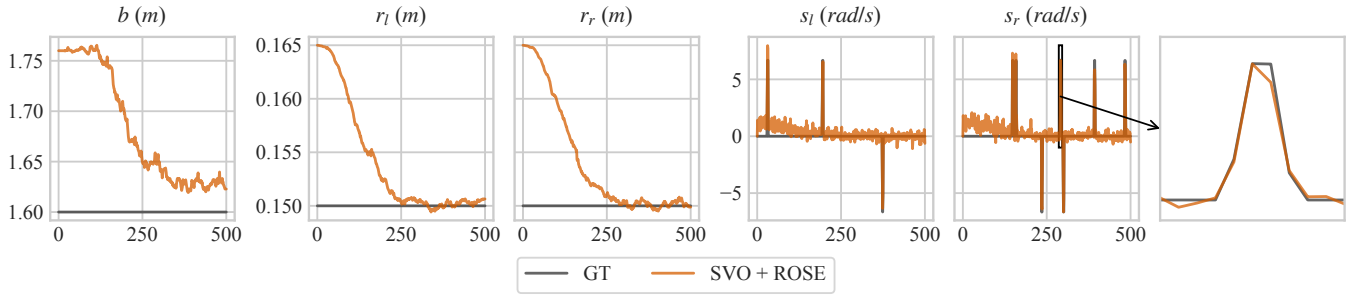

Fig. 6: Convergence of wheel intrinsics and wheel slip events for one of the Monte Carlo simulation runs showing ROSE fused with stereo visual odometry (SVO). ROSE is able to observe wheel slip and wheel intrinsics accurately online.

V. Experiments

To evaluate the effectiveness of our preintegrated off-road wheel odometry, experiments were done on both simulated and real-world datasets. While our method can be fused with any other sensor, for our experiments we chose to fuse with stereo cameras and an IMU for VWO and VIWO systems.

For simplicity, we use a basic stereo camera measurement $\begin{bmatrix} u_l & u_r & v \end{bmatrix}^\top$, where u_l and u_r represent the left and right pixel column coordinate respectively, and v represents the shared row coordinate of the rectified pair. Using this measurement, a residual is defined using the difference between the measured pixel coordinates and the expected coordinates from projection using calibrated camera intrinsics and extrinsics.

A. Baseline

For our baseline, we compare against a basic wheel integration scheme that integrates in 2D – namely providing constraints in x, y, and yaw. Further, to constrain the remaining dimensions, and as is common in previous literature [12], [13], we additionally add a prior on the z, pitch, and roll dimensions. The 2D integration with planar priors we denote as the "Planar" method moving forward. We compare against stereo visual odometry (SVO), integrated wheel odometry (WO), our base 6D integration method (ROSE--) to justify our 6D integration, and our full method with both slip and intrinsics estimation (ROSE). The factors for each of these are shown in Fig. 4.

B. Simulation Validation

The simulation experiments were done by integrating wheel odometry moving across the surface shown in Fig. 5. This surface was created by a combination of low frequency $\cos$ functions. Further note, the surface has z-range of about $1.5\,\mathrm{m}$ and a pitch and roll maximum value of $2.5°$ for a fair comparison to the planar methods. In many off-road scenarios, there will potentially be a much larger gap that will naturally favor our method.

Stereo vision was simulated by sampling 3D points near the trajectory and projecting them using the camera intrinsics. More points were sampled if less than 50 points were found in a single frame. Stereo cameras were simulated at $5\,\mathrm{Hz}$ and wheel odometry at $100\,\mathrm{Hz}$. The vehicle was simulated as moving for $100\,\mathrm{s}$ for a total distance of $200\,\mathrm{m}$.

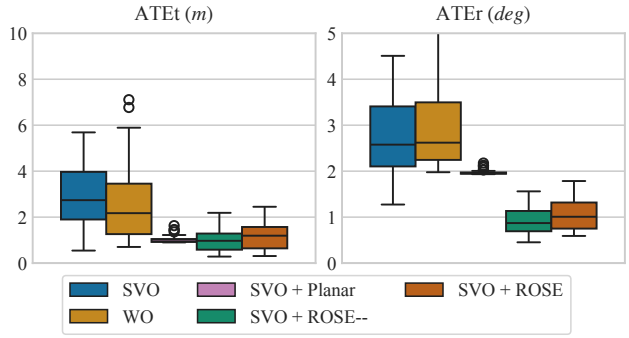

Fig. 7: Results of a Monte Carlo simulation with 50 runs in idealized conditions, i.e. correct wheel intrinsics and no wheel slippage. The 6D integration used in our method (ROSE) performs more accurately than previous work.

1) Idealized Run: First, to justify our usage of our 3D integration, we compare in an idealized scenario with exactly correct intrinsics and no wheel slippage. For completeness, we ran the simulation 50 times and computed both the translational and rotational average trajectory error (ATEt and ATEr, respectively). Results are shown in Fig. 7.

It can be seen that ROSE-- without the extra estimation performed comparably to the baseline in ATEt, while it outperformed it in ATEr by a fair margin. This is likely due to our method additionally estimating pitch and roll, while the baseline assumes they are 0.

2) Corrupted Run: Next, to test the advantages of the additional wheel slip and intrinsics estimation we perturb the initial wheel intrinsics estimate, by 10% and randomly introduce 10 wheel slip events each with a duration of $0.5\,\mathrm{s}$ and length of $0.5\,\mathrm{m}$. The Monte Carlo simulation results can be seen in Fig. 9 and the convergence of wheel intrinsics and slip estimation from a single run is shown in Fig. 6.

The observability of both the wheel intrinsics and slip when fused with stereo vision can be seen in the convergence in Fig. 6. Notice that each slip event was accurately estimated by our method to a fairly high degree. Also of note, at the beginning of the estimation, the system initially estimates error in the system as slip; however, after the wheel intrinsics converge, the slip and intrinsics are estimated much more accurately.

The significant impact that wheel slip events have can also been seen in Fig. 9. In all of the scenarios where wheel slip was not being estimated, the estimate ended up being poorer than

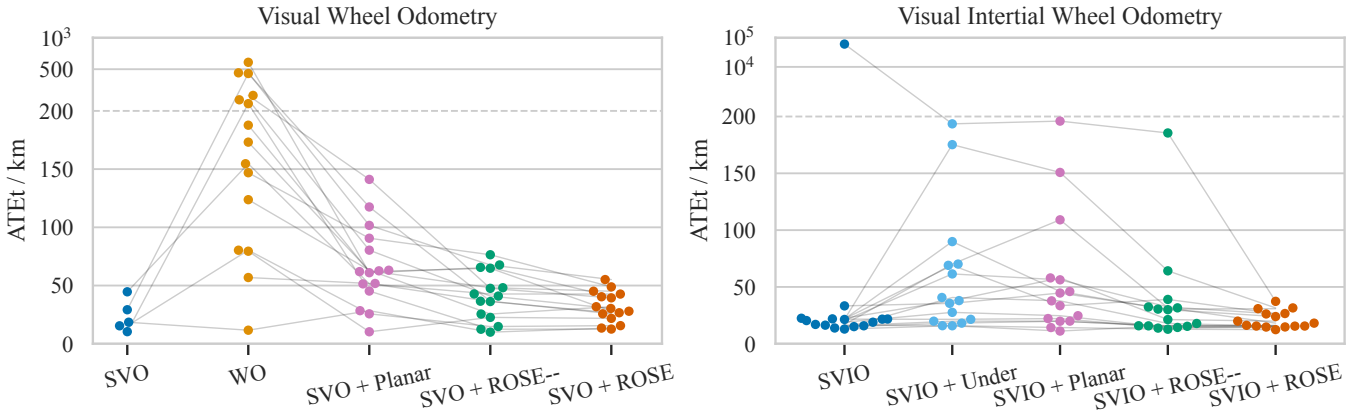

Fig. 8: Results from 15 trajectory runs taken from an off-road vehicle in Penn Hills, Pennsylvania. Each connected line represents a single trajectory run. On the left are the results from the VWO systems and VIWO systems on the right. The underconstrained case isn't included on the left as it often failed due to lacking constraints. Additionally, the stereo odometry failed on the majority of trajectories due to failing to gather matches when moving at high speeds. The y-axis scales switch from linear to logarithmic at 200 m. Notice that our 6D preintegration (ROSE-- and ROSE) always outperformed the baselines and the addition of slip and intrinsics estimation (ROSE) often improved results for the scenarios where there was significant slippage.

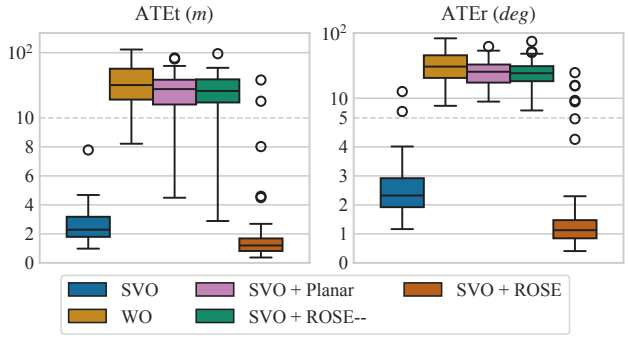

Fig. 9: Results of a Monte Carlo simulation with 50 runs in real world conditions. Wheel intrinsics were initialized 10% too large, and 10 slip events occurred randomly, each with a length of 0.5 s and length 0.5 m. Note the y-axis scales change from linear to logarithmic at 1.0 m and 5°, respectively. This is in contrast to the scales of Fig. 7. ROSE is able to estimate wheel slip events and wheel intrinsics online, resulting in a more robust system.

just using stereo odometry alone, explaining the limited usage of wheel encoders in state estimation today. However, when the wheel slip is taken into account, ROSE is able to improve on the stereo only methods, and achieves similar performance as in the idealized scenario.

C. Real World Experiments

Real world experiments were done on a six-wheeled skid steer off-road vehicle equipped with wide-angle stereo cameras running at 8 Hz, a KVH-branded inertial measurement unit (IMU) at 100 Hz, and two wheel encoders equipped on the middle wheels on the left and right side at 20 Hz.

For the stereo vision, features were found using FAST features [30], left and right image features were matched using a rectified disparity image, and features were temporally propagated using KLT tracking [31]. We additionally use IMU preintegration [32] when IMU fusion is done.

The wheel encoder values were linearly interpolated to generate synthetic measurements at the same timesteps as the stereo cameras. Fifteen experimental trajectories were taken ranging in length from 0.4 to 1.7 km in Penn Hills, Pennsylvania. An example is shown in Fig. 1 and results are shown in Fig. 8. We divided the ATEt error by the length of the trajectory to get a better comparison across experiments. Covariances used are shown in Table I.

TABLE I: Experiment Noise Statistics

Covariance	Value
Wheel Encoder	$5 \times 10^{-2}\,\mathrm{rad/sec}/\sqrt{\mathrm{Hz}}$
Angular Velocity	$5 \times 10^{-2}\,\mathrm{rad/sec}/\sqrt{\mathrm{Hz}}$
Linear Velocity	$1 \times 10^{-2}\,\mathrm{m/sec}/\sqrt{\mathrm{Hz}}$
Intr. Radius Bet.	$7 \times 10^{-6}\,\mathrm{m}$
Intr. Baseline Bet.	$9 \times 10^{-4}\,\mathrm{m}$
Slip Prior	$1 \times 10^{-2}\,\mathrm{rad/sec}$
Slip Kernel c	1.0

For the scenarios where an IMU is used, we fuse with stereo visual inertial odometry (SVIO). In this scenario, we also compare against the 2D integration without the planar priors also known as the underconstrained (Under) case.

Due to the high-speed nature of many of the trajectories, there were a number of timesteps that failed to gather reliable feature matches between images, resulting in a failure for visual odometry and justifying the use of sensor fusion. Additionally in the VWO case, the underconstrained baseline often failed due to degenerate systems and is not included. Further, notice that our base 6D preintegration (ROSE--) performed as well or better than the planar scenario in the majority of trajectories, while the slip and intrinsics estimation version clearly performed the best. In some scenarios it appears the slip and intrinsics estimation didn't help significantly; this is to be expected as in experiments without extensive slip occurring or with minimal changes to wheel intrinsics, ROSE will likely perform similarly to the base method, ROSE--.

When IMU measurements are included, it can be seen that our method does not worsen the accuracy of the SVIO method, but does improve it slightly and naturally adds robustness to sensor failures. The underconstrained and planar versions combined with SVIO mostly degraded the state estimation. Our methods, both ROSE and ROSE--, seemed to stabilize

results a fair amount. One trajectory had outlier wheel odometry measurements resulting in enormous errors in the other VIWO methods. Another trajectory began while moving, in which case the SVIO method failed to accurately initialize and estimate IMU biases and resulted in extremely large ATEt. In both of these scenarios, ROSE was still able to accurately perform state estimation, demonstrating the robustness it can add to existing sensor fusion methods. While other methods do exist to solve IMU bias estimation, we find incorporating ROSE into existing SVIO systems to be a straightforward and effective solution.

VI. Conclusion & Future Work

These experiments show that our method is able to avoid many of the problems that wheel odometry encounters in off-road scenarios. The standard approaches often fail due to being underconstrained, planar assumptions, wheel slippage, or an inability to function in the case of other sensor failures. ROSE avoids these issues by leveraging a full 6D preintegration scheme and estimating wheel intrinsics and slippage online. This provides an added level of robustness to the state estimation as well as potential for increased accuracy.

Future work should be focused on reducing the number of tunable covariances of the system as this is the largest current drawback of this method. This could be done by implementing a heuristic to auto-tune the covariances or via some form of online estimation.

References

[1] Z. Zhang, H. Liu, J. Qi, K. Ji, G. Xiong, and J. Gong, "A tightly coupled LIDAR-IMU SLAM in off-road environment," in *2019 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, (Cairo, Egypt), pp. 1–6, IEEE, Sept. 2019.

[2] G. Huang, "Visual-inertial navigation: A concise review," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9572–9582, May 2019.

[3] J. Y. Wong, *Theory of Ground Vehicles*. John Wiley & Sons, Mar. 2001.

[4] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," Sept. 2012.

[5] Y. He, Z. Chai, X. Liu, Z. Li, H. Luo, and F. Zhao, "Tightly-coupled vision-gyro-wheel odometry for ground vehicle with online extrinsic calibration," in *2020 3rd International Conference on Intelligent Autonomous Systems (ICoIAS)*, pp. 99–106, Feb. 2020.

[6] S. Bai, J. Lai, P. Lyu, Y. Cen, B. Wang, and X. Sun, "Graph-optimisation-based self-calibration method for IMU/odometer using preintegration theory," *The Journal of Navigation*, vol. 75, pp. 594–613, May 2022.

[7] J. Liu, W. Gao, C. Xie, and Z. Hu, "Implementation and observability analysis of visual-inertial-wheel odometry with robust initialization and online extrinsic calibration," *Robotics and Autonomous Systems*, vol. 176, p. 104686, June 2024.

[8] X. Yu, S. Teng, T. Chakhachiro, W. Tong, T. Li, T.-Y. Lin, S. Koehler, M. Ahumada, J. M. Walls, and M. Ghaffari, "Fully proprioceptive slip-velocity-aware state estimation for mobile robots via invariant Kalman filtering and disturbance observer," Sept. 2023.

[9] T.-Y. Lin, T. Li, W. Tong, and M. Ghaffari, "Proprioceptive invariant robot state estimation," Feb. 2024.

[10] W. Lee, K. Eckenhoff, Y. Yang, P. Geneva, and G. Huang, "Visual-inertial-wheel odometry with online calibration," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4559–4566, Oct. 2020.

[11] Z. Dang, T. Wang, and F. Pang, "Tightly-coupled data fusion of VINS and odometer based on wheel slip estimation," in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 1613–1619, Dec. 2018.

[12] M. Quan, S. Piao, M. Tan, and S.-S. Huang, "Tightly-coupled monocular visual-odometric SLAM using wheels and a MEMS gyroscope," *IEEE Access*, vol. 7, pp. 97374–97389, 2019.

[13] K. J. Wu, C. X. Guo, G. Georgiou, and S. I. Roumeliotis, "VINS on wheels," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5155–5162, May 2017.

[14] T. Hua, T. Li, and L. Pei, "PIEKF-VIWO: Visual-inertial-wheel odometry using partial invariant extended kalman filter," Mar. 2023.

[15] M. Zhang, Y. Chen, and M. Li, "Vision-aided localization for ground robots," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2455–2461, Nov. 2019.

[16] M. Zhang, X. Zuo, Y. Chen, Y. Liu, and M. Li, "Pose estimation for ground robots: On manifold representation, integration, reparameterization, and optimization," *IEEE Transactions on Robotics*, vol. 37, pp. 1081–1099, Aug. 2021.

[17] P. Zhou, Y. Liu, P. Gu, J. Liu, and Z. Meng, "Visual localization and mapping leveraging the constraints of local ground manifolds," *IEEE Robotics and Automation Letters*, vol. 7, pp. 4196–4203, Apr. 2022.

[18] Y. Tian, N. Sidek, and N. Sarkar, "Modeling and control of a nonholonomic wheeled mobile robot with wheel slip dynamics," in *2009 IEEE Symposium on Computational Intelligence in Control and Automation*, pp. 7–14, Mar. 2009.

[19] R. Balakrishna and A. Ghosal, "Modeling of slip for wheeled mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 11, pp. 126–132, Feb. 1995.

[20] C. C. Ward and K. Iagnemma, "A dynamic-model-based wheel slip detector for mobile robots on outdoor terrain," *IEEE Transactions on Robotics*, vol. 24, pp. 821–831, Aug. 2008.

[21] F. Rogers-Marcovitz, M. George, N. Seegmiller, and A. Kelly, "Aiding off-road inertial navigation with high performance models of wheel slip," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 215–222, Oct. 2012.

[22] A. Censi, A. Franchi, L. Marchionni, and G. Oriolo, "Simultaneous calibration of odometry and sensor parameters for mobile robots," *IEEE Transactions on Robotics*, vol. 29, pp. 475–492, Apr. 2013.

[23] C. X. Guo, F. M. Mirzaei, and S. I. Roumeliotis, "An analytical least-squares solution to the odometer-camera extrinsic calibration problem," in *2012 IEEE International Conference on Robotics and Automation*, pp. 3962–3968, May 2012.

[24] L. Heng, B. Li, and M. Pollefeys, "CamOdoCal: Automatic intrinsic and extrinsic calibration of a rig with multiple generic cameras and odometry," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1793–1800, Nov. 2013.

[25] J. Deray, J. Solà, and J. Andrade-Cetto, "Joint on-manifold self-calibration of odometry model and sensor extrinsics using pre-integration," in *2019 European Conference on Mobile Robots (ECMR)*, pp. 1–6, Sept. 2019.

[26] J. Solà, J. Deray, and D. Atchuthan, "A micro Lie theory for state estimation in robotics," *arXiv:1812.01537 [cs]*, Dec. 2021.

[27] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, and A. Zanna, "Lie-group methods," *Acta Numerica*, vol. 9, pp. 215–365, Jan. 2000.

[28] "Gtsam/doc/ImuFactor.pdf at 4.2.0 · borglab/gtsam · GitHub." https://github.com/borglab/gtsam/blob/4.2.0/doc/ImuFactor.pdf.

[29] F. Kschischang, B. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, pp. 498–519, Feb. 2001.

[30] E. Rosten, R. Porter, and T. Drummond, "Faster and better: A machine learning approach to corner detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 105–119, Jan. 2010.

[31] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *IJCAI'81: 7th International Joint Conference on Artificial Intelligence*, vol. 2, (Vancouver, Canada), pp. 674–679, Aug. 1981.

[32] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-manifold preintegration for real-time visual–inertial odometry," *IEEE Transactions on Robotics*, vol. 33, pp. 1–21, Feb. 2017.