

Invariant Kalman Filtering

Axel Barrau¹ and Silvère Bonnabel²¹Safran Tech, Groupe Safran, 78772 Magny Les Hameaux CEDEX, France;
email: axel.barrau@safrangroup.com²MINES ParisTech, PSL Research University, Centre for Robotics, 75006 Paris,
France; email: silvere.bonnabel@mines-paristech.frAnnu. Rev. Control Robot. Auton. Syst. 2018.
1:237–57First published as a Review in Advance on
December 20, 2017The *Annual Review of Control, Robotics, and
Autonomous Systems* is online at
control.annualreviews.org<https://doi.org/10.1146/annurev-control-060117-105010>Copyright © 2018 by Annual Reviews.
All rights reserved

Keywords

Kalman filter, geometry, Lie groups, localization, autonomous navigation, simultaneous localization and mapping, SLAM

Abstract

The Kalman filter—or, more precisely, the extended Kalman filter (EKF)—is a fundamental engineering tool that is pervasively used in control and robotics and for various estimation tasks in autonomous systems. The recently developed field of invariant extended Kalman filtering uses the geometric structure of the state space and the dynamics to improve the EKF, notably in terms of mathematical guarantees. The methodology essentially applies in the fields of localization, navigation, and simultaneous localization and mapping (SLAM). Although it was created only recently, its remarkable robustness properties have already motivated a real industrial implementation in the aerospace field. This review aims to provide an accessible introduction to the methodology of invariant Kalman filtering and to allow readers to gain insight into the relevance of the method as well as its important differences with the conventional EKF. This should be of interest to readers intrigued by the practical application of mathematical theories and those interested in finding robust, simple-to-implement filters for localization, navigation, and SLAM, notably for autonomous vehicle guidance.

**ANNUAL
REVIEWS Further**Click here to view this article's
online features:

- Download figures as PPT slides
- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

1. INTRODUCTION

The goal of a filter is to estimate the state of a dynamical system by combining an evolution model and some sensor measurements that bring partial information about the state. Unfortunately, models are inherently inaccurate, and sensors are subject to noise that corrupts the measurements. The idea of filtering is to explicitly include both sources of uncertainty in the model and compute the best estimates of the state that can be inferred from the available information.

Even though the mathematical theory is now well understood, it is still challenging to design filters in practice for control, robotics, and autonomous systems. The extended Kalman filter (EKF) appeared in the 1960s with the advent of computers and was first implemented by NASA in the Apollo program to estimate the trajectory of the space capsule in real time. The past two decades have also witnessed the development of particle filters, with great advances in both theory and practice. However, these filters rely on extensive numerical computations that are not always suited to real-time onboard implementations, and there are theoretical caveats, especially when process noise is low (typically for static parameter estimation). The robotics community has also more recently turned to optimization-based techniques for filtering (the problem is typically formulated as nonlinear least squares), but the computational demands are extensive, and their robustness to erroneous initializations is not yet clearly established. Because robots and many control systems are real-time systems, the amount of computation is limited, and the EKF is still a widespread tool in control and robotics, along with its more recent variant, the unscented Kalman filter (UKF). In the aerospace industry, the EKF remains the reference filter.

1.1. Extended Kalman Filtering

Consider a general dynamical system in discrete time whose state is described by the vector variable $X_n \in \mathbb{R}^d$. We associate a sequence of observations $(Y_n)_{n \geq 0} \in \mathbb{R}^p$, which are measurement data returned by sensors. The trusted evolution model is

$$X_n = f(X_{n-1}, u_n, w_n), \quad 1.$$

where f is the function encoding the evolution of the system; w_n is the (unknown) process noise, which is a centered random variable with covariance matrix Q_n ; the vector $u_n \in \mathbb{R}^m$ is a control input; and the observation consists of partial measurements of the state at time n :

$$Y_n = b(X_n) + V_n, \quad 2.$$

with b the observation function and V_n the (unknown) measurement noise that accounts for sensors' limitations.

The EKF computes in real time an approximation $\hat{X}_{n|n}$ to the best state estimate given the observations. Let $Y_{1:n}$ denote the collection of past measurements Y_1, Y_2, \dots, Y_n , and let $u_{1:n}$ be similarly defined. To be more precise, the EKF represents the belief $\mathbb{P}(X_n | u_{1:n}, Y_{1:n})$, which assigns a probability to each possible value of the true state in light of all the information collected so far, by a mean $\hat{X}_{n|n}$ and covariance matrix $P_{n|n}$. Indeed, the rationale is to use the following Gaussian approximation: $\mathbb{P}(X_n | u_{1:n}, Y_{1:n}) \approx \mathcal{N}(\hat{X}_{n|n}, P_{n|n})$. To compute the mean and covariance, the EKF uses a two-step procedure:

1. *Propagation*: The estimate $\hat{X}_{n-1|n-1}$ obtained after the observation Y_{n-1} is propagated through the deterministic part of Equation 1:

$$\hat{X}_{n|n-1} = f(\hat{X}_{n-1|n-1}, u_n, 0). \quad 3.$$

To compute the associated covariance, introduce the estimation errors defined as

$$e_{n-1|n-1} = X_{n-1} - \hat{X}_{n-1|n-1}, \quad e_{n|n-1} = X_n - \hat{X}_{n|n-1}. \quad 4.$$

The key idea underlying the EKF is to linearize the error system through a first-order Taylor expansion of the nonlinear functions f and b at the estimate $\hat{X}_{n-1|n-1}$. Indeed, using the Jacobians $F_n = \frac{\partial f}{\partial X}(\hat{X}_{n-1|n-1}, u_n, 0)$, $G_n = \frac{\partial f}{\partial w}(\hat{X}_{n-1|n-1}, u_n, 0)$, and $H_n = \frac{\partial b}{\partial X}(\hat{X}_{n|n-1})$, the combination of Equations 1–3 yields the following first-order expansion of the error system:

$$e_{n|n-1} = F_n e_{n-1|n-1} + G_n w_n, \quad 5.$$

$$Y_n - b(\hat{X}_{n|n-1}) = H_n e_{n|n-1} + V_n, \quad 6.$$

where the second-order terms—that is, terms of order $O(\|e\|^2, \|w\|^2, \|e\|\|w\|)$ —have been removed (see, e.g., Reference 1). $P_{n-1|n-1}$ is an approximation to the true covariance $\mathbb{E}(e_{n-1|n-1} e_{n-1|n-1}^T)$, and it is propagated through the linearized model Equation 5 so that $P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + G_n Q_n G_n^T$ is an approximation of $\mathbb{E}(e_{n|n-1} e_{n|n-1}^T)$, and we have $\mathbb{P}(X_n | u_{1:n}, Y_{1:n-1}) \approx \mathcal{N}(\hat{X}_{n|n-1}, P_{n|n-1})$.

2. *Measurement update:* To account for the measurement Y_n , we let $z_n = Y_n - b(\hat{X}_{n|n-1})$, and z_n is called the innovation. Assuming that $e_{n|n-1} \sim \mathcal{N}(0, P_{n|n-1})$ and that the approximation from Equation 6 is exact, the linear Kalman filter equations ensure that the updated error $e_{n|n} = X_n - \hat{X}_{n|n}$ satisfies $e_{n|n} \sim \mathcal{N}(0, P_{n|n})$, where

$$\hat{X}_{n|n} = \hat{X}_{n|n-1} + K_n z_n \quad \text{and} \quad P_{n|n} = [I - K_n H_n] P_{n|n-1}, \quad 7.$$

with K_n , called the Kalman gain, computed in Algorithm 1. Of course, the belief after update $\mathcal{N}(\hat{X}_{n|n}, P_{n|n})$ is only an approximation to $\mathbb{P}(X_n | u_{1:n}, Y_{1:n})$ owing to the linearizations. In practice, those linearizations may lead the filter to inconsistencies and sometimes even divergence.

Algorithm 1 (extended Kalman filter). Choose an initial estimate $\hat{X}_{0|0}$ and uncertainty matrix $P_{0|0}$.

loop

Define F_n , G_n , and H_n through Equations 5 and 6

Define Q_n as $\text{Cov}(w_n)$ and R_n as $\text{Cov}(V_n)$

Propagation

$$\hat{X}_{n|n-1} = f(\hat{X}_{n-1|n-1}, u_n, 0)$$

$$P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + G_n Q_n G_n^T$$

Measurement update

Compute $z_n = Y_n - b(\hat{X}_{n|n-1})$, $S_n = H_n P_{n|n-1} H_n^T + R_n$, and

$$K_n = P_{n|n-1} H_n^T S_n^{-1}$$

$$P_{n|n} = [I - K_n H_n] P_{n|n-1}$$

$$\hat{X}_{n|n} = \hat{X}_{n|n-1} + K_n z_n$$

end loop

1.2. Motivation for the Use of Geometry

Users who are facing the filtering problem defined by the system of Equations 1 and 2, with f and b two nonlinear maps, are free to choose a different coordinate system to design an EKF. For instance, in radar tracking, one can choose not only a frame attached to the target or to the radar, but also range and bearing as an alternative to Cartesian coordinates. In general, it is unclear what the best coordinates for EKF design are. However, some systems have a natural choice of

coordinates. For instance, when facing a linear Gaussian system, the EKF boils down to the linear Kalman filter, which is optimal, and it would be nonsense to work with alternative coordinates that would make the system nonlinear. In this article, we advocate that, for a large class of systems defined on matrix Lie groups, the machinery of geometry provides coordinates that are unarguably more suited to the problem. In those cases, the invariant extended Kalman filter (IEKF) theory is useful, as the original problem is often formulated using coordinates that do not match the group structure, leading to degraded performance of the conventional EKF.

For systems on Lie groups, the IEKF was originally introduced in Reference 2 and continued in References 3–7. Reference 7 recently described the complete methodology along with the convergence properties of the IEKF. More generally, the use of Lie groups for state estimation dates back to the 1970s (8–10) and has recently spanned a range of applications and a rich stream of theoretical results (see, e.g., 6, 11–16).

Following the work of, among others, Chirikjian (17), the robotics community has increasingly recognized that using probability distributions properly defined on Lie groups is important, notably for pose estimation (see, e.g., 17–22). Moreover, Reference 23 recently showed that the use of the IEKF over the conventional EKF solves the inconsistency issues of EKF-based simultaneous localization and mapping (SLAM).

1.3. Outline

Section 2 consists of geometry preliminaries. Section 3 reviews the methodology of invariant Kalman filtering. Section 4 is concerned with the mathematical guarantees that come with the IEKF. Section 5 presents some real industrial applications in the field of inertial navigation. Section 6 reviews the inconsistency of EKF-based SLAM and the interest of IEKF-based SLAM. The presentation of the present article is freely inspired by a tutorial article that uses methods rooted in differential geometry to improve Monte Carlo schemes (24).

2. LIE GROUPS AND PROBABILITY

2.1. Matrix Lie Groups¹

In this section, we provide readers with the bare minimum of Lie group theory that is required to introduce the IEKF methodology. A matrix Lie group G is a subset of square invertible $N \times N$ matrices $\mathcal{M}_N(\mathbb{R})$ verifying the following properties:

$$I_N \in G, \quad \forall \chi \in G, \chi^{-1} \in G, \quad \forall \chi_1, \chi_2 \in G, \chi_1 \chi_2 \in G,$$

where I_N is the identity matrix of \mathbb{R}^N . The subset G is generally not a vector space and can thus be viewed as a curved space (see **Figure 1**). To every point $\chi \in G$, one can associate a vector space $T_\chi G$ that is called the tangent space at χ and is defined as all the matrices that write $\frac{d}{dt}\gamma(0)$, where $\gamma: \mathbb{R} \rightarrow G$ is a smooth curve of G that satisfies $\gamma(0) = \chi$. The elements of this space are called tangent vectors.

The tangent space $T_{I_N} G$ at the identity I_N is called the Lie algebra and plays a very specific role. It is denoted \mathfrak{g} , and its dimension d defines a dimension for the group G itself, where d is generally much smaller than the dimension N^2 of the ambient space. We have $\mathfrak{g} \subset \mathcal{M}_N(\mathbb{R})$,

¹All of the results carry over to general abstract Lie groups, but the (less general) matrix Lie groups are well suited to tutorial and computational purposes and encompass all of the applications discussed.




Figure 1

G is a curved space. Left and right multiplications offer two ways to identify the tangent space $T_\chi G$ at χ with the tangent space at identity $T_{I_N} G$, called the Lie algebra \mathfrak{g} . In turn, the application $\xi \mapsto \xi^\wedge$ provides a linear bijection between the Euclidean space \mathbb{R}^d and \mathfrak{g} . (Lie groups are homogeneous spaces, which somehow “look the same everywhere.” As such, the figure may be slightly misleading, since the curved surface representing G seems irregular. Yet it seems to us that representing a Lie group by, e.g., a sphere would be an oversimplification and to some extent equally misleading.)

but there is always an invertible map $\mathbb{R}^d \rightarrow \mathfrak{g}$ that allows identifying \mathfrak{g} to \mathbb{R}^d . For $\xi \in \mathbb{R}^d$, we denote $\xi^\wedge \in \mathfrak{g}$ the corresponding element of \mathfrak{g} , and we recall that $\xi \mapsto \xi^\wedge$ is a linear map. There are then two canonical ways to identify \mathbb{R}^d and the tangent space $T_\chi G$ at any $\chi \in G$: through left and right multiplications, which are generally different. Indeed, for any $\xi \in \mathbb{R}^d$, the vectors $\chi(\xi^\wedge)$ and $(\xi^\wedge)\chi$ are both tangent vectors at χ . Of course, they are generally different owing to the noncommutativity of matrix multiplication.

The usual matrix exponential map $\exp_m : \mathfrak{g} \rightarrow G$ constitutes a bijection from a neighborhood $V \subset \mathfrak{g}$ of 0 to a neighborhood of the identity I_N in G . In this article, we call the Lie exponential map the map $\exp : \mathbb{R}^d \rightarrow G$, defined by $\exp(\xi) = \exp_m(\xi^\wedge)$, which is a bijection in a neighborhood of $0 \in \mathbb{R}^d$, with $\exp(\xi^\wedge)^{-1} = \exp(-\xi^\wedge)$. Moreover, $\chi \exp(\cdot)$ and $\exp(\cdot)\chi$ provide two distinct bijections between a neighborhood of 0 in \mathbb{R}^d and a neighborhood of χ in G .

The Baker-Campbell-Hausdorff (BCH) formula gives a series expansion for the image in \mathfrak{g} of the product on G : $\text{BCH}(\xi, \zeta) = \exp^{-1}(\exp(\xi)\exp(\zeta))$. In particular, it ensures $\exp(\xi)\exp(\zeta) = \exp(\xi + \zeta + T)$, where T is of the order $O(\|\xi\|^2, \|\zeta\|^2, \|\xi\|\|\zeta\|)$.

Example 1. Consider the group of rotation matrices $G = \text{SO}(3)$, which describes the orientation (attitude) of a body in space. It is the subset of matrices R of $\mathcal{M}_3(\mathbb{R})$ such that $RR^T = I_3$ and $\det(R) = 1$. Each $R \in G$ can be viewed as the rotation that maps vectors expressed in the body frame to vectors expressed in the earth-fixed frame. We have $d = 3$, and for any $\xi \in \mathbb{R}^3$, $\xi^\wedge \in \mathbb{R}^{3 \times 3}$ is the skew symmetric matrix associated with the cross-product operator $\zeta \rightarrow \xi \times \zeta$. The Lie algebra is the set of skew symmetric matrices. Any tangent vector $U \in T_R G$ at R represents an infinitesimal shift of R in $\text{SO}(3)$. It can be written as $U = R\xi^\wedge$ for some vector $\xi \in \mathbb{R}^3$ or, alternatively, as $U = \zeta^\wedge R$ for some $\zeta \in \mathbb{R}^3$. Both ξ and ζ can be seen as angular velocity vectors that represent the same infinitesimal rotation, but ξ is a vector of the body frame, whereas ζ is expressed in the earth-fixed frame. It is easy to prove that indeed $\zeta = R\xi$.

2.2. Uncertainty Representation on Matrix Lie Groups

To define random variables on Lie groups, we cannot apply the usual approach of additive noise for $\chi \in G$ because G is not a vector space. By contrast, we define the probability distribution

$\chi \sim \mathcal{N}_L(\bar{\chi}, P)$ as the probability law of the random variable $\chi \in G$ defined as

$$\chi = \bar{\chi} \exp(\xi), \quad \xi \sim \mathcal{N}(0, P), \quad 8.$$

where $\mathcal{N}(\cdot, \cdot)$ is the classical Gaussian distribution in Euclidean space \mathbb{R}^d identified to the Lie algebra \mathfrak{g} , and $P \in \mathbb{R}^{d \times d}$ is a covariance matrix. In Equation 8, the original Gaussian vector ξ of the Lie algebra is moved over by left multiplication to be centered at $\bar{\chi} \in G$, and we similarly define the distribution $\chi \sim \mathcal{N}_R(\bar{\chi}, P)$ for right multiplication of $\bar{\chi}$ through the random variable

$$\chi = \exp(\xi) \bar{\chi}, \quad \xi \sim \mathcal{N}(0, P). \quad 9.$$

In Equations 8 and 9, $\bar{\chi}$ is deterministic and can take any value, whereas P is the covariance of the small, noisy perturbation ξ . We stress that we have defined these probability density functions directly in the vector space \mathbb{R}^d and that both $\mathcal{N}_L(\cdot, \cdot)$ and $\mathcal{N}_R(\cdot, \cdot)$ are not Gaussian distributions. These kinds of distributions, introduced to the best of our knowledge in References 17, 20, and 21, were advocated and studied in Reference 18 for pose estimation and leveraged for Kalman filtering in References 6, 7, and 25. They are sometimes referred to as concentrated Gaussians on Lie groups. Note that χ in Equations 8 and 9 is a well-defined random variable of G . Computing its distribution on G is not trivial, but for the purposes we pursue in invariant filtering, it is not necessary.

Several approaches to filtering for systems possessing a geometric structure have been developed in the literature. For stochastic processes on Riemannian manifolds (26), some results have been derived (see, e.g., 27). The specific situation where the process evolves in a vector space but the observations belong to a manifold has also been considered (see, e.g., 10, 28, and more recently 29). For systems on Lie groups, powerful tools to study the filtering equations—such as harmonic analysis (30–32)—have been used, notably in the case of bilinear systems (9) and estimation of the initial condition of a Brownian motion (33). A somewhat different but related approach to filtering consists of finding the path that best fits the data in a deterministic setting. This approach is thus related to optimal control theory, where geometric methods have long played an important role (8). A certain class of least squares problems on the Euclidean group has been tackled in Reference 34 (see also 35).

3. INVARIANT KALMAN FILTERING

In invariant Kalman filtering, the state space is assumed to be a matrix Lie group. As a result, it comes with the nonlinear machinery of Section 2.1. However, we emphasize that it is the coupling of the choice of the Lie group structure and the dynamical model used in Equation 1 that allows the derivation of properties of the filter. Originally, only dynamics that were invariant to the group action were considered—hence, the name IEKF is rooted in References 36 and 37—and the class of dynamics to be considered has recently been much extended, as presented in Section 3.1. The remainder of this section is as follows: Section 3.2 is devoted to the methodology of the IEKF, Section 3.3 to the geometric vision that underlies it, and Section 3.4 to its recent unscented version.

3.1. Group Affine Systems

This section reviews some results from Reference 7, which are derived in the continuous-time context, and we present here their discrete-time counterpart. Consider a deterministic dynamical system on a matrix Lie group G :

$$\chi_n = f(\chi_{n-1}, u_n), \quad 10.$$

with f a smooth map on the group and $u_n \in \mathbb{R}^m$, or possibly $u_n \in G$ (in which case we use a capital U_n). In invariant Kalman filtering, function f is required to verify

$$\forall \chi, \nu \in G, \forall u, \quad f(\chi \nu, u) = f(\chi, u) f(I_N, u)^{-1} f(\nu, u). \quad 11.$$

Systems that satisfy this condition are called group affine systems in Reference 7. Indeed, if the group is merely a vector space with addition as the group composition law, we recover the affine functions, as Equation 11 then becomes $f(a + b, u) = f(a, u) - f(0, u) + f(b, u)$, which implies $f(x, u)$ is necessarily affine if it is smooth. Note that the continuous version of Equation 11 introduced in Reference 7 is similar to the notion of linearity on a group defined in a different context in Reference 38.

Theorem 1 (fundamental property of invariant filtering). We have the following equivalence:

$$f \text{ satisfies Equation 11} \Leftrightarrow \text{There exists a map } g \text{ such that } \forall \chi, \nu, u, \\ f(\nu, u)^{-1} f(\chi, u) = g(\nu^{-1} \chi, u).$$

Moreover, in this case, for each $u \in \mathbb{R}^m$ there exists $F \in \mathbb{R}^{d \times d}$ such that $\forall \xi \in \mathbb{R}^d$, $g(\exp(\xi), u) = \exp(F\xi)$; that is, the function $g(\cdot, u)$ is wholly encoded in a simple matrix F .

This result is pivotal, as will be explained in Section 4. The fact that $g(\cdot, u)$ may be described by a mere matrix F is not a first-order approximation; it is true at all orders, which is a quite surprising result.

Example 2. Let us continue our simple example of rotation matrices. Consider a motion R_0, R_1, \dots on $\text{SO}(3)$ with $R_n = R_{n-1} U_n$. The matrix $U_n \in \text{SO}(3)$ represents the relative rotation undergone by a body in space between time steps $n - 1$ and n and can be measured by (flawless) gyroscopes; it can thus be viewed as an input. The dynamics can be cast into the form of Equation 10, letting $f(R, U) = RU$. Thus, $f(\chi \nu, U) = \chi \nu U = f(\chi, U) f(I_3, U)^{-1} f(\nu, U)$, so Equation 11 is satisfied, and we have

$$f(\nu, U)^{-1} f(\chi, u) = (\nu U)^{-1} \chi U = U^{-1} \nu^{-1} \chi U = U^{-1} (\nu^{-1} \chi) U.$$

We see that the last term is indeed a function of $\nu^{-1} \chi$ and U , as predicted by the theorem (take indeed $g(\chi, U) = U^{-1} \chi U$). This has been long known in the $\text{SO}(3)$ case and was heavily exploited for attitude observer design (see, e.g., 6, 11–16, 37). Moreover, it is well known from Lie group theory (see, e.g., 18) that for any $\xi \in \mathbb{R}^d$, we have $U^{-1} \exp(\xi) U = \exp(\text{Ad}_U \xi)$, where $\text{Ad}_U \xi = U \xi$ in the $\text{SO}(3)$ case. Letting $F = U$, we see that this agrees again with the second part of Theorem 1.

3.2. Invariant Extended Kalman Filter Methodology

This section summarizes the IEKF methodology described in Reference 7. Consider a general dynamical system $\chi_n \in G$ associated with a sequence of observations $(Y_n)_{n \geq 0} \in \mathbb{R}^p$ as follows:

$$\begin{aligned} \chi_n &= f(\chi_{n-1}, u_n) \exp(w_n), & 12. \\ Y_n &= b(\chi_n) + V_n, & 13. \end{aligned}$$

where $u_n \in \mathbb{R}^d$ is a control input, $w_n \in \mathbb{R}^d$ is an unknown vector encoding the process noise, $V_n \in \mathbb{R}^p$ is the (unknown) measurement noise, and f satisfies Equation 11. The (left) IEKF² computes in real time an approximation to the posterior using the uncertainty representation Equation 8; that is, it computes at each step two parameters, $\hat{\chi}_{n|n} \in G$ and $P_{n|n} \in \mathbb{R}^{d \times d}$, and makes the approximation $\mathbb{P}(\chi_n | u_{1:n}, Y_{1:n}) \approx \mathcal{N}_L(\hat{\chi}_{n|n}, P_{n|n})$. The two-step procedure is then as follows:

1. *Propagation*: The IEKF propagates an estimate obtained after the previous observation Y_{n-1} through the deterministic part of Equation 12, i.e., by setting $w_n = 0$:

$$\hat{\chi}_{n|n-1} = f(\hat{\chi}_{n-1|n-1}, u_n). \quad 14.$$

To compute the associated covariance, introduce the left-invariant estimation errors defined as

$$\eta_{n-1|n-1} := \chi_{n-1}^{-1} \hat{\chi}_{n-1|n-1}, \quad \eta_{n|n-1} := \chi_n^{-1} \hat{\chi}_{n|n-1}. \quad 15.$$

This error is indeed invariant to left multiplications $(\chi, \hat{\chi}) \mapsto (\Gamma\chi, \Gamma\hat{\chi})$ for any group element $\Gamma \in G$, and right-invariant errors can be analogously defined. This is a natural way of measuring the discrepancy between the true state χ and the estimate $\hat{\chi}$ in a Lie group context where the usual linear error $\hat{\chi} - \chi$ is not even an element of G . Because f satisfies Equation 11, Theorem 1 yields

$$\eta_{n|n-1} = [f(\chi_{n-1}, u_n) \exp(w_n)]^{-1} f(\hat{\chi}_{n-1|n-1}, u_n) = \exp(-w_n) g(\eta_{n-1|n-1}, u_n), \quad 16.$$

where we have also used the fact that $\exp(w_n)^{-1} = \exp(-w_n)$. Along the lines of the EKF, we want to linearize the error system through a first-order Taylor expansion of the nonlinear functions f and b at the estimate $\hat{\chi}$. However, in contrast to Equation 4, the estimation errors defined in Equation 15 are elements of G , that is, square matrices rather than vectors. But when χ and $\hat{\chi}$ are close, the invariant error $\eta = \chi^{-1} \hat{\chi}$ is close to the identity matrix I_N . Using the fact that the Lie exponential map provides a bijection between a neighborhood of \mathbb{R}^d and a neighborhood of I_N , the estimation error can be locally approximated by an element of \mathbb{R}^d ; that is, we let $\xi_{n-1|n-1}, \xi_{n|n-1} \in \mathbb{R}^d$ be defined by

$$\eta_{n-1|n-1} = \exp(\xi_{n-1|n-1}), \quad \eta_{n|n-1} = \exp(\xi_{n|n-1}). \quad 17.$$

Let $F_n \in \mathbb{R}^{d \times d}$ be defined by $\forall \xi \in \mathbb{R}^d$, $g(\exp(\xi), u_n) = \exp(F_n \xi)$. This matrix exists thanks to Theorem 1, and it is not defined through first-order approximation, as in the conventional EKF case. According to the BCH formula, we have $\exp(-w_n) g(\exp(\xi_{n-1|n-1}), u_n) = \exp(F_n \xi_{n-1|n-1} - w_n)$ up to terms of order $\|\xi_{n-1|n-1}\|^2$, $\|w_n\|^2$, $\|\xi_{n-1|n-1}\| \|w_n\|$. Recalling Equations 16 and 17, this means that, up to the first order, $\exp(\xi_{n|n-1}) = \exp(F_n \xi_{n-1|n-1} - w_n)$, so we get the following linearized equation, which is the Lie group counterpart of Equation 5:

$$\xi_{n|n-1} = F_n \xi_{n-1|n-1} - w_n. \quad 18.$$

We have $\exp(\xi_{n|n-1}) = \eta_{n|n-1} = \chi_n^{-1} \hat{\chi}_{n|n-1}$, which implies that $\chi_n = \hat{\chi}_{n|n-1} \exp(-\xi_{n|n-1})$. Resorting to the uncertainty representation of Equation 8 and using Equation 18, we have just proved that if $\mathbb{P}(\chi_{n-1} | u_{1:n-1}, Y_{1:n-1}) \approx \mathcal{N}_L(\hat{\chi}_{n-1|n-1}, P_{n-1|n-1})$, then we have approximately the propagated distribution $\mathbb{P}(\chi_n | u_{1:n}, Y_{1:n-1}) \approx \mathcal{N}_L(\hat{\chi}_{n|n-1}, P_{n|n-1})$, where $P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + Q_n$ and $Q_n = \text{Cov}(w_n) = \text{Cov}(-w_n)$.

²For simplicity of exposure, the IEKF theory is presented using the left-invariant error $\eta = \chi^{-1} \hat{\chi}$ and the uncertainty representation shown in Equation 8. Swapping the role of left and right multiplications, we can similarly define an alternative filter, the right-invariant EKF.

Remark 1. Even though the state belongs to a nonlinear space and is a matrix and not a vector, the Lie exponential map allows us to linearize the error system in \mathbb{R}^d as in the conventional EKF methodology. Moreover, note that F_n depends on u_n but not on $\hat{\chi}_{n-1|n-1}$, which is in contrast to the conventional EKF, where F_n in Equation 5 depends on $\hat{X}_{n-1|n-1}$. This is a consequence of Equation 11 and Theorem 1 and will play an important role in Section 4.

2. *Measurement update:* To account for the new measurement, we let $z_n = Y_n - b(\hat{\chi}_{n|n-1})$ be the innovation. It is a vector of \mathbb{R}^p . We have $z_n = b(\chi_n) - b(\hat{\chi}_{n|n-1}) + V_n = b(\hat{\chi}_{n|n-1} \exp(\xi_{n|n-1})) - b(\hat{\chi}_{n|n-1}) + V_n$. As $\xi_{n|n-1}$ is assumed to be small, and as $\exp(0) = I_N$, a first-order Taylor expansion in $\xi \in \mathbb{R}^d$ arbitrary allows defining H_n as

$$b(\hat{\chi}_{n|n-1} \exp(\xi)) - b(\hat{\chi}_{n|n-1}) := H_n \xi + O(\|\xi\|^2). \quad 19.$$

Now that we have obtained a linearized system in \mathbb{R}^d akin to Equations 5 and 6, the conventional Kalman theory can be applied to derive the Kalman gain K_n and the updated covariance matrix $P_{n|n}$. The term $K_n z_n$ is a corrective shift computed on the linearized system. Thus, it should act on the linearized error $\xi_{n|n-1}$. As our estimation errors on the group are of the form $\exp(\xi) = \chi^{-1} \hat{\chi}$ —that is, $\chi = \hat{\chi} \exp(\xi)$ —an approximation to the best estimate of χ_n after observation Y_n that is consistent with Equation 15 is obtained through the Lie group counterpart $\hat{\chi}_{n|n} = \hat{\chi}_{n|n-1} \exp(K_n z_n)$ of the linear update shown in Equation 7. The equations of the filter are detailed in Algorithm 2.

Algorithm 2 (invariant extended Kalman filter). Choose initial $\hat{\chi}_{0|0} \in G$ and $P_{0|0} \in \mathbb{R}^{d \times d} = \text{Cov}(\xi_{0|0})$.

loop

Define H_n as in Equation 19 and F_n through the expansion $g(\exp(\xi), u_n) = \exp(F_n \xi)$

Define Q_n as $\text{Cov}(w_n)$ and R_n as $\text{Cov}(V_n)$

Propagation

$$\hat{\chi}_{n|n-1} = f(\hat{\chi}_{n-1|n-1}, u_n)$$

$$P_{n|n-1} = F_n P_{n-1|n-1} F_n^{-1} + Q_n$$

Measurement update

Compute $z_n = Y_n - b(\hat{\chi}_{n|n-1})$, $S_n = H_n P_{n|n-1} H_n^T + R_n$, and

$$K_n = P_{n|n-1} H_n^T S_n^{-1}$$

$$P_{n|n} = [I - K_n H_n] P_{n|n-1}$$

$$\hat{\chi}_{n|n} = \hat{\chi}_{n|n-1} \exp(K_n z_n)$$

end loop

Remark 2. Many systems of interest are such that Equation 13 is of the form $Y_n = \chi_n(b + V_n)$, where $b \in \mathbb{R}^N$ is a known vector. Indeed, this means that one measures some combination of the matrix entries, i.e., partial measurements of the state. In this case, we use an alternative innovation $z_n = \hat{\chi}_{n|n-1}^{-1} Y_n - b$, which is equal to $\eta_{n|n-1}^{-1} b - b + V_n$. Then, since $\eta_{n|n-1}^{-1} b = \exp(-\xi_{n|n-1}) b$, the linearized output H_n corresponding to innovation z_n is defined similarly to Equation 19, through the first-order expansion $z_n = \exp(-\xi)^{-1} b - b := H_n \xi + O(\|\xi\|^2)$. We then see that, along the lines of Remark 1, the matrix H_n does not depend on $\hat{\chi}_{n|n-1}$ either. Note that the choice of the alternative innovation above is mostly tutorial; a rigorous application of Algorithm 2 to the present case would in fact lead to entirely identical estimates.

Table 1 Differences between a conventional extended Kalman filter (EKF) and an invariant extended Kalman filter (IEKF)

	EKF		IEKF	
	Expression	Nature	Expression	Nature
State	X_n	Vector of \mathbb{R}^d	$\hat{\chi}_n$	Matrix of $G \subset \mathbb{R}^{N \times N}$
Uncertainty representation	$\mathcal{N}(\hat{X}, P)$	Gaussian in \mathbb{R}^d	$\mathcal{N}_L(\hat{\chi}, P)$	Random matrix (see Equation 8)
Linearized error	$e_{n n}$	Vector of \mathbb{R}^d	$\xi_{n n}$	Vector of \mathbb{R}^d
Nonlinear error	$e_{n n}$	Vector of \mathbb{R}^d	$\eta_{n n} = \exp(\xi_{n n})$	Matrix of $G \subset \mathbb{R}^{N \times N}$
Covariance matrix	P	Matrix of $\mathbb{R}^{d \times d}$	P	Matrix of $\mathbb{R}^{d \times d}$
Correction term	$K_n z_n$	Shift in \mathbb{R}^d	$K_n z_n$	Shift in \mathbb{R}^d

3.3. Geometric Insight

Table 1 compares the main features of both EKF and IEKF methodologies. The IEKF features admit a geometric interpretation illustrated by **Figure 2**.

3.4. Unscented Version

When applying Algorithm 2 to a particular system, matrices F_n, H_n must be computed, which might prove a little difficult to some practitioners. This is one of the main reasons why the UKF of Reference 39 has become a popular alternative to the EKF. Brossard et al. (25) proposed an unscented version of the IEKF that spares practitioners a computation of the Jacobians for systems modeled using discrete-time dynamics.

There have been other attempts to use (partially) the Lie group structure for UKF design. Condomines et al. (40, 41) presented a symmetry-preserving observer design based on the unscented transform, Loianno et al. (22) used the Lie group structure of SE(3) for a drone navigation application, and Forbes & Zotnik (42) exploited related ideas.

4. THEORETICAL CONVERGENCE GUARANTEES

Over the past decade, a large body of literature has been devoted to deterministic nonlinear observer design on Lie groups in the control community. We partially review this field in Section 4.2.




Figure 2

In the invariant extended Kalman filter methodology, the covariance matrix P represents a dispersion in \mathbb{R}^d (illustrated by the red 99% confidence ellipsoid in the *left plot*). In turn, using the bijection $\xi \mapsto \xi^\wedge$, P corresponds to a dispersion in the Lie algebra $T_{I_N}G = \mathfrak{g}$ (*central plot*). By adopting the uncertainty representation of Equation 8, we implicitly posit that the left-invariant error $\chi^{-1}\hat{\chi}$ is the exponential of a Gaussian. This means that we implicitly use left multiplication to move the ellipsoid over from $T_{I_N}G$ to $T_{\hat{\chi}_n}G$ (*right plot*). Thus, the choice of a particular nonlinear estimation error rules the way that confidence ellipsoids “turn” when moved over in G . Using Equation 9 instead—that is, right-invariant errors (and thus right multiplication $\xi \mapsto \xi^\wedge \hat{\chi}_{n|n}$) or, more generally, the original coordinates in which the system is modeled—we would get a quite different ellipsoid in $T_{\hat{\chi}_{n|n}}G$.




Figure 3

Architecture of the extended Kalman filter (Algorithm 1). The gain depends on the estimated state.

But first, we review the convergence properties of the IEKF when it is used as a candidate nonlinear observer.

4.1. The Invariant Extended Kalman Filter as a Stable Nonlinear Observer

Consider the model of Equations 1 and 2 and turn the noise off. This yields a deterministic system of the form $X_n = f(X_{n-1}, u_n)$, $Y_n = b(X_n)$. A stable nonlinear observer is a dynamical system of the form $\hat{X}_n = \hat{f}(\hat{X}_{n-1}, u_n, Y_n)$, where \hat{f} is designed to achieve asymptotic convergence of the estimation error; that is, $\hat{X}_n - X_n \rightarrow 0$ when $n \rightarrow \infty$. Although no noise perturbs the system, the state X_n is assumed to be unknown, and the system's equations yield only partial information about it. As a result, designing stable nonlinear observers is often a great challenge.

4.1.1. The conventional extended Kalman filter as a nonlinear observer. When facing a nonlinear observer problem, any EKF can readily be used as a candidate asymptotic observer by choosing arbitrary matrices Q_n, R_n , which are then viewed as tuning parameters. Unfortunately, there are not many guarantees that the EKF will asymptotically converge when used as an observer. The main results (43, 44) rely on strong assumptions about the filter's behavior, and the EKF can actually fail to converge sometimes, even for small initial estimation errors. To understand why, consider the block diagram of **Figure 3**, which illustrates the architecture of the EKF (Algorithm 1).

The important point is that the computation of the gain K_n relies on the matrices F_n, H_n introduced at Equations 5 and 6, since the system is linearized at the estimate. This creates a loop between the estimate and the gain computation that can destabilize the filter. Indeed, if the estimate is not sufficiently close to the true state, the gain will be erroneous, and the correction applied by the filter Equation 7 will amplify the estimation error. This positive feedback may lead to divergence.

4.1.2. The invariant extended Kalman filter as a nonlinear observer. The IEKF does not suffer from this drawback. Consider Equation 12 with the noise turned off—that is, $\chi_n = f(\chi_{n-1}, u_n)$ —and assume that f satisfies Equation 11. As in Remark 2, also assume that the observation (with the noise turned off) writes $Y_n = b(\chi_n) = \chi_n b$, with $b \in \mathbb{R}^N$ a known vector. As emphasized by Remarks 1 and 2, neither F_n nor H_n then depends on the estimates. This makes the gain computation independent from the estimate and spares the filter potentially harmful positive




Figure 4

Architecture of the invariant extended Kalman filter (Algorithm 2). The gain does not depend on the estimates.

feedback. Indeed, as illustrated by the block diagram of **Figure 4**, the inner loop of the EKF has disappeared.³

The independence of the gain computation to the trajectory means that if the system were linearized at the (unknown) true state, the gain K_n would be identical to the one obtained by linearizing at the estimate. Thanks to this strong property, the stability of the IEKF as an observer could be proved in Reference 7. Although the guaranteed convergence properties are only local, this is in sharp contrast to the conventional EKF, where no such guarantees exist.

4.2. Nonlinear Observers on Lie Groups

There has been a huge body of research devoted to nonlinear observers on Lie groups over the past decade, especially for attitude estimation. Although the literature is too broad to be covered here, we can point to a few significant references. Early developments in estimation on Lie groups were covered in Section 2.2. In the early 2000s, there were essentially two streams of research that bolstered the development of observers on Lie groups. The first was initiated by Aghannan & Rouchon (45) and seeks to design nonlinear observers that share the symmetries of the original system. The theory was formalized and developed in Reference 36 and applied to estimation on Lie groups in Reference 37. At the same time, the complementary filter on $SO(3)$ for attitude estimation was introduced in Reference 11. This filter makes extensive use of left-invariant errors on $SO(3)$ and the autonomy properties of the error equation. Owing to its simplicity and global convergence guarantees, it has become a renowned attitude estimator and has proved useful for quadrotor unmanned aerial vehicle (UAV) control. Papers that study observers that are akin to the complementary filter on $SO(3)$ for attitude estimation include References 6 and 12–16, and similar ideas have been applied to pose estimation (e.g., in 46). The idea to use similar techniques for noisy (instead of deterministic) systems on Lie groups and the subsequent theory of invariant Kalman filtering is slightly more recent (see, e.g., 2, 4–7, 40).

5. INDUSTRIAL APPLICATION: DRONE NAVIGATION

Hybrid inertial navigation was among the first applications of the EKF in the 1960s and a driver for scientific breakthroughs and reformulations that led from the original work of Kalman to a

³In the general IEKF theory of Reference 7, there are in fact situations where the gain may partially depend on the trajectory, but in a way that proves to be harmless.



Figure 5

The Euroflir 410 (*left*), the last generation of ultra-long-range electro-optical system commercialized by Safran. Its navigation system includes the first commercial implementation of an invariant extended Kalman filter. It is embedded in particular in the Patroller drone (*right*).

widespread industrial tool (see, e.g., 47). The company Safran Electronics & Defense (formerly known as Sagem), which is the largest European manufacturer of inertial navigation systems (INSs), chose to invest in invariant filtering. The Euroflir 410 (**Figure 5**) is the first commercial product implementing invariant Kalman filtering. It is notably embedded in the recently developed Patroller drone, which is a long-endurance UAV used for military purposes (the French army recently purchased 14 Patroller drones) and various other tasks, such as detection of forest fires.

Figure 6 was created from real in-flight experimental data for Global Positioning System/INS (GPS/INS) hybrid navigation. The initial heading error was deliberately set to an extreme value (90°) to obtain experimental confirmation that the convergence is not affected regardless of the




Figure 6

Heading estimated by an invariant extended Kalman filter–based inertial navigation system started in flight with an extremely large initial error (90°) (panel *a*). Owing to the properties highlighted in Section 4, convergence is not perturbed despite a large initial error and complicated trajectory (panel *b*).




Figure 7

Heading error for an invariant extended Kalman filter (IEKF) (*left*) and a conventional extended Kalman filter (EKF) (*right*) for Global Positioning System/inertial navigation system (GPS/INS) hybrid navigation. The filters have identical performance for very small initial estimation errors (5°). However, much larger initial estimation errors (90°) completely degrade the time required by the EKF to converge, whereas they only marginally degrade the IEKF’s performance. The small-angles hypothesis thus seems to be crucial to the EKF but not to the IEKF.

actual trajectory followed (a fact that agrees with the insights of Section 4). The situation is entirely different when using a conventional EKF, as shown in **Figure 7**, which was also created from real experimental data. The conventional EKF and IEKF indeed share similar behavior for small angles (around 5°), but the EKF rapidly deteriorates for larger angles, while the IEKF keeps converging almost as fast.

Remark 3. Commercial navigation systems such as the one equipping the Euroflir 410 have to model a very high-dimensional state that includes sensor bias, scale factors, and several other uncertainty sources, so that the state space is not a Lie group and the dynamics are not strictly group affine. However, the IEKF methodology can easily be adapted and used by appending the other variables to the state and treating them along the lines of the EKF methodology. The success of this approach stems from the fact that navigation systems are clearly much closer to group affine systems than to linear systems, as assumed by the conventional EKF methodology.

6. APPLICATION TO SIMULTANEOUS LOCALIZATION AND MAPPING

The SLAM problem has a rich history in robotics and autonomous navigation over the past two decades (see, e.g., 48, 49). It can be posed as a filtering problem, and EKF-based SLAM (EKF-SLAM) was one of the first algorithms used in this field. It was mostly abandoned, though, owing to the inconsistencies of its estimates. However, it was recently proved in Reference 23 that invariant Kalman filtering resolves the inconsistency issues of the EKF.

The IEKF-based SLAM (IEKF-SLAM) algorithm was introduced and studied in 2015 (23), and preliminary ideas date back to 2012 (50). Zhang et al. (51) derived some complementary properties of the IEKF-SLAM algorithm and successfully applied them to visual-inertial SLAM (52).




Figure 8

(a) In the simultaneous localization and mapping (SLAM) problem, a robot (*triangle*) is moving and makes measurements relative to some landmarks (*circles*). The goal is to simultaneously estimate the robot's trajectory and the position of the landmarks, i.e., the map. (b) The problem is unobservable because a global rotation of the robot and the landmarks is impossible to observe; the robot makes only relative measurements of landmarks and thus cannot distinguish between the left and right configurations shown in this panel.

6.1. The Inconsistency of Extended Kalman Filter–Based Simultaneous Localization and Mapping and Benefits of the Invariant Extended Kalman Filter

The issue of EKF-SLAM inconsistency has been the subject of many papers (see, e.g., 53–55), which have accumulated empirical evidence and theoretical explanations in particular situations. In this context, inconsistency refers to the inability of the filter's output covariance matrix $P_{n|n}$ to correctly reflect the true error dispersion $\mathbb{E}(X_n - \hat{X}_{n|n})(X_n - \hat{X}_{n|n})^T$. In particular, it was proved that the orientation uncertainty is a key feature in the inconsistency and that this derives from the linearization process. A little later, Huang et al. (56, 57) provided a complementary sound theoretical analysis of EKF-SLAM inconsistency, insisting on the EKF's inability to correctly reflect the unobservabilities inherent to the SLAM problem.

Although all of the results below carry over to three-dimensional SLAM, we focus on the two-dimensional case for simplicity. In the SLAM problem in two dimensions, a robot moves and observes some fixed features (or landmarks) of the environment. At time n , let $x_n \in \mathbb{R}^2$ be the position of the robot, let θ_n be its orientation with respect to a global frame, and let $p_1, p_2, \dots, p_K \in \mathbb{R}^2$ be the positions of the K landmarks in the global frame (see **Figure 8a**). The state is the vector $X_n = (x_n, \theta_n, p_1, \dots, p_K)$, and the goal is to estimate it from observations of landmarks relative to the robot's frame through sensors (lidars or cameras) attached to the robot. The collection of landmarks p_1, p_2, \dots, p_K constitutes the map (hence the term SLAM). Let $R(\alpha)$ denote a planar rotation of angle α and consider the transformation on the state space $\Psi_\alpha : (x_n, \theta_n, p_1, \dots, p_K) \mapsto (R(\alpha)x_n, \theta_n + \alpha, R(\alpha)p_1, \dots, R(\alpha)p_K)$. As simply illustrated in **Figure 8b**, the state $X_n = (x_n, \theta_n, p_1, \dots, p_K)$ and the state $\Psi_\alpha(X_n)$ are indistinguishable; that is, when departing from either configuration, regardless of the motions of the robot, it is impossible for the robot to tell that its own position and orientation and the landmarks' positions are actually different in both cases.

The main source of the inconsistency of the EKF-SLAM algorithm derives from this unobservability, as illustrated by **Figure 9**. The linearity of the update step Equation 7 in the original variables combined with the uncertainty representation $\mathcal{N}(\hat{X}_{n|n}, P_{n|n-1})$ results in a mere translation of the confidence ellipsoid at the update step, which does not match with unobservabilities of the system. By contrast, the SLAM state space can be endowed with a Lie group structure, as first noted in Reference 50. As proved in Reference 23, this defines local coordinates such that




Figure 9

Difference between the conventional extended Kalman filter (EKF) and invariant extended Kalman filter (IEKF) update steps with respect to unobservable directions. (a) EKF-based simultaneous localization and mapping (SLAM). $\Psi_\alpha(\hat{X}_{n|n-1})$ and $\hat{X}_{n|n-1}$ correspond to the two undistinguishable configurations of **Figure 8b**, so $\{\Psi_\alpha(\hat{X}_{n|n-1}), \alpha \in \mathbb{R}\}$ defines a continuous curve of states undistinguishable from $\hat{X}_{n|n-1}$ in the state space. In terms of a linearized model, this means that $\Psi'_\alpha(X_n) := \frac{d}{d\alpha} \Psi_\alpha(X_n)$ defines a direction in the state space at X_n along which no information can ever be gained through measurement data. The problem with the EKF is as follows. The covariance $P_{n|n}$ is computed with local information (i.e., the linearized model) at the propagated step $\hat{X}_{n|n-1}$. Even if the updated belief correctly reflects the unobservability there—i.e., the confidence ellipsoid defined by $P_{n|n}$ is very elongated in the unobservable direction $\Psi'_\alpha(\hat{X}_{n|n-1})$ —it is then moved over to updated state $\hat{X}_{n|n}$ through a simple translation. But $\Psi'_\alpha(\hat{X}_{n|n-1}) \neq \Psi'_\alpha(\hat{X}_{n|n})$, which means that the updated belief $\mathcal{N}(X_{n|n}, P_{n|n})$ is elongated in the direction $\Psi'_\alpha(\hat{X}_{n|n-1})$, which does not match with the actual unobservable direction $\Psi'_\alpha(\hat{X}_{n|n})$ at $\hat{X}_{n|n}$. (b) IEKF-based SLAM. In the coordinates induced by the Lie group structure of the state space, the unobservable directions $\Psi'_\alpha(\hat{X})$ become independent of \hat{X} . The effect, back in the original variables, is that confidence ellipsoids “turn” in a way that matches with unobservable directions.

the unobservable directions become independent of the linearization point. As a result, an EKF-SLAM algorithm based on those coordinates—namely, an IEKF-SLAM algorithm—computes beliefs that always match with the unobservable directions, which is what allowed the consistency properties to be proven in Reference 23.

6.2. Experimental Results

The incremental smoothing and mapping (iSAM) algorithm presented in Reference 58 can be implemented using the package available at <https://svn.csail.mit.edu/isam>. The method finds its roots in work by Thrun & Montemerlo (59). At each time step, it returns the maximum likelihood value of the state given all past measurements. When it converges, iSAM can be viewed as optimal and provides a reference estimate. This is why, over the past few years, the SLAM community has gradually turned to optimization-based techniques.

The conventional EKF, described in Algorithm 1; the IEKF, which is a right-multiplication-based version of Algorithm 2; and the iSAM algorithm have been tested on the publicly available real Victoria Park data set described by Guivant et al. (60). This data set resulted from a large-scale experiment in which a vehicle drove through a park, where hundreds of trees constituted the landmarks. **Figure 10** shows the estimated trajectories from different SLAM algorithms along with the actual GPS-measured trajectory. The algorithms use the odometers, steering angle, and laser sensors’ measurements, and the GPS measurements are used solely for evaluation. **Table 2** shows the error with respect to the GPS measurements for various tunings of covariance matrices (because the noise characteristics are not precisely known, various choices are possible).

The table is instructive regarding the robustness of the filters: The results of the IEKF and iSAM are almost insensitive to retained tuning parameters [root-mean-square error (RMSE) stays between 5 and 11 m], while the conventional EKF shows erratic performance when parameter




Figure 10

(a) The utility car used for the experiment that generated the Victoria Park data set. The vehicle was equipped with a SICK laser range and bearing sensor, a linear variable differential transformer sensor for the steering, and a back-wheel odometer. Image courtesy of J. Guivant and E. Nebot, Australian Centre for Field Robotics. (b) The trajectory of the car estimated by several simultaneous localization and mapping (SLAM) algorithms along with Global Positioning System (GPS) measurements. Algorithm abbreviations: EKF, extended Kalman filter; IEKF, invariant extended Kalman filter; iSAM, incremental smoothing and mapping.

tuning is not optimized (RMSE increases up to 123 m). Moreover, the IEKF and iSAM results are always very close to each other. This is interesting, as there is already a large corpus of experience and method in the industry for EKF implementation in the field of navigation, which is not yet the case for optimization-based filtering techniques. In particular, the IEKF has already been industrially implemented on a flying device (as mentioned in Section 5), and invariant Kalman filtering may thus open up for novel, robust industrial SLAM algorithms.

Table 2 Root-mean-square position errors on the Victoria Park data set for several noise-tuning parameters

	σ		
	1%	4%	8%
σ_V	EKF		
1 m	6.54 m	123.01 m	34.5 m
10 m	6.32 m	10.04 m	13.2 m
σ_V	IEKF		
1 m	6.50 m	9.07 m	10.84 m
10 m	5.95 m	6.42 m	6.45 m
σ_V	iSAM		
1 m	6.56 m	9.61 m	10.50 m
10 m	5.52 m	5.99 m	6.29 m

Parameter σ encodes model noise, and parameter σ_V encodes observation noise. The trajectories returned by the filters were rotated and translated to best match the Global Positioning System (GPS) measurements, and the residual root-mean-square errors are reported. Note that the results achieved by the IEKF and iSAM algorithms are very close and are weakly sensitive to (large) variations in the tuning parameters, in contrast to the results of the conventional EKF. This is another argument in favor of the robustness of invariant Kalman filtering. Algorithm abbreviations: EKF, extended Kalman filter; IEKF, invariant extended Kalman filter; iSAM, incremental smoothing and mapping.

7. CONCLUSION

The use of differential geometry and, more precisely, continuous symmetries for estimation emerged in the nonlinear automatic control field decades ago. Attitude estimation, essentially for the control of UAVs, has been a strong driver for the development of Lie group–based estimators over the past 10 years, with the most popular attitude estimator probably being the nonlinear complementary filter described by Mahony et al. (11).

Invariant Kalman filtering was introduced a decade ago in Reference 2 and has motivated many developments. Since then, the IEKF has been mathematically shown to come with convergence properties, has been implemented industrially for drone navigation, and has been shown to resolve a major issue of the EKF when applied to SLAM. SLAM is an important application that is sometimes considered the holy grail of robotics, as it enables a robot to be truly autonomous (49). Invariant Kalman filtering opens the way to novel implementations, as the EKF is a proven algorithm in terms of implementation and validation, but previous attempts to apply the EKF to SLAM were not satisfactory.

This review has aimed to provide an accessible introduction to the methodology of invariant Kalman filtering and to allow readers to gain insight into the relevance of the method and the important differences with conventional extended Kalman filtering. This should be of interest to readers intrigued by the application of mathematical theories to practical applications and those interested in finding robust, simple-to-implement filters for localization, navigation, and SLAM, notably for autonomous vehicle guidance.

SUMMARY POINTS

1. Invariant Kalman filtering is a recent methodology to design extended Kalman filters (EKFs) based on alternative coordinates dictated by geometry.
2. The invariant extended Kalman filter (IEKF) comes with convergence, stability, and robustness properties that the conventional EKF lacks. It is, however, reserved for a class of systems on Lie groups or systems that are close to this class.
3. The IEKF is particularly well suited to the localization and navigation of autonomous vehicles and has been successfully implemented in an industrial product of Safran Electronics & Defense, which is the largest European manufacturer of inertial navigation systems.
4. The IEKF has recently been proved to resolve the inconsistencies of the EKF for the important application of simultaneous localization and mapping (SLAM) in robotics. Although the standard EKF is the traditional algorithm for SLAM, it had been largely abandoned because of its inconsistencies.

DISCLOSURE STATEMENT

A.B. works for Safran, which has registered three patents for applications of invariant filtering.

LITERATURE CITED

1. Stengel RF. 1986. *Optimal Control and Estimation*. New York: Courier
2. Bonnabel S. 2007. Left-invariant extended Kalman filter and attitude estimation. In *2007 46th IEEE Conference on Decision and Control*, pp. 1027–32. New York: IEEE

3. Bonnabel S, Martin P, Salaün E. 2009. Invariant extended Kalman filter: theory and application to a velocity-aided attitude estimation problem. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pp. 1297–304. New York: IEEE
4. Martin P, Salaün E. 2010. Generalized multiplicative extended Kalman filter for aided attitude and heading reference system. In *ALAA Guidance, Navigation, and Control Conference*, chap. 2010-8301. Reston, VA: Am. Inst. Aeronaut. Astronaut.
5. Barczyk M, Lynch AF. 2013. Invariant observer design for a helicopter UAV aided inertial navigation system. *IEEE Trans. Control Syst. Technol.* 21:791–806
6. Bourmaud G, Mégret R, Arnaudon M, Giremus A. 2015. Continuous-discrete extended Kalman filter on matrix Lie groups using concentrated Gaussian distributions. *J. Math. Imaging Vis.* 51:209–28
7. Barrau A, Bonnabel S. 2017. The invariant extended Kalman filter as a stable observer. *IEEE Trans. Autom. Control* 62:1797–812
8. Brockett RW. 1973. Lie algebras and Lie groups in control theory. In *Geometric Methods in System Theory*, ed. DQ Mayne, RW Brockett, pp. 43–82. Dordrecht, Neth.: D. Reidel
9. Willsky AS, Marcus SI. 1975. Estimation for bilinear stochastic systems. In *Variable Structure Systems with Application to Economics and Biology*, ed. RR Mohler, A Ruberti, pp. 116–37. Berlin: Springer
10. Duncan TE. 1977. Some filtering results in Riemann manifolds. *Inform. Control* 35:182–95
11. Mahony R, Hamel T, Pfimlin JM. 2008. Nonlinear complementary filters on the special orthogonal group. *IEEE Trans. Autom. Control* 53:1203–18
12. Tayebi A, McGilvray S, Roberts A, Moallem M. 2007. Attitude estimation and stabilization of a rigid body using low-cost sensors. In *2007 46th IEEE Conference on Decision and Control*, pp. 6424–29. New York: IEEE
13. Bohn J, Sanyal AK. 2012. Unscented state estimation for rigid body motion on SE(3). In *2012 IEEE 51st Annual Conference on Decision and Control (CDC)*, pp. 7498–503. New York: IEEE
14. Batista P, Silvestre C, Oliveira P. 2012. A GES attitude observer with single vector observations. *Automatica* 48:388–95
15. Khosravian A, Trumpf J, Mahony R, Lageman C. 2015. Observers for invariant systems on Lie groups with biased input measurements and homogeneous outputs. *Automatica* 55:19–26
16. Lee T. 2016. Global unscented attitude estimation via the matrix Fisher distributions on SO(3). In *2016 American Control Conference (ACC)*, pp. 4942–47. New York: IEEE
17. Chirikjian GS. 2009. *Stochastic Models, Information Theory, and Lie Groups*, Vol. 1: *Classical Results and Geometric Methods*. Boston: Birkhäuser
18. Barfoot TD, Furgale PT. 2014. Associating uncertainty with three-dimensional poses for use in estimation problems. *IEEE Trans. Robot.* 30:679–93
19. Chirikjian GS, Kobilarov M. 2014. Gaussian approximation of non-linear measurement models on Lie groups. In *2014 IEEE 53rd Annual Conference on Decision and Control (CDC)*, pp. 6401–6. New York: IEEE
20. Chirikjian GS. 2011. *Stochastic Models, Information Theory, and Lie Groups*, Vol. 2: *Analytic Methods and Modern Applications*. Boston: Birkhäuser
21. Wolfe K, Mashner M, Chirikjian GS. 2011. Bayesian Fusion on Lie groups. *J. Algebraic Stat.* 2:75–97
22. Loianno G, Watterson M, Kumar V. 2016. Visual inertial odometry for quadrotors on SE(3). In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1544–51. New York: IEEE
23. Barrau A, Bonnabel S. 2015. An EKF-SLAM algorithm with consistency properties. arXiv:1510.06263
24. Barp A, Briol FX, Kennedy AD, Girolami M. 2018. Geometry and dynamics for Markov chain Monte Carlo. *Annu. Rev. Stat.* 5:451–71
25. Brossard M, Bonnabel S, Condomines J. 2017. Unscented Kalman filtering on Lie groups. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. New York: IEEE. In press
26. Itô K. 1950. Stochastic differential equations in a differentiable manifold. *Nagoya Math. J.* 1:35–47
27. Ng S, Caines P. 1985. Nonlinear filtering in Riemannian manifolds. *IMA J. Math. Control Inform.* 2:25–36
28. Pontier M, Szpirglas J. 1987. Filtering on manifolds. In *Stochastic Modelling and Filtering*, ed. E Germani, pp. 147–60. Berlin: Springer
29. Said S, Manton JH. 2013. On filtering with observation in a manifold: reduction to a classical filtering problem. *SIAM J. Control Optim.* 51:767–83

30. Willsky A. 1973. *Dynamical systems defined on groups: structural properties and estimation*. PhD Thesis, Dep. Aeronaut. Astronaut., Mass. Inst. Technol., Cambridge
31. Park W, Liu Y, Zhou Y, Moses M, Chirikjian GS, et al. 2008. Kinematic state estimation and motion planning for stochastic nonholonomic systems using the exponential map. *Robotica* 26:419–34
32. Willsky AS. 1973. Some estimation problems on Lie groups. In *Geometric Methods in System Theory*, ed. DQ Mayne, RW Brockett, pp. 305–14. Dordrecht, Neth.: D. Reidel
33. Duncan T. 1988. An estimation problem in compact Lie groups. *Syst. Control Lett.* 10:257–63
34. Han Y, Park F. 2001. Least squares tracking on the Euclidean group. *IEEE Trans. Autom. Control* 46:1127–32
35. Zamani M, Trumppf J, Mahony R. 2013. Minimum-energy filtering for attitude estimation. *IEEE Trans. Autom. Control* 58:2917–21
36. Bonnabel S, Martin P, Rouchon P. 2008. Symmetry-preserving observers. *IEEE Trans. Autom. Control* 53:2514–26
37. Bonnabel S, Martin P, Rouchon P. 2009. Non-linear symmetry-preserving observers on Lie groups. *IEEE Trans. Autom. Control* 54:1709–13
38. Ayala V, Tirao J. 1999. Linear control systems on Lie groups and controllability. In *Differential Geometry and Control*, ed. G Ferreyra, R Gardner, H Hermes, H Sussmann, pp. 47–64. Proc. Symp. Pure Math. Vol. 64. Providence, RI: Am. Math. Soc.
39. Julier SJ, Uhlmann JK. 1997. New extension of the Kalman filter to nonlinear systems. In *Signal Processing, Sensor Fusion, and Target Recognition VI*, ed. I Kadar. Proc. SPIE Vol. 3068. Bellingham, WA: SPIE. <https://doi.org/10.1117/12.280797>
40. Condomines JP, Seren C, Hattenberger G. 2013. Nonlinear state estimation using an invariant unscented Kalman filter. In *AIAA Guidance Navigation and Control Conference*, chap. 2013-4869. Reston, VA: Am. Inst. Aeronaut. Astronaut.
41. Condomines JP, Seren C, Hattenberger G. 2014. Pi-invariant unscented Kalman filter for sensor fusion. In *2014 IEEE 53rd Annual Conference on Decision and Control (CDC)*, pp. 1035–40. New York: IEEE
42. Forbes J, Zotnik DE. 2017. Sigma point Kalman filtering on matrix Lie groups applied to the SLAM problem. In *Geometric Science of Information: Third International Conference, GSI 2017*, ed. F Nielsen, F Barbaresco, pp. 318–28. Cham, Switz.: Springer
43. Boutayeb M, Rafaralahy H, Darouach M. 1997. Convergence analysis of the extended Kalman filter used as an observer for nonlinear deterministic discrete-time systems. *IEEE Trans. Autom. Control* 42:581–86
44. Song Y, Grizzle JW. 1992. The extended Kalman filter as a local asymptotic observer for nonlinear discrete-time systems. In *1992 American Control Conference*, pp. 3365–69. New York: IEEE
45. Aghannan N, Rouchon P. 2002. On invariant asymptotic observers. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002*, Vol. 2, pp. 1479–84. New York: IEEE
46. Scandaroli GG, Morin P, Silveira G. 2011. A nonlinear observer approach for concurrent estimation of pose, IMU bias and camera-to-IMU rotation. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3335–41. New York: IEEE
47. McGee LA, Schmidt SF. 1985. *Discovery of the Kalman filter as a practical tool for aerospace and industry*. Tech. Memo. 86847, NASA, Washington, DC
48. Dissanayake G, Newman P, Durrant-Whyte H, Clark S, Csobra M. 2001. A solution to the simultaneous localisation and mapping (SLAM) problem. *IEEE Trans. Robot. Automat.* 17:229–41
49. Durrant-Whyte H, Bailey T. 2006. Simultaneous localization and mapping: part I. *IEEE Robot. Autom. Mag.* 13:99–110
50. Bonnabel S. 2012. Symmetries in observer design: review of some recent results and applications to EKF-based SLAM. In *Robot Motion and Control 2011*, ed. K Kozłowski, pp. 3–15. London: Springer
51. Zhang T, Wu K, Su D, Huang S, Dissanayake G. 2017. An invariant-EKF VINS algorithm for improving consistency. arXiv:1702.07920
52. Zhang T, Wu K, Song J, Huang S, Dissanayake G. 2017. Convergence and consistency analysis for a 3-D invariant-EKF SLAM. *IEEE Robot. Autom. Lett.* 2:733–40
53. Julier SJ, Uhlmann JK. 2001. A counter example to the theory of simultaneous localization and map building. In *Proceedings 2001 ICRA: IEEE International Conference on Robotics and Automation*, Vol. 4, pp. 4238–43. New York: IEEE

54. Castellanos JA, Neira J, Tardós JD. 2004. Limits to the consistency of EKF-based SLAM. *IFAC Proc. Vol. 37*:716–21
55. Huang S, Dissanayake G. 2007. Convergence and consistency analysis for extended Kalman filter based SLAM. *IEEE Trans. Robot.* 23:1036–49
56. Huang GP, Mourikis AI, Roumeliotis SI. 2008. Analysis and improvement of the consistency of extended Kalman filter based SLAM. In *2008 IEEE International Conference on Robotics and Automation*, pp. 473–79. New York: IEEE
57. Huang GP, Mourikis AI, Roumeliotis SI. 2010. Observability-based rules for designing consistent EKF SLAM estimators. *Int. J. Robot. Res.* 29:502–28
58. Kaess M, Ranganathan A, Dellaert F. 2008. iSAM: incremental smoothing and mapping. *IEEE Trans. Robot.* 24:1365–78
59. Thrun S, Montemerlo M. 2006. The graph SLAM algorithm with applications to large-scale mapping of urban structures. *Int. J. Robot. Res.* 25:403–29
60. Guivant J, Nebot E, Durrant-Whyte H. 2000. Simultaneous localization and map building using natural features in outdoor environments. *Intell. Auton. Syst.* 6:581–86