Systems & Control Letters 106 (2017) 87–95

Contents lists available at ScienceDirect

Systems & Control Letters

journal homepage: www.elsevier.com/locate/sysconle

Towards consistent filtering for discrete-time partially-observable nonlinear systems☆

Guoquan Huang

Department of Mechanical Engineering, University of Delaware, Newark, DE 19716, USA

ARTICLE INFO

Article history:
Received 25 October 2016
Received in revised form 10 June 2017
Accepted 21 June 2017

Keywords:
Nonlinear filtering
Extended Kalman filter
Consistency
Fisher information matrix
Observability matrix
Nonlinear systems

ABSTRACT

In this paper, we study the filter consistency of discrete-time nonlinear systems with partially-observable measurements, where the full state is not reconstructable from the available measurements at each time step. Linearized filters such as the extended Kalman filter (EKF) which are realized based on the corresponding linearized systems, may become inconsistent. Relying on a novel *decomposition* of the observability matrix based on different measurement sources, we show that the filter acquires spurious information from the measurements of each source, which erroneously reduces the uncertainty of the state estimates and hence causes inconsistency. Based on this key insight, we propose an *information-aware* methodology and develop two novel EKF algorithms of computing filter Jacobians which ensure that all decompositions of the observability matrix have nullspace of correct dimensions. In the first, the linearization points are selected so as to minimize their expected linearization errors under the constraints that the decompositions of the observability matrix have correct nullspace. In the second, we project the canonical measurement Jacobian onto the actual information-available directions. The proposed approaches are shown to significantly outperform the canonical EKF in the particular application of radar-based target tracking.

1. Introduction

An issue of concern with nonlinear estimation problems such as target tracking [1] is the filter inconsistency—that is, no provably consistent estimator can be constructed for a nonlinear system, and the consistency of every filter has to be evaluated experimentally. As defined in [2], a filter is consistent if its estimation errors are zero-mean and its covariance is equal to the true covariance. Consistency is one of the primary criteria for evaluating the performance of any filter, since if the filter is inconsistent, then its estimation accuracy is unknown, which in turn makes the filter unreliable. To date, the problem of estimation inconsistency has been studied primarily in robotics for simultaneous localization and mapping (SLAM) systems that are *unobservable* [3–6]. In particular, in our prior work [5], we have identified the observability mismatch between the EKF linearized system and the underlying nonlinear SLAM system as one main cause for the filter inconsistency; and based on this, have developed an observability-constrained (OC)-EKF to improve the particular EKF-SLAM consistency. However, little is known about the causes of filter inconsistency for general *observable* systems with partially-observable measurements.

To bridge this gap, in this paper, building upon our prior conference publication [7], we investigate the problem of filter inconsistency for a broad class of discrete-time nonlinear systems from a new *information* perspective, i.e., by examining the directions of the state space along which the information is actually available from the measurements of each source (sensor). Based on this analysis, we propose a novel *information-aware* methodology to improve consistency by ensuring that the filter acquires the information from each source's measurements along the correct directions of the state space.

Specifically, the Fisher information matrix (FIM) [2] of a set of measurements encapsulates all the available information about the entire state of a stochastic system. By marginalizing all but the initial state, we obtain the corresponding FIM that contains all the information available in the measurements for determining the initial state. Studying the FIM's structure reveals the directions along which the information is (un)available from the measurements. These can be exploited in the design of consistent filtering algorithms, i.e., by enforcing filters to gain information from measurements only along the "correct" directions (where information is actually available). Moreover, we *analytically* show that the FIM of the initial state can be factorized in terms of the observability

☆ This work was partially supported by the University of Delaware (UD) College of Engineering, UD Cybersecurity Initiative (UDCSI), the Delaware NASA/EPSCoR Seed Grant, the NSF (IIS-1566129), and the DTRA (HDTRA1-16-1-0039). Part of the paper was presented in the American Control Conference, Washington, DC, June 17–19, 2013.

E-mail address: ghuang@udel.edu.

http://dx.doi.org/10.1016/j.sysconle.2017.06.006
0167-6911/

matrix of the corresponding deterministic system. Based on this key finding, and to improve consistency, we propose to impose the constraint of acquiring information along the correct directions on the resulting decompositions (with respect to different measurement sources) of the observability matrix, instead of the FIM. To this end, we introduce two different EKFs that compute the appropriate filter Jacobians, either directly (i.e., by projecting the canonical Jacobians onto the information-available subspace) or indirectly (i.e., by first finding optimal linearization points). As a result, the filters only gain the information actually available from each source's measurements, and thus substantially improve consistency and accuracy, as opposed to the canonical EKF.

The remainder of the paper is structured as follows: After presenting our information-aware methodology in the next section, we demonstrate the proposed approach in the radar-tracking application in Section 3, and its performance is validated through Monte-Carlo simulations. Finally, Section 4 outlines the main conclusions of this work.

2. Information-aware methodology

Consider the following discrete-time nonlinear system:

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k \tag{1}$$

$$\mathbf{z}_{i,k} = \mathbf{h}(\mathbf{x}_k, \mathbf{s}_{i,k}) + \mathbf{v}_{i,k}\,, \quad i \in \{1, \ldots, s\} \tag{2}$$

where $\mathbf{x}_k \in \mathcal{R}^n$ denotes the state of the system, $\mathbf{u}_k \in \mathcal{R}^\eta$ is the control input, and $\mathbf{w}_k \in \mathcal{R}^n$ is zero-mean white Gaussian process noise, i.e., $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$. $\mathbf{z}_{i,k} \in \mathcal{R}^m$ is the measurement taken from the ith ($i \in \{1, \ldots, s\}$) measurement source (e.g., sensor), and is generally (although not necessarily) of *lower* dimension than the state vector (i.e., $m < n$), which is the partially-observable case as considered in this work. The parameter $\mathbf{s}_{i,k}$ denotes the known parameters of the ith measurement source, such as the sensor's location or a binary indicator of the availability of the ith measurement. The random variable $\mathbf{v}_{i,k} \in \mathcal{R}^m$ is zero-mean white Gaussian measurement noise, i.e., $\mathbf{v}_{i,k} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_{i,k})$.

We employ the EKF to recursively compute the state estimate and error covariance. Specifically, we linearize the nonlinear system at the linearization points, $\mathbf{x}^\star_{k|k-1}$ and $\mathbf{x}^\star_{k|k}$ (i.e., the linearization points before and after the update at time-step k) [see (1) and (2)], and obtain the linearized error-state system[1]:

$$\tilde{\mathbf{x}}_{k+1|k} = \boldsymbol{\Phi}_k \tilde{\mathbf{x}}_{k|k} + \mathbf{w}_k \tag{3}$$

$$\tilde{\mathbf{z}}_{i,k|k-1} = \mathbf{H}_{i,k} \tilde{\mathbf{x}}_{k|k-1} + \mathbf{v}_{i,k}\,, \quad i \in \{1, \ldots, s\} \tag{4}$$

where

$$\boldsymbol{\Phi}_k = \nabla_{\mathbf{x}_k} \mathbf{f}\Big|_{\{\mathbf{x}^\star_{k|k}, \mathbf{x}^\star_{k+1|k}\}}, \quad \mathbf{H}_{i,k} = \nabla_{\mathbf{x}_k} \mathbf{h}\Big|_{\{\mathbf{x}^\star_{k|k-1}\}}. \tag{5}$$

The canonical choice of linearization point is the latest state estimate, which, however, is not necessarily the best choice as we will show later. Once the propagation and measurement Jacobians are computed, we propagate and update the state estimate and covariance, respectively, as follows [2]:

$$\hat{\mathbf{x}}_{k+1|k} = \mathbf{f}(\hat{\mathbf{x}}_{k|k}, \mathbf{u}_k) \tag{6}$$

$$\mathbf{P}_{k+1|k} = \boldsymbol{\Phi}_k \mathbf{P}_{k|k} \boldsymbol{\Phi}_k^T + \mathbf{Q}_k \tag{7}$$

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \mathbf{r}_k \tag{8}$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k \mathbf{S}_k \mathbf{K}_k^T \tag{9}$$

where $\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}_{i,k}^T\mathbf{S}_k^{-1}$ is the Kalman gain, $\mathbf{r}_k = \mathbf{z}_{i,k} - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1}, \mathbf{s}_{i,k})$ is the residual, and $\mathbf{S}_k = \mathbf{H}_{i,k}\mathbf{P}_{k|k-1}\mathbf{H}_{i,k}^T + \mathbf{R}_{i,k}$ is the corresponding residual covariance.

2.1. Observability and Fisher information

Since the EKF is constructed based on the linearized system [see (3) and (4)], it is important to study the observability properties of the corresponding deterministic system (i.e., noise free). For a *deterministic* system, observability examines whether the information provided by the available measurements is sufficient for estimating the initial state without ambiguity. This however does not guarantee, while offering a hope for, viable estimation for the corresponding *stochastic* system. In particular, the observability matrix for the linearized system (3)–(4) during the time interval $[0, k]$ is defined by [8,9]:

$$\mathbf{M} = \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \boldsymbol{\Phi}_0 \\ \vdots \\ \mathbf{H}_k \boldsymbol{\Phi}_{k-1} \cdots \boldsymbol{\Phi}_0 \end{bmatrix}. \tag{10}$$

If the linear system is observable, then the corresponding observability matrix $\mathbf{M}$ is full-rank.

The FIM is closely related to the system observability and precisely describes the information available in the measurements [2]. Thus, by studying its properties, we can gain insight about the directions in the state space along which information is actually available. To this end, we examine the structure of the Hessian (information) matrix of the corresponding batch maximum a posteriori (MAP) estimator over the time interval $[0, k]$, which is known to be optimal [10]. In what follows, we show that the FIM of the initial state $\mathbf{x}_0$ (obtained by marginalizing or integrating over all other states) has the same rank properties as the observability matrix. This motivates us to instead examine the observability matrix, rather than the FIM, to enforce proper information acquisition from measurements, in the ensuing analysis.

The optimal batch-MAP estimator utilizes all available information to estimate the *entire* state trajectory that is formed by stacking all states in the time interval $[0, k]$:

$$\mathbf{x}_{0:k} = \begin{bmatrix} \mathbf{x}_0^T & \mathbf{x}_1^T & \cdots & \mathbf{x}_k^T \end{bmatrix}^T. \tag{11}$$

Specifically, the batch-MAP estimator seeks to determine the entire state-space trajectory estimate $\hat{\mathbf{x}}_{0:k|k}$ by maximizing the following posterior pdf (assuming *no* prior is available):

$$p(\mathbf{x}_{0:k}|\mathbf{z}_{0:k}) \propto \prod_{\kappa=0}^{k-1} p(\mathbf{x}_{\kappa+1}|\mathbf{x}_\kappa) \prod_{\kappa=0}^{k} p(\mathbf{z}_{i,\kappa}|\mathbf{x}_\kappa) \tag{12}$$

where $\mathbf{z}_{0:k}$ denotes all the sensor measurements in the time interval $[0, k]$. In the above expression, we have employed the assumption of independent state and measurement noise and the Markovian property of the system dynamics [see (1), and (2), respectively]. Moreover, using the assumption of Gaussian noise, the posterior pdf (12) can be written as:

$$p(\mathbf{x}_{0:k}|\mathbf{z}_{0:k}) \propto \prod_{\kappa=0}^{k-1} \frac{1}{\sqrt{|2\pi \mathbf{Q}_\kappa|}} \exp\left(-\frac{1}{2}\|\mathbf{x}_{\kappa+1} - \mathbf{f}(\mathbf{x}_\kappa, \mathbf{u}_\kappa)\|^2_{\mathbf{Q}_\kappa}\right) \times \prod_{\kappa=0}^{k} \frac{1}{\sqrt{|2\pi \mathbf{R}_{i,\kappa}|}} \exp\left(-\frac{1}{2}\|\mathbf{z}_{i,\kappa} - \mathbf{h}(\mathbf{x}_\kappa, \mathbf{s}_{i,\kappa})\|^2_{\mathbf{R}_{i,\kappa}}\right) \tag{13}$$

where we have employed the notation, $\|\mathbf{a}\|^2_{\Lambda} \triangleq \mathbf{a}^T \Lambda^{-1} \mathbf{a}$. Due to the monotonicity of the negative logarithm, maximizing (13) is

[1] Throughout this paper, the subscript $\ell|j$ refers to the estimate of a quantity at time step ℓ, after all measurements up to time step j have been processed. $\hat{x}$ is used to denote the estimate of a random variable x, while $\tilde{x} = x - \hat{x}$ is the error in this estimate. $\mathbf{0}_{m\times n}$ denotes $m \times n$ matrices of zeros, and $\mathbf{I}_n$ is the $n \times n$ identity matrix.