

The Complex-Step Derivative Approximation on Matrix Lie Groups

Charles Champagne Cossette¹, Alex Walsh², and James Richard Forbes³

Abstract—The complex-step derivative approximation is a numerical differentiation technique that can achieve analytical accuracy, to machine precision, with a single function evaluation. In this paper, the complex-step derivative approximation is extended to be compatible with elements of matrix Lie groups. As with the standard complex-step derivative, the method is still able to achieve analytical accuracy, up to machine precision, with a single function evaluation. Compared to a central-difference scheme, the proposed complex-step approach is shown to have superior accuracy. The approach is applied to two different pose estimation problems, and is able to recover the same results as an analytical method when available.

I. Introduction

Attitude and pose, ubiquitous entities of interest in robotics problems, are most naturally represented as elements of matrix Lie groups. Path planning, state estimation, and control algorithms often require Jacobian computations with respect to attitude and pose. Often these Jacobians are computed analytically, by hand, via a Taylor-series expansion while adhering to the matrix Lie group structure of the problem [Barfoot2019]. However, in some cases analytical computation of Jacobians may be impractical, necessitating a numerical procedure. Numerical computation of Jacobians is also useful for quickly comparing algorithms that require Jacobians, before investing effort into one specific algorithm and the associated analytically derived Jacobians. Numerical Jacobians can also be used to verify Jacobians that are derived by hand.

A variety of numerical differentiation techniques appropriate for matrix Lie groups can be found in the literature. In [Absil2008] a forward-difference method is described for general matrix manifolds, a method that is used in the open-source software MANOPT [Boumal2014]. A central-difference method is employed in the open-source software GTSAM [GTSAM2019a], and algorithmic differentiation methods are presented in [Robenack2011, Sommer2013]. The Python-based software PYMANOPT [Townsend2016] is an open-source optimization toolbox for matrix manifolds that employs algorithmic differentiation. SOPHUS [Sophus2019] is

another open-source C++ package that exploits the automatic differentiation functionality available in CERES [Ceres2019], a nonlinear least-squares library developed by Google. However, algorithmic differentiation can be time consuming to implement and finite-differencing is prone to subtractive cancellation errors, thus limiting precision [Martins2003]. The complex-step derivative approximation is a numerical method for computing first derivatives that does not suffer from subtractive cancellation errors [Martins2003]. One of the earlier appearances of the complex-step derivative can be found in [Squire1998], where the derivatives of scalar functions of real variables are evaluated. In [Martins2003], the complex-step derivative is investigated further, along with its use in Fortran, C/C++, and other languages. An application to a multidisciplinary design optimization problem is also shown. This method has gained popularity due to its ability to realize machine-precision accuracy of derivative computations, and doing so without tuning the step size, since it can be reduced to an arbitrarily small value. The complex-step derivative also requires only one complex function evaluation, which is beneficial compared to central-differencing when the function is expensive to evaluate. The complex-step derivative is straight-forward to implement, especially in Matlab, where the default variable type is complex.

This paper considers the formulation and application of the complex-step derivative approximation to functions of matrix Lie group elements. The aforementioned advantages of the standard complex-step derivative remain present, while the proposed method can be used to compute both left and right Jacobians. Various examples are presented, demonstrating the utility and advantages of the matrix Lie group version of the complex-step derivative. In particular, pose estimation problems are considered, one using the ETH Zürich EuRoC dataset [Burri2016], where analytical Jacobians are available for comparison, and one using the ‘Lost in the Woods’ dataset [Barfoot2009], where computation of analytical Jacobians is possible, but time consuming. When solving for the maximum a posteriori (MAP) estimate of the pose using a Gauss-Newton algorithm, it is shown that computing the Jacobians using the complex-step derivative realizes the same accuracy and convergence properties as when analytical Jacobians are used.

II. Preliminaries

A. Matrix Lie Groups

A matrix Lie group \mathcal{G} is a Lie group that consists of the set of $m \times m$ invertible matrices, where the group

*This work was supported by NSERC Collaborative Research and Development, Engage, and Discovery Grant programs.

¹Ph.D. Candidate, Department of Mechanical Engineering, McGill University, 817 Sherbrooke St. W., Montreal, QC, Canada, H3A 0C3. e-mail: charles.cossette@mail.mcgill.ca

²Postdoctoral Researcher, Department of Mechanical Engineering, McGill University, 817 Sherbrooke St. W., Montreal, QC, Canada, H3A 0C3. e-mail: alex.walsh@mail.mcgill.ca

³Associate Professor, Department of Mechanical Engineering, McGill University, 817 Sherbrooke St. W., Montreal, QC, Canada, H3A 0C3. e-mail: james.richard.forbes@mcgill.ca

operation is matrix multiplication [Chirikjian2009]. From the definition of a group, a matrix Lie group is closed under matrix multiplication. That is, given $\mathbf{X}, \mathbf{Y} \in \mathcal{G}$, it follows that $\mathbf{XY} \in \mathcal{G}$. A matrix Lie group is a closed subgroup of the general linear group defined by [Hall2015]

$$GL(m, \mathbb{C}) = \{\mathbf{X} \in \mathbb{C}^{m \times m} \mid \det(\mathbf{X}) \neq 0\},$$

which is also a matrix Lie group. The matrix Lie algebra of \mathcal{G} is denoted \mathfrak{g} , and is defined as [Hall2015],

$$\mathfrak{g} = \{\Xi \mid \exp(t\Xi) \in \mathcal{G}, \forall t \in \mathbb{R}\}. \quad (1)$$

It can be shown that the matrix Lie algebra defined by (1) is a valid Lie algebra [Hall2015], and is a vector space closed under the operation of the Lie bracket $[\cdot, \cdot]$, which can be computed by $[\mathbf{A}, \mathbf{B}] = \mathbf{AB} - \mathbf{BA} \in \mathfrak{g}$, for all $\mathbf{A}, \mathbf{B} \in \mathfrak{g}$. The wedge operator $(\cdot)^\wedge : \mathbb{R}^n \rightarrow \mathfrak{g}$ maps a column matrix to the matrix Lie algebra. The exponential map $\exp(\cdot) : \mathfrak{g} \rightarrow \mathcal{G}$ maps an element of the matrix Lie algebra to the matrix Lie group, and is computed using the matrix exponential. The only matrix Lie group elements $\mathbf{X} \in \mathcal{G}$ that are considered in this paper are those that can be written as

$$\mathbf{X} = \exp(\xi^\wedge),$$

where $\xi \in \mathbb{R}^n$. The “vee” operator $(\cdot)^\vee : \mathfrak{g} \rightarrow \mathbb{R}^n$ maps an element of the matrix Lie algebra to a column matrix. The logarithmic map $\ln(\cdot) : \mathcal{G} \rightarrow \mathfrak{g}$ maps an element of the matrix Lie group to the matrix Lie algebra, and is computed by the matrix logarithm. A parameterization of the group \mathcal{G} can be retrieved from \mathbf{X} via

$$\xi = \ln(\mathbf{X})^\vee,$$

when the matrix logarithm is well defined. The adjoint representation of \mathbf{X} is denoted $\text{Ad}(\mathbf{X})$, such that $(\text{Ad}(\mathbf{X})\zeta)^\wedge = \mathbf{X}\zeta^\wedge\mathbf{X}^{-1}$, $\zeta \in \mathbb{R}^n$. This leads to the identity

$$\exp((\text{Ad}(\mathbf{X})\zeta)^\wedge) = \mathbf{X} \exp(\zeta^\wedge) \mathbf{X}^{-1}.$$

The Baker-Campbell-Hausdorff (BCH) formula is the solution to

$$\mathbf{z} = \ln(\exp(\xi_1^\wedge) \exp(\xi_2^\wedge)),$$

and the exact solution is an infinite sum [Barfoot2019]. A first-order approximation to the BCH formula is

$$\ln(\exp(\xi_1^\wedge) \exp(\xi_2^\wedge)) = \xi_1^\wedge + \xi_2^\wedge,$$

which is exact in the event that $[\xi_1^\wedge, \xi_2^\wedge] = \mathbf{0}$. Such an approximation is typically used when both ξ_1 and ξ_2 are assumed to be small. The details of the special Euclidean groups $SE(2)$, $SE(3)$, and the group of double direct isometries $SE_2(3)$ can be found in the appendix.

B. Gauss-Newton Algorithm

The Gauss-Newton algorithm is an optimization algorithm appropriate for nonlinear least-squares functions of the form

$$J(\mathbf{x}) = \frac{1}{2} \mathbf{e}(\mathbf{x})^\top \mathbf{W} \mathbf{e}(\mathbf{x}), \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{q \times q}$ is a symmetric positive definite weight matrix and $\mathbf{e} : \mathbb{R}^p \rightarrow \mathbb{R}^q$ is some error function. Employing Newton’s method directly on (2) requires the Hessian of $J(\mathbf{x})$, which is potentially difficult to obtain. An alternate strategy is to substitute a first-order approximation of $\mathbf{e}(\mathbf{x})$ about some nominal $\bar{\mathbf{x}}$, given by [Barfoot2019]

$$\mathbf{e}(\bar{\mathbf{x}} + \delta\mathbf{x}) \approx \mathbf{e}(\bar{\mathbf{x}}) + \left. \frac{\partial \mathbf{e}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\bar{\mathbf{x}}} \delta\mathbf{x},$$

into (2), thus yielding the Jacobian and a Hessian approximation of $J(\mathbf{x})$,

$$J(\mathbf{x}) \approx \frac{1}{2} \mathbf{e}(\bar{\mathbf{x}})^\top \mathbf{W} \mathbf{e}(\bar{\mathbf{x}}) + \underbrace{\mathbf{e}(\bar{\mathbf{x}})^\top \mathbf{W} \frac{\partial \mathbf{e}(\mathbf{x})}{\partial \mathbf{x}}}_{\left. \frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} \right|_{\bar{\mathbf{x}}}} \delta\mathbf{x} + \frac{1}{2} \delta\mathbf{x}^\top \underbrace{\left(\frac{\partial \mathbf{e}(\mathbf{x})}{\partial \mathbf{x}} \right)^\top \mathbf{W} \left(\frac{\partial \mathbf{e}(\mathbf{x})}{\partial \mathbf{x}} \right)}_{\left. \frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \right|_{\bar{\mathbf{x}}}} \delta\mathbf{x}.$$

The Gauss-Newton algorithm then proceeds identically to Newton’s method, where the nominal point is iterated by $\bar{\mathbf{x}}_\ell = \bar{\mathbf{x}}_{\ell-1} + \delta\mathbf{x}_{\ell-1}$. The step $\delta\mathbf{x}_{\ell-1}$ is calculated as

$$\delta\mathbf{x}_{\ell-1} = - \left(\left. \frac{\partial J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \right|_{\bar{\mathbf{x}}_{\ell-1}} \right)^{-1} \left(\left. \frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} \right|_{\bar{\mathbf{x}}_{\ell-1}} \right)^\top.$$

III. The Complex-Step Derivative Approximation

A. Review

Consider the complex-differentiable function $f : \mathbb{C} \rightarrow \mathbb{C}$ perturbed about the nominal point \bar{x} by jh where $\bar{x}, h \in \mathbb{R}$ and $j = \sqrt{-1}$. A Taylor series expansion yields

$$f(\bar{x} + jh) = f(\bar{x}) + \left. \frac{\partial f(z)}{\partial z} \right|_{z=\bar{x}} jh - \frac{1}{2} \left. \frac{\partial^2 f(z)}{\partial z^2} \right|_{z=\bar{x}} h^2 - \frac{1}{3!} \left. \frac{\partial^3 f(z)}{\partial z^3} \right|_{z=\bar{x}} jh^3 \dots \quad (3)$$

If $f(\bar{x})$ is assumed to be real for all real \bar{x} , then, to first order, taking the imaginary portion of (3) yields [Squire1998]

$$\left. \frac{\partial f(z)}{\partial z} \right|_{z=\bar{x}} = \frac{\text{Im}\{f(\bar{x} + jh)\}}{h} + \mathcal{O}(h^2).$$

This is valid as long as $f(\bar{x}) \in \mathbb{R}$ for all $\bar{x} \in \mathbb{R}$, and that derivatives are evaluated at strictly real nominal points. From a practical standpoint, a user is often attempting to find derivatives of $f : \mathbb{R} \rightarrow \mathbb{R}$. Providing that this can be extended to $f : \mathbb{C} \rightarrow \mathbb{C}$ such that f is complex-differentiable, then with a minor abuse of notation, this

can construct a derivative approximation for $f(x)$ as written in [Squire1998, Martins2003],

$$\frac{\partial f(x)}{\partial x} \approx \frac{\text{Im}\{f(x + jh)\}}{h}.$$

Since there are no subtractive cancellation errors, the complex-step derivative approximation can produce machine-precision approximations by reducing h to an arbitrarily small step size.

B. The Complex-Step Derivative on Matrix Lie Groups

Consider a complex-differentiable function $f : \mathcal{G} \rightarrow \mathbb{C}$ where $\mathcal{G} \subset GL(m, \mathbb{C})$, $\mathbf{X}(\epsilon^R) = \bar{\mathbf{X}} \exp(\epsilon^R \wedge)$ is parametrizable by a perturbation $\epsilon^R = [\epsilon_1^R \ \epsilon_2^R \ \dots \ \epsilon_n^R]^\top \in \mathbb{C}^n$ on the right, and $\bar{\mathbf{X}} \in \mathbb{R}^{m \times m}$ is some nominal value of \mathbf{X} . Consider perturbing $f(\mathbf{X}(\epsilon^R))$ by $\epsilon^R = \mathbf{0} + jh\mathbf{1}_i$, where $\mathbf{1}_i$ is the i^{th} column of the appropriately-dimensional identity matrix $\mathbf{1}$. The composition $f(\mathbf{X}(\epsilon^R))$ has essentially recast f as $f : \mathbb{C}^n \rightarrow \mathbb{C}$, from which a Taylor series expansion yields [Chirikjian2009]

$$\begin{aligned} f(\bar{\mathbf{X}} \exp((jh\mathbf{1}_i) \wedge)) &= f(\bar{\mathbf{X}}) + \left. \frac{\partial f(\mathbf{X}(\epsilon^R))}{\partial \epsilon_i^R} \right|_{\epsilon^R=\mathbf{0}} jh \\ &\quad - \frac{1}{2} \left. \frac{\partial^2 f(\mathbf{X}(\epsilon^R))}{\partial \epsilon_i^{R^2}} \right|_{\epsilon^R=\mathbf{0}} h^2 + \mathcal{O}(h^3). \end{aligned} \quad (4)$$

Since it is assumed that $f(\bar{\mathbf{X}}) \in \mathbb{R}$, taking the imaginary component of (4) yields an approximation for the derivative

$$\frac{\partial f(\mathbf{X}(\epsilon^R))}{\partial \epsilon_i^R} \approx \frac{\text{Im}\{f(\bar{\mathbf{X}} \exp((jh\mathbf{1}_i) \wedge))\}}{h}. \quad (5)$$

The right Jacobian $\partial f(\mathbf{X}(\epsilon^R))/\partial \epsilon^R$ can be obtained by individually computing the derivatives using (5) with $i = 1, 2, \dots, n$. The left Jacobian can identically be obtained by instead parametrizing \mathbf{X} with $\mathbf{X}(\epsilon^L) = \exp(\epsilon^L \wedge) \bar{\mathbf{X}}$. This leads to

$$\frac{\partial f(\mathbf{X}(\epsilon^L))}{\partial \epsilon_i^L} \approx \frac{\text{Im}\{f(\exp((jh\mathbf{1}_i) \wedge) \bar{\mathbf{X}})\}}{h}. \quad (6)$$

Note that the superscripts on ϵ^R and ϵ^L are simply labels that correspond to right and left perturbations, respectively, as opposed to exponents.

Example 1: Consider the function

$$f(\mathbf{T}) = \mathbf{v}^\top \mathbf{T} \mathbf{y},$$

where $\mathbf{T} \in SE(3)$ and $\mathbf{v}, \mathbf{y} \in \mathbb{R}^4$. The left Jacobian can be determined analytically using the first-order approximation $\mathbf{T} = \exp(\epsilon^L \wedge) \bar{\mathbf{T}} \approx (\mathbf{1} + \epsilon^L \wedge) \bar{\mathbf{T}}$ and a Taylor series expansion. To this end,

$$\begin{aligned} f(\exp(\epsilon^L \wedge) \bar{\mathbf{T}}) &= \mathbf{v}^\top \exp(\epsilon^L \wedge) \bar{\mathbf{T}} \mathbf{y} \\ &\approx \mathbf{v}^\top (\mathbf{1} + \epsilon^L \wedge) \bar{\mathbf{T}} \mathbf{y} \\ &= \mathbf{v}^\top \bar{\mathbf{T}} \mathbf{y} + \underbrace{\mathbf{v}^\top (\bar{\mathbf{T}} \mathbf{y})^\odot}_{\left. \frac{\partial f(\mathbf{T}(\epsilon^L))}{\partial \epsilon^L} \right|_{\epsilon^L=\mathbf{0}}} \epsilon^L, \end{aligned} \quad (7)$$

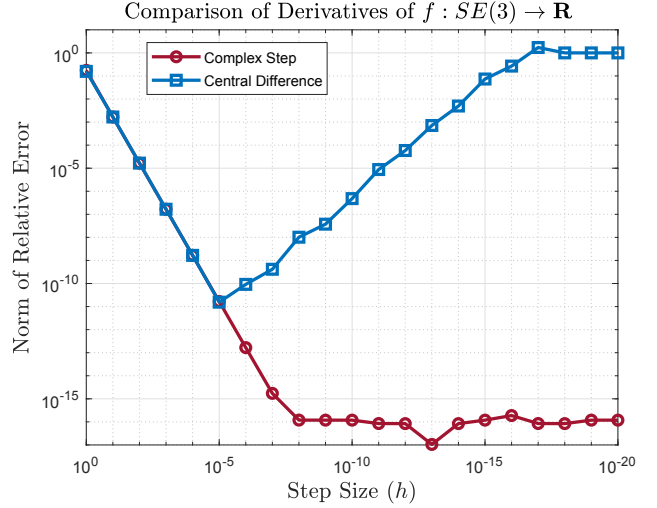


Fig. 1. Variation of relative error in gradient of $f : SE(3) \rightarrow \mathbb{R}$ with step size, for both complex-step and central-difference methods. Machine precision is achievable with a sufficient reduction in step size.

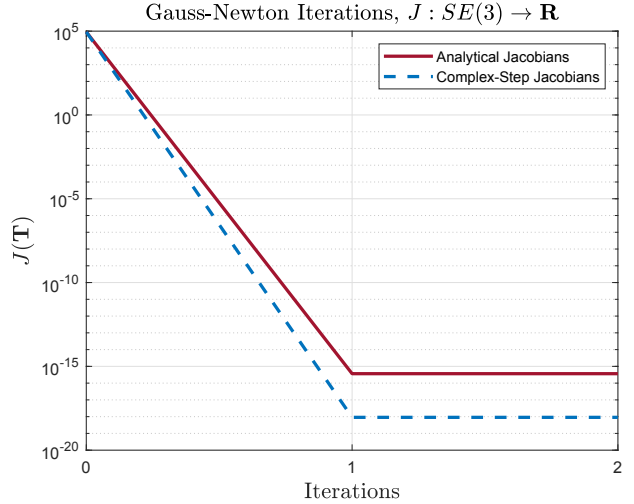


Fig. 2. Convergence history of a Gauss-Newton optimization algorithm on a simple nonlinear least-squares problem. The analytical Jacobians require an approximation to be tractable. The complex-step can calculate Jacobians down to machine precision, hence providing a more accurate first step.

where the $(\cdot)^\odot$ operator is defined in the appendix. The elements of $\partial f(\mathbf{T}(\epsilon^L))/\partial \epsilon^L$ are computed using (6) with varying step sizes h , and the results are compared with a central-difference scheme in Fig. 1. The error is computed by taking the relative 2-norm of the difference between the analytical and numerical solutions. Like the standard complex-step derivative, the complex-step derivative tailored to the matrix Lie group $SE(3)$ is able to achieve analytic accuracy, up to machine precision, for small enough h , while the central-difference derivative is not.

Example 2: Consider the nonlinear least-squares function

$$J(\mathbf{T}) = \frac{1}{2} \mathbf{e}(\mathbf{T})^\top \mathbf{W} \mathbf{e}(\mathbf{T}), \quad (8)$$

where $\mathbf{T} = \exp(\boldsymbol{\epsilon}^{L\wedge})\bar{\mathbf{T}} \in SE(3)$, $\mathbf{W} \in \mathbb{R}^{6 \times 6}$ is a symmetric positive definite weight matrix, and the error is given by

$$\mathbf{e}(\mathbf{T}) = \ln(\mathbf{T}^{-1}\mathbf{T}^{\text{ref}})^\vee.$$

The matrix $\mathbf{T}^{\text{ref}} \in SE(3)$ is some reference point used to construct the error. The Jacobian $\partial \mathbf{e}(\mathbf{T}(\boldsymbol{\epsilon}^L))/\partial \boldsymbol{\epsilon}^L$ can be used to construct Jacobian and Hessian approximations of $J(\mathbf{x})$, which are used in the Gauss-Newton algorithm. Like in Example 1, the analytical left Jacobian can be determined by perturbing \mathbf{T} on the left,

$$\begin{aligned} \mathbf{e}(\exp(\boldsymbol{\epsilon}^{L\wedge})\bar{\mathbf{T}}) &= \ln(\bar{\mathbf{T}}^{-1} \exp(-\boldsymbol{\epsilon}^{L\wedge})\mathbf{T}^{\text{ref}})^\vee \\ &= \ln(\exp((-Ad(\bar{\mathbf{T}}^{-1})\boldsymbol{\epsilon}^L)^\wedge) \underbrace{\bar{\mathbf{T}}^{-1}\mathbf{T}^{\text{ref}}}_{\exp(\mathbf{e}(\bar{\mathbf{T}})^\wedge)})^\vee \\ &\approx \mathbf{e}(\bar{\mathbf{T}}) + \underbrace{(-Ad(\bar{\mathbf{T}}^{-1}))}_{\left. \frac{\partial \mathbf{e}(\mathbf{T}(\boldsymbol{\epsilon}^L))}{\partial \boldsymbol{\epsilon}^L} \right|_{\boldsymbol{\epsilon}^L=0}} \boldsymbol{\epsilon}^L, \end{aligned}$$

where, in the last line, a first-order approximation to the BCH formula has been used.

The elements of the Jacobian $\partial \mathbf{e}(\mathbf{T}(\boldsymbol{\epsilon}^L))/\partial \boldsymbol{\epsilon}^L$ were also calculated using (6) with a step size of $h = 10^{-20}$. An optimization was performed with both Jacobian calculation methods, where the Gauss-Newton step $\delta \boldsymbol{\epsilon}_{\ell-1}$ is determined from

$$\delta \boldsymbol{\epsilon}_{\ell-1} = \left[\left(\frac{\partial \mathbf{e}}{\partial \boldsymbol{\epsilon}^L} \right)^\top \mathbf{W} \left(\frac{\partial \mathbf{e}}{\partial \boldsymbol{\epsilon}^L} \right) \right]^{-1} \left[- \left(\frac{\partial \mathbf{e}}{\partial \boldsymbol{\epsilon}^L} \right)^\top \mathbf{W} \mathbf{e}(\bar{\mathbf{T}}) \right],$$

and the argument of $\mathbf{e}(\mathbf{T}(\boldsymbol{\epsilon}^L))$ is dropped for conciseness. The point is updated by

$$\bar{\mathbf{T}}_\ell = \exp(\delta \boldsymbol{\epsilon}_{\ell-1}^\wedge) \bar{\mathbf{T}}_{\ell-1}.$$

As shown in Fig. 2 using both an analytic Jacobian or a complex-step Jacobian results in an optimum being reached in a single step. Note that calculating Jacobians using the complex-step is shown to have a minor improvement in cost function reduction as compared to the analytical method. The reason is that the analytical method uses a first-order BCH approximation, which is ultimately slightly less accurate than the machine-precision complex-step Jacobian calculations.

IV. Batch Estimation

The methodology of Example 2 is now applied to a practical state estimation problem. Consider the task of estimating the position and attitude of a rigid body at different points in time t_0, t_1, \dots, t_K using various measurements. The state of the rigid body at a discrete point in time t_k can be represented by the matrix Lie group element $\mathbf{T}_k \in \mathcal{G}$, where \mathcal{G} will depend on the estimation task.

A. Maximum A Posteriori Estimation

The MAP approach [Barfoot2019] to estimate the states in a batch framework results in the minimization of the least-squares cost function shown in (8), where the errors to be minimized are

$$\mathbf{e}(\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_K) = \begin{bmatrix} \mathbf{e}_{u,0} \\ \mathbf{e}_{u,1} \\ \vdots \\ \mathbf{e}_{u,K} \\ \mathbf{e}_{y,0} \\ \vdots \\ \mathbf{e}_{y,K} \end{bmatrix}.$$

The error term $\mathbf{e}_{u,0}$ represents an error between the known initial state $\bar{\mathbf{T}}_0$, with uncertainty, and the estimated initial state \mathbf{T}_0 . This term is computed as

$$\mathbf{e}_{u,0} = \ln(\mathbf{T}_0^{-1}\bar{\mathbf{T}}_0)^\vee.$$

The process error terms $\mathbf{e}_{u,1}, \dots, \mathbf{e}_{u,K}$ are a function of a discrete-time process model of the form $\mathbf{T}_k = \mathbf{F}(\mathbf{T}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1})$ where \mathbf{u}_{k-1} and \mathbf{w}_{k-1} are the input and zero-mean process noise at time t_{k-1} , respectively. These error terms are calculated as

$$\mathbf{e}_{u,k} = \ln(\mathbf{T}_k^{-1}\mathbf{F}(\mathbf{T}_{k-1}, \mathbf{u}_{k-1}, \mathbf{0}))^\vee.$$

Finally, the terms $\mathbf{e}_{y,0}, \dots, \mathbf{e}_{y,K}$ correspond to the errors between measurements, and a measurement model of the form $\mathbf{y}_k = \mathbf{g}(\mathbf{T}_k, \boldsymbol{\nu}_k)$, where $\boldsymbol{\nu}_k$ is zero-mean measurement noise. Hence, the measurement errors are

$$\mathbf{e}_{y,k} = \mathbf{y}_k - \mathbf{g}(\mathbf{T}_k, \mathbf{0}).$$

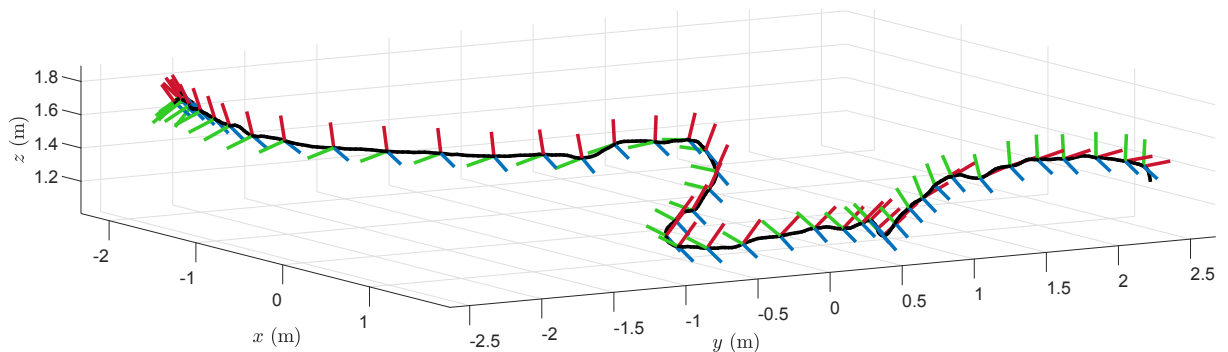


Fig. 3. Trajectory visualization of a batch-estimation solution from the EuRoC Dataset.

Following the MAP formulation the weight in (8) is

$$\mathbf{W} = \text{diag}(\mathbf{P}_0^{-1}, \mathbf{Q}_1^{-1}, \dots, \mathbf{Q}_K^{-1}, \mathbf{R}_0^{-1}, \dots, \mathbf{R}_K^{-1}),$$

where the matrix \mathbf{P}_0 is a covariance matrix associated with the uncertainty in the initial state, \mathbf{T}_0 . The matrices \mathbf{Q}_k and \mathbf{R}_k are covariance matrices associated with the process and measurement noises, respectively.

The goal is to find $\mathbf{T}_0, \dots, \mathbf{T}_K$ that minimize the least-squares cost function given by (8). To use a Gauss-Newton algorithm, the right (or left) Jacobian $\partial \mathbf{e}(\mathbf{T}(\boldsymbol{\epsilon}^R)) / \partial \boldsymbol{\epsilon}^R$ is needed, where $\boldsymbol{\epsilon}^R = [\boldsymbol{\epsilon}_0^{R^T} \dots \boldsymbol{\epsilon}_K^{R^T}]^T$ is a matrix that consists of perturbations to the individual estimated states. Since the error $\mathbf{e}(\mathbf{T}_0, \dots, \mathbf{T}_K)$ is a function of K different Lie group elements, it is worth mentioning a simple technique that allows a user to treat the same function as a function of a single matrix Lie group element, as shown next in Section IV-B. The Jacobians can then be computed using (5) or (6).

B. Recasting $f(\mathbf{X}_0, \dots, \mathbf{X}_K)$ as $f(\mathbf{X})$

Consider a function $f(\mathbf{X}_0, \dots, \mathbf{X}_K) \in \mathbb{R}$ where $\mathbf{X}_0, \dots, \mathbf{X}_K \in \mathcal{G}$. Let $\mathbf{X}_i = \bar{\mathbf{X}}_i \exp(\boldsymbol{\epsilon}_i^{R^\wedge})$. Define $\mathbf{X} \triangleq \text{diag}(\mathbf{X}_0, \dots, \mathbf{X}_K)$, thus leading to

$$\begin{aligned} \mathbf{X} &= \begin{bmatrix} \bar{\mathbf{X}}_0 & & & \\ & \ddots & & \\ & & \bar{\mathbf{X}}_K & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} \exp(\boldsymbol{\epsilon}_0^{R^\wedge}) & & & \\ & \ddots & & \\ & & \exp(\boldsymbol{\epsilon}_K^{R^\wedge}) & \\ & & & \ddots \end{bmatrix} \\ &= \begin{bmatrix} \bar{\mathbf{X}}_0 & & & \\ & \ddots & & \\ & & \bar{\mathbf{X}}_K & \\ & & & \ddots \end{bmatrix} \exp \begin{bmatrix} \boldsymbol{\epsilon}_0^{R^\wedge} & & & \\ & \ddots & & \\ & & \boldsymbol{\epsilon}_K^{R^\wedge} & \\ & & & \ddots \end{bmatrix}. \end{aligned} \quad (9)$$

By defining $\boldsymbol{\epsilon}^R \triangleq [\boldsymbol{\epsilon}_0^{R^T} \dots \boldsymbol{\epsilon}_K^{R^T}]^T$, $\bar{\mathbf{X}} \triangleq \text{diag}(\bar{\mathbf{X}}_0, \dots, \bar{\mathbf{X}}_K)$ along with a new operator $(\cdot)^\Delta$ such that $\boldsymbol{\epsilon}^\Delta \triangleq \text{diag}(\boldsymbol{\epsilon}_0^\wedge, \dots, \boldsymbol{\epsilon}_K^\wedge)$, equation (9) becomes

$$\mathbf{X} = \bar{\mathbf{X}} \exp(\boldsymbol{\epsilon}^{R^\Delta}).$$

Therefore, a collection of matrix Lie group elements can be packaged into a single element of a new group. This can be done similarly with left perturbations.

C. The EuRoC Dataset

The EuRoC micro aerial vehicle dataset collected by the Autonomous Systems Laboratory at ETH Zürich, Switzerland [Burri2016] includes accelerometer and gyroscope measurements, as well as ground truth position data. To simulate position measurements akin to GPS or UWB measurements, normally distributed random noise is added to the provided ground position data. The state of the rigid body can be represented by the matrix Lie group element $\mathbf{T}_k \in SE_2(3)$, and as such the velocity is also estimated. The accelerometer measurements $\mathbf{u}_k^{\text{acc}}$ and gyroscope measurements $\mathbf{u}_k^{\text{gyro}}$ are treated as process-model inputs $\mathbf{u}_k = [\mathbf{u}_k^{\text{acc}^T} \mathbf{u}_k^{\text{gyro}^T}]^T$, while the position measurements $\mathbf{y}_k^{\text{pos}}$ are treated as measurement-model outputs.

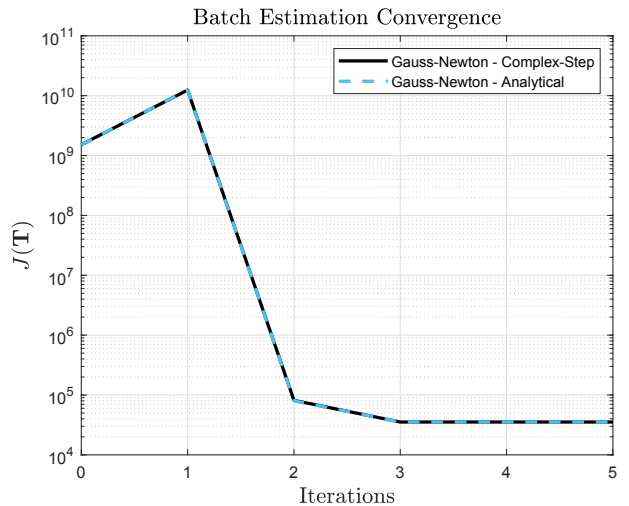


Fig. 4. Convergence history of a Gauss-Newton algorithm on EuRoC Dataset. Virtually identical performance is achieved to the analytical solution, and using a central-difference method. However, the complex-step requires only 1 function evaluation, and no step-size tuning was needed.

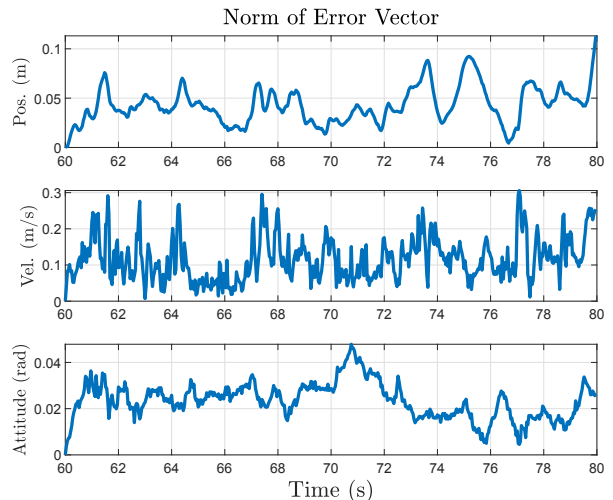


Fig. 5. Magnitude of errors in position, velocity, and attitude resulting from the optimal batch-estimation solution using the complex-step.

For this problem, the analytical expression for the Jacobian $\partial \mathbf{e}(\mathbf{T}) / \partial \boldsymbol{\epsilon}^R$ can be obtained, and the details of the derivation can be found in [Arsenaute2019]. Henceforth, the arguments of functions of multiple matrix Lie group elements will be consolidated under \mathbf{T} , as described in Section IV-B. The right Jacobian is

$$\frac{\partial \mathbf{e}(\mathbf{T})}{\partial \boldsymbol{\epsilon}^R} \approx \begin{bmatrix} -\mathbf{1} & & & \\ \mathbf{F}_0 & \ddots & & \\ & \ddots & \ddots & \\ & & & -\mathbf{1} \\ & & & \mathbf{F}_K \\ \hline \mathbf{H}_0 & & & & \\ & \ddots & & & \\ & & & & \mathbf{H}_K \end{bmatrix},$$

where

$$\begin{aligned} \mathbf{F}_k &= \text{Ad}(\mathbf{T}_k^{-1} \mathbf{F}_{k-1}^{\text{op}}) \mathbf{B}, \\ \mathbf{F}_{k-1}^{\text{op}} &= \begin{bmatrix} \mathbf{C}_{k-1} & \mathbf{v}_{k-1} + T\mathbf{g} & \mathbf{r}_{k-1} + T\mathbf{v}_{k-1} \\ & 1 & \\ & & 1 \end{bmatrix}, \\ \mathbf{B} &= \begin{bmatrix} \mathbf{1} & & \\ & \mathbf{1} & \\ & T\mathbf{1} & \mathbf{1} \end{bmatrix}, \\ \mathbf{H}_k &= [\mathbf{1} \ \mathbf{0} \ \mathbf{0}] \mathbf{T}_k \mathbf{p}^\odot, \end{aligned}$$

where $T = t_k - t_{k-1}$, $\mathbf{p} = [\mathbf{0} \ 1]^\top$, and \mathbf{g} is the gravity vector resolved in the datum frame. These expressions require first-order approximations to the BCH formula, similar to Example 2. This is common procedure, as the approximation becomes more accurate as errors become small [Arsenault2019, Barfoot2019].

A Gauss-Newton optimization is performed on the MH_03_medium dataset. For simplicity, the accelerometer and gyroscope measurements are downsampled from the original 200 Hz in order to reduce the amount of variables in the optimization procedure. An alternative to downsampling is to perform IMU preintegration as described in [Forster2017], but this is beyond the scope of this paper. The specifications of the batch-estimation problem are shown in Table I. The process covariance matrix was set to,

$$\mathbf{Q}_k = \text{diag}(1.6 \cdot 10^{-7} \cdot \mathbf{1}, 2 \cdot 10^{-6} \cdot \mathbf{1}, 10^{-10} \mathbf{1}). \quad (10)$$

TABLE I
EuRoC Estimation Scenario Specifications

Specification	Value	Units
Accelerometer meas. freq.	25	Hz
Gyroscope meas. freq.	25	Hz
Position meas. freq.	10	Hz
Data time span	60 - 80	s
Number of states estimated	500	-
Std. deviation of position meas.	0.1	m
Initial state guess covariance \mathbf{P}_0	$10^{-10} \cdot \mathbf{1}$	$[\text{rad}^2, (\text{m/s})^2, \text{m}^2]$
Process covariance \mathbf{Q}_k	See eqn. (10)	$[\text{rad}^2, (\text{m/s})^2, \text{m}^2]$
Measurement covariance \mathbf{R}_k	$0.1^2 \cdot \mathbf{1}$	m^2
Complex-step der. step size h	10^{-20}	-

The initial state, $\check{\mathbf{T}}_0$, is set to the ground truth, and hence the diagonal of \mathbf{P}_0 is given arbitrarily small numbers. The matrix \mathbf{Q}_k was further tuned to yield better performance, after obtaining the nominal noise values provided in the EuRoC dataset. Using the initial state, the process model is directly integrated using the accelerometer and gyroscope measurements, which then provides an initial guess for the poses at all the discrete time points. This dead reckoning solution is then used to initialize the Gauss-Newton algorithm.

Figure 3 shows a visualization of the trajectory once the optimization procedure has converged. Figure 4 shows the value of the cost function $J(\mathbf{T})$ across the iterations of the Gauss-Newton algorithm. Since the



Fig. 6. Experimental setup of the ‘Lost in the Woods’ dataset, courtesy of [Barfoot2009]. Truth measurements are obtained from a motion capture system.

initial guess for the states is obtained by dead reckoning, this sets all the process errors $\mathbf{e}_{u,1}, \dots, \mathbf{e}_{u,K}$ to zero. The first iteration attempts to decrease the measurement errors, resulting in an increase in process errors, and hence an increase in the overall cost function.

In this example, BCH approximations in the analytical Jacobians did not create any difference in the convergence history since the errors are initialized to be small in the the dead reckoning step. A central-difference scheme was also used to calculate Jacobians, and after multiple trial-and-error attempts with different step sizes, an identical convergence history to what is shown in Fig. 4 was obtained. However, the central-difference method requires twice as many function evaluations as the complex-step method, and therefore required approximately twice the total computing time. Finally, Fig. 5 shows the 2-norm of the difference between the batch-estimation solution and the ground truth. The errors are small, indicating the MAP framework has converged close to the ground truth.

D. The ‘Lost in the Woods’ Dataset

The ‘Lost in the Woods’ dataset consists of a mobile wheeled robot navigating through a “forest” of tubes [Barfoot2009], as seen in Figure 6. The robot is equipped with wheel odometry providing forward velocity measurements, denoted u_k^{vel} , and angular velocity measurements, denoted u_k^{ang} . Furthermore, the robot has a laser range finder that provides range and bearing measurements to pre-identified landmarks (the tubes shown in Figure 6), denoted r_k^ℓ, ϕ_k^ℓ for landmark ℓ at t_k , respectively. The positions of the landmarks in a datum reference frame are known in advance, and are denoted \mathbf{r}^ℓ . The state of the robot can be represented by $\mathbf{T}_k \in SE(2)$.

The process model consists of the nonholonomic vehicle kinematics. Written in the form $\mathbf{T}_k = \mathbf{F}(\mathbf{T}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1})$,

$$\mathbf{T}_k = \mathbf{T}_{k-1} \Psi_{k-1},$$

TABLE II

‘Lost in the Woods’ Estimation Scenario Specifications

Specification	Value	Units
Wheel odometry freq.	5	Hz
Laser range finder freq.	5	Hz
Data time span	500 - 620	s
Number of states estimated	600	-
Initial state guess covariance \mathbf{P}_0	$\mathbf{1}$	[rad ² , m ² , m ²]
Complex-step der. step size h	10^{-20}	-

where

$$\Psi_{k-1} = \begin{bmatrix} \exp(T(u_{k-1}^{\text{ang}} + w_{k-1}^{\text{ang}})^{\wedge}) & T(u_{k-1}^{\text{vel}} + w_{k-1}^{\text{vel}})\mathbf{1}_1 \\ \mathbf{0} & 1 \end{bmatrix},$$

$T = t_k - t_{k-1}$, and $w_{k-1}^{\text{vel}}, w_{k-1}^{\text{ang}}$ are zero-mean normally distributed noises associated with the velocity and angular velocity measurements, respectively. The measurement model consists of the range and bearing measurements for each landmark. Written as $\mathbf{y} = \mathbf{g}(\mathbf{T}_k, \boldsymbol{\nu}_k)$, the measurement model is

$$\begin{bmatrix} r_k^\ell \\ \phi_k^\ell \end{bmatrix} = \begin{bmatrix} \sqrt{(\mathbf{r}^\ell - \mathbf{D}\mathbf{T}_k\mathbf{p})^\top (\mathbf{r}^\ell - \mathbf{D}\mathbf{T}_k\mathbf{p})} \\ (\text{atan2}(\mathbf{1}_2^\top (\mathbf{r}^\ell - \mathbf{D}\mathbf{T}_k\mathbf{p}), \mathbf{1}_1^\top (\mathbf{r}^\ell - \mathbf{D}\mathbf{T}_k\mathbf{p}))) \\ -\mathbf{1}_1^\top \ln(\mathbf{T}_k)^\vee \end{bmatrix} + \boldsymbol{\nu}_k,$$

where $\boldsymbol{\nu}_k$ is zero-mean normally distributed measurement noise, $\mathbf{D} = [\mathbf{1} \ \mathbf{0}]$, $\mathbf{p} = [d \ 0 \ 1]^\top$, and d is the distance between the laser range finder and the reference point on the robot.

Computing the Jacobians associated with the measurement model by hand is, although not impossible, laborious due to the $\text{atan2}(\cdot, \cdot)$ term. Hence, the complex-step derivative is used to directly evaluate the right Jacobian $\partial \mathbf{e}(\mathbf{T}) / \partial \boldsymbol{\epsilon}^R$ for use in the Gauss-Newton optimization.

Dead reckoning was performed using wheel odometry in order to generate an initial guess for the Gauss-Newton optimization. All measurements were downsampled from the original 10 Hz to 5 Hz in order to limit the number of variables in the optimization procedure. The \mathbf{Q}_k and \mathbf{R}_k matrices were directly formed from the discrete-time covariances provided in the dataset [Barfoot2009]. The initial state \mathbf{T}_0 was set to be a random perturbation from the ground truth.

The algorithm converged in 6 iterations, and produced a trajectory visualizable in Figure 7. The errors are shown in Figure 8, which show good performance when compared to the ground truth position and attitude data. This can also be achieved with central-difference, but again, the computation time is significantly longer, and the step sized must be tuned.

V. Conclusion

This paper has shown that the complex-step derivative can successfully be used to obtain Jacobians of functions that have matrix Lie group elements as arguments. Machine-precision can be achieved with a single complex function evaluation. To use the complex-step, functions must be programmed to accept complex numbers, which

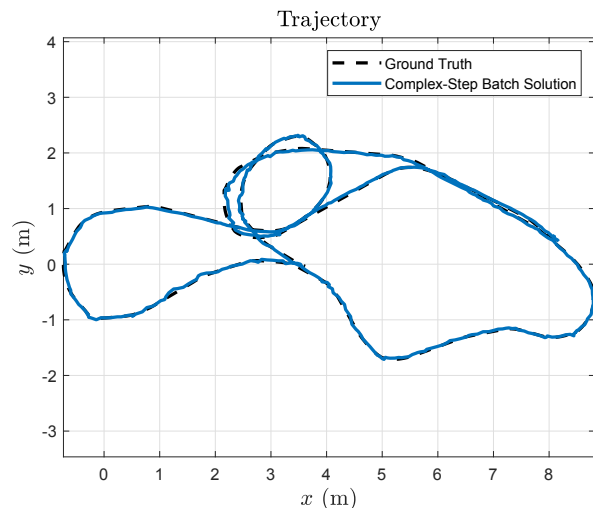


Fig. 7. 2D trajectory trace for the ‘Lost in the Woods’ dataset. The solution using the complex-step derivative shows excellent agreement with the ground truth data.

is occasionally time consuming. In Matlab, it is critical to use the (\cdot) transpose operator as opposed to the $(\cdot)'$ conjugate transpose, and also to redefine the $\text{abs}()$, $\text{max}()$, and $\text{min}()$ functions. A guide to proper implementation in various other programming languages can be found in [MDOUMich2017].

There is a multitude of other potential applications for this tool, such as numerical linearization of high-fidelity dynamics models, real-time state estimation and Kalman filtering [Vittaldev2012], and the training of matrix Lie group-based neural networks [Peretroukhin2018]. For second derivatives, the complex-step is unfortunately unable to realize machine-precision accuracy. However, methods are available to improve the accuracy [Lai2008], which are likely extendible to matrix Lie groups. Furthermore, if an analytical Jacobian is known, the Hessian can be determined with machine precision using the complex-step [Martins2003].

Appendix

A. The Special Euclidean Group $SE(2)$

The group $SE(2)$ is defined as [Chirikjian2009]

$$SE(2) = \left\{ \mathbf{T} = \begin{bmatrix} \mathbf{C} & \mathbf{r} \\ \mathbf{0} & 1 \end{bmatrix} \in \mathbb{R}^{3 \times 3} \mid \mathbf{C} \in SO(2), \mathbf{r} \in \mathbb{R}^2 \right\},$$

where $SO(n)$ refers to the Special Orthogonal Group consisting of orthonormal matrices with unit determinant. The matrix Lie algebra associated with $SE(2)$ is

$$\mathfrak{se}(2) = \{ \boldsymbol{\Xi} = \boldsymbol{\xi}^\wedge \in \mathbb{R}^{3 \times 3} \mid \boldsymbol{\xi} \in \mathbb{R}^3 \},$$

where

$$\boldsymbol{\xi}^\wedge = \begin{bmatrix} \xi^\phi \\ \xi_1^r \\ \xi_2^r \end{bmatrix}^\wedge = \begin{bmatrix} 0 & -\xi^\phi & \xi_1^r \\ \xi^\phi & 0 & \xi_2^r \\ 0 & 0 & 0 \end{bmatrix}.$$

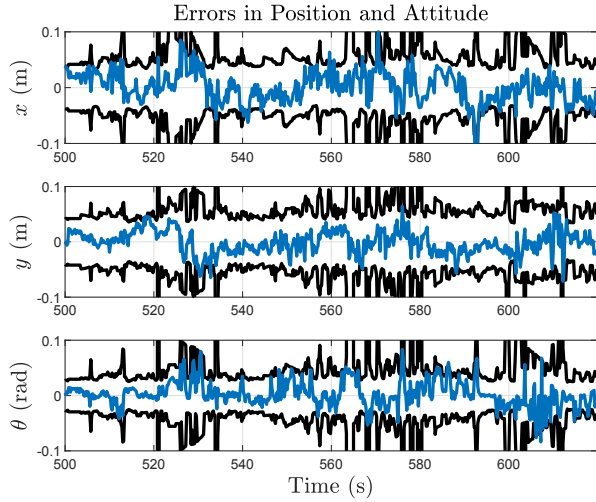


Fig. 8. Error in position x, y and attitude θ between estimated solution and ground truth (blue), along with ± 3 standard deviation bounds (black). There is less than 10 cm of position error, and less than 0.1 rad of attitude error.

The closed-form expression for the exponential map $\exp : \mathfrak{se}(2) \rightarrow SE(2)$ is

$$\exp(\xi^\wedge) = \begin{bmatrix} \mathbf{C} & \mathbf{J}_\ell \xi^r \\ \mathbf{0} & 1 \end{bmatrix},$$

where $\xi^r = [\xi_1^r \ \xi_2^r]^\top$ and

$$\mathbf{J}_\ell = \frac{1}{\xi^\phi} \begin{bmatrix} \sin(\xi^\phi) & -(1 - \cos(\xi^\phi)) \\ (1 - \cos(\xi^\phi)) & \sin(\xi^\phi) \end{bmatrix}.$$

B. The Special Euclidean Group $SE(3)$

The matrix Lie group $SE(3)$ is defined as [Barfoot2019]

$$SE(3) = \left\{ \mathbf{T} = \begin{bmatrix} \mathbf{C} & \mathbf{r} \\ \mathbf{0} & 1 \end{bmatrix} \in \mathbb{R}^{4 \times 4} \mid \mathbf{C} \in SO(3), \mathbf{r} \in \mathbb{R}^3 \right\}.$$

The matrix Lie algebra associated with $SE(3)$ is

$$\mathfrak{se}(3) = \{ \Xi = \xi^\wedge \in \mathbb{R}^{4 \times 4} \mid \xi \in \mathbb{R}^6 \},$$

where

$$\xi^\wedge = \begin{bmatrix} \xi^\phi \\ \xi^r \end{bmatrix}^\wedge = \begin{bmatrix} \xi^{\phi^\times} & \xi^r \\ \mathbf{0} & 0 \end{bmatrix}, \quad \xi^\phi, \xi^r \in \mathbb{R}^3,$$

and

$$\xi^{\phi^\times} = \begin{bmatrix} \xi_1^\phi \\ \xi_2^\phi \\ \xi_3^\phi \end{bmatrix}^\times = \begin{bmatrix} 0 & -\xi_3^\phi & \xi_2^\phi \\ \xi_3^\phi & 0 & -\xi_1^\phi \\ -\xi_2^\phi & \xi_1^\phi & 0 \end{bmatrix}.$$

The closed-form expression for the exponential map $\exp : \mathfrak{se}(3) \rightarrow SE(3)$ is

$$\exp(\xi^\wedge) = \begin{bmatrix} \exp(\xi^{\phi^\times}) & \mathbf{J}_\ell \xi^r \\ \mathbf{0} & 1 \end{bmatrix},$$

where

$$\mathbf{J}_\ell = \frac{\sin(\phi)}{\phi} \mathbf{1} + \left(1 - \frac{\sin(\phi)}{\phi} \right) \mathbf{a} \mathbf{a}^\top + \frac{1 - \cos(\phi)}{\phi} \mathbf{a}^\times,$$

$$\exp(\xi^{\phi^\times}) = \cos(\phi) \mathbf{1} + (1 - \cos(\phi)) \mathbf{a} \mathbf{a}^\top + \sin(\phi) \mathbf{a}^\times,$$

and $\phi = \|\xi^\phi\|$ and $\mathbf{a} = \xi^\phi / \phi$. The matrix \mathbf{J}_ℓ is known as the left Jacobian of the group $SO(3)$. It is also useful to define the operator [Barfoot2019]

$$\mathbf{p}^\odot = \begin{bmatrix} \varepsilon \\ \eta \end{bmatrix}^\odot = \begin{bmatrix} -\varepsilon^\times & \eta \mathbf{1} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \varepsilon \in \mathbb{R}^3, \eta \in \mathbb{R},$$

such that $\mathbf{x}^\wedge \mathbf{p} = \mathbf{p}^\odot \mathbf{x}$ holds.

C. The Group of Double Direct Isometries $SE_2(3)$

The matrix Lie group $SE_2(3)$ is defined as

$$SE_2(3) = \left\{ \mathbf{T} = \begin{bmatrix} \mathbf{C} & \mathbf{v} & \mathbf{r} \\ \mathbf{0} & 1 & 0 \\ \mathbf{0} & 0 & 1 \end{bmatrix} \mid \mathbf{C} \in SO(3), \mathbf{v}, \mathbf{r} \in \mathbb{R}^3 \right\}.$$

The matrix Lie algebra associated with $SE_2(3)$ is

$$\mathfrak{se}_2(3) = \{ \Xi = \xi^\wedge \in \mathbb{R}^{5 \times 5} \mid \xi \in \mathbb{R}^9 \},$$

where

$$\xi^\wedge = \begin{bmatrix} \xi^\phi \\ \xi^v \\ \xi^r \end{bmatrix}^\wedge = \begin{bmatrix} \xi^{\phi^\times} & \xi^v & \xi^r \\ \mathbf{0} & 0 & 0 \\ \mathbf{0} & 0 & 0 \end{bmatrix}, \quad \xi^\phi, \xi^v, \xi^r \in \mathbb{R}^3.$$

The closed-form expression for the exponential map $\exp : \mathfrak{se}_2(3) \rightarrow SE_2(3)$ is

$$\exp(\xi^\wedge) = \begin{bmatrix} \exp(\xi^{\phi^\times}) & \mathbf{J}_\ell \xi^v & \mathbf{J}_\ell \xi^r \\ \mathbf{0} & 1 & 0 \\ \mathbf{0} & 0 & 1 \end{bmatrix}.$$

It is also useful to define the operator

$$\mathbf{p}^\odot = \begin{bmatrix} \varepsilon \\ \eta_1 \\ \eta_2 \end{bmatrix}^\odot = \begin{bmatrix} -\varepsilon^\times & \eta_1 \mathbf{1} & \eta_2 \mathbf{1} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

where $\varepsilon \in \mathbb{R}^3$ and $\eta_1, \eta_2 \in \mathbb{R}$, such that $\mathbf{x}^\wedge \mathbf{p} = \mathbf{p}^\odot \mathbf{x}$ holds.

Acknowledgment

The authors would like to thank Jonathan Arsenault and Thomas Hitchcox for their many helpful discussions.