Robotics and Autonomous Systems 57 (2009) 1198–1210

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

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

Covariance recovery from a square root information matrix for data association

Michael Kaess [a,*], Frank Dellaert [b]

[a] CSAIL, Massachusetts Institute of Technology, 32 Vassar Street, Cambridge, MA 02139, United States

[b] College of Computing, Georgia Institute of Technology, 801 Atlantic Dr., Atlanta, GA 30332, United States

ARTICLE INFO

Article history:
Available online 3 July 2009

Keywords:
Data association
Smoothing
Simultaneous localization and mapping

ABSTRACT

Data association is one of the core problems of simultaneous localization and mapping (SLAM), and it requires knowledge about the uncertainties of the estimation problem in the form of marginal covariances. However, it is often difficult to access these quantities without calculating the full and dense covariance matrix, which is prohibitively expensive. We present a dynamic programming algorithm for efficient recovery of the marginal covariances needed for data association. As input we use a square root information matrix as maintained by our incremental smoothing and mapping (iSAM) algorithm. The contributions beyond our previous work are an improved algorithm for recovering the marginal covariances and a more thorough treatment of data association, now including the joint compatibility branch and bound (JCBB) algorithm. We further show how to make information theoretic decisions about measurements before actually taking the measurement, therefore allowing a reduction in estimation complexity by omitting uninformative measurements. We evaluate our work on simulated and real-world data.

1. Introduction

Data association is an essential component of simultaneous localization and mapping (SLAM) [1]. The data association problem in SLAM, which is also known as the correspondence problem, consists of matching the current measurements with their corresponding previous observations. Correspondences can be obtained directly between measurements taken at different times, or by matching the current measurements to landmarks already contained in the map based on earlier measurements. A solution to the correspondence problem provides frame-to-frame matching, but also allows for closing large loops in the trajectory. Such loops are more difficult to handle as the estimation uncertainty is much larger than between successive frames, and the measurements might even be taken from a different direction.

Performing data association can be difficult especially in ambiguous situations, but is greatly simplified when the state estimation uncertainties are known. Parts of the overall SLAM state estimate uncertainty are needed to make a probabilistic decision based on the maximum likelihood (ML) criterion or when using the joint compatibility branch and bound (JCBB) algorithm by Neira and Tardos [2], a popular algorithm for SLAM [3,1]. The parts that are required are so-called marginal covariances that represent the uncertainties between a relevant subset of the variables, for example a pose and a landmark pair.

However, it is generally difficult to recover the exact marginal covariances in real time. But as mobile robot applications require decisions to be made in real time, we need an efficient solution for recovering the marginal covariances. While it is trivial to recover the covariances from an Extended Kalman Filter (EKF), its uncertainties are inconsistent when non-linear measurement functions are present, which is typically the case. Other solutions to SLAM, for example based on iterative equation solvers such as [4–7], cannot directly access the marginal covariances. An alternative is to use conservative estimates of the marginal covariances as in [8]; however, they will provide fewer constraints for ambiguous data association decisions and therefore fail earlier.

Our solution provides efficient access to the marginal covariances based on the square root information matrix. Such a factored information matrix is maintained by our incremental smoothing and mapping (iSAM) algorithm [9], which efficiently updates the factored representation when new measurements arrive. Our solution consists of a dynamic programming algorithm that recovers only parts of the full covariance matrix based on the square root information matrix, thereby avoiding having to calculate the full dense covariance matrix, which contains a number of entries that is quadratic in the number of variables.

The contributions over our previous work [10,9] are an improved marginal covariance recovery algorithm and a detailed discussion of the algorithm. We also add the JCBB algorithm to our discussion of data association techniques, and use a uniform

* Corresponding author. Tel.: +1 678 365 1243.
E-mail addresses: kaess@mit.edu (M. Kaess), dellaert@cc.gatech.edu (F. Dellaert).

0921-8890/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.robot.2009.06.008

mathematical presentation to contrast the presented methods. Beyond typical data association work, we further show how to use these marginal covariances to determine the value of a specific measurement, allowing one to drop redundant or uninformative measurements in order to increase estimation efficiency. We present detailed evaluations on simulated and real-world data. And finally we provide insights into using JCBB versus the RANSAC algorithm by Fischler and Bolles [11].

2. Covariance recovery

We show how to efficiently obtain selected parts of the covariance matrix, the so-called marginal covariances, based on a square root information matrix. But first we briefly introduce the square root information matrix and an efficient algorithm for calculating it.

2.1. Square root information matrix

The square root information matrix appears in the context of smoothing and mapping (SAM) [12], a smoothing formulation of the SLAM problem. The smoothing formulation includes the complete robot trajectory, that is all poses $\mathbf{x}_i$ ($i \in \{0 \ldots M\}$) in addition to the landmarks $\mathbf{l}_j$ ($j \in \{1 \ldots N\}$). This is in contrast to typical filtering methods that only keep the most recent pose by marginalizing out previous poses. Smoothing provides the advantage of a sparse information matrix, therefore allowing one to efficiently solve [12] the equation system.

The SLAM problem typically contains non-linear functions (through robot orientation and bearing measurements) and therefore requires iterative linearization and solution steps. Please see [12,9] for a detailed treatment of the process and measurement models, and their linearization and combination into one large least-squares system. One step of the resulting linearized SLAM problem can be written as

$$\arg\min_{\mathbf{x}} \|A\mathbf{x} - \mathbf{b}\|^2 \tag{1}$$

where A is the measurement Jacobian of the SLAM problem at the current linearization point, $\mathbf{x}$ the unknown state vector combining poses and landmarks, and $\mathbf{b}$ the so-called right-hand side that is irrelevant in this work. Solutions to the state vector $\mathbf{x}$ in (1) can be found based on the *square root information matrix* R, an upper triangular matrix that is found by Cholesky factorization of the information matrix $\mathcal{I} := A^{\mathrm{T}}A = R^{\mathrm{T}}R$ or directly by QR factorization of the measurement Jacobian $A = Q\begin{bmatrix} R \\ 0 \end{bmatrix}$. The upper triangular shape of the square root information matrix allows efficient solution of the SLAM problem by back-substitution.

In practice it is too expensive to refactor the information matrix each time a new measurement arrives. Instead, our incremental smoothing and mapping (iSAM) algorithm [9] updates the square root information matrix directly with the new measurements. Periodic variable reordering keeps the square root information matrix sparse, allowing efficient solution by back-substitution as well as efficient access to marginal covariances, which is described next.

2.2. Recovering marginal covariances

Knowledge of the relative uncertainties between subsets $\{j_1, \ldots, j_K\}$ of the SLAM variables are needed for data association. In particular, the marginal covariances

$$\begin{bmatrix} \Sigma_{j_1j_1} & \Sigma_{j_2j_1}^{\mathrm{T}} & \cdots & \Sigma_{j_Kj_1}^{\mathrm{T}} \\ \Sigma_{j_2j_1} & \Sigma_{j_2j_2} & \cdots & \Sigma_{j_Kj_2}^{\mathrm{T}} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{j_Kj_1} & \Sigma_{j_Kj_2} & \cdots & \Sigma_{j_Kj_K} \end{bmatrix} \tag{2}$$

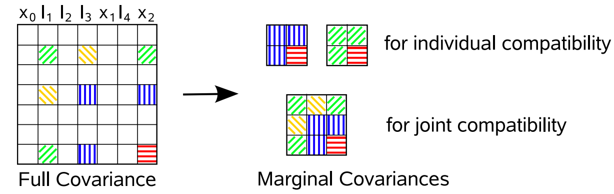

Fig. 1. Only a small number of entries of the dense covariance matrix are of interest for data association. In this example, both the individual and the combined marginals between the landmarks $\mathbf{l}_1$ and $\mathbf{l}_3$ and the latest pose $\mathbf{x}_2$ are retrieved. As we show here, these entries can be obtained without calculating the full dense covariance matrix.

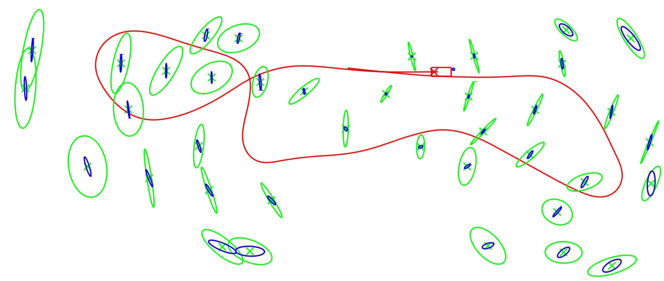

Fig. 2. Marginal covariances *projected into the current robot frame* (robot indicated by red rectangle) for a short trajectory (red curve) and some landmarks (green crosses). The exact covariances (blue, smaller ellipses) obtained by our fast algorithm coincide with the exact covariances based on full inversion (orange, mostly hidden by blue). Note the much larger conservative covariance estimates (green, large ellipses) as recovered in our previous work [9]. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

are the basis for advanced data association techniques that we discuss in detail in Section 3, as well as for information theoretic decisions about the value of measurements as discussed in Section 4. This marginal covariance matrix contains various blocks from the full covariance matrix, as is shown in Fig. 1. Calculating the full covariance matrix to recover these entries is not an option because the covariance matrix is always densely populated with n^2 entries, where n is the number of variables. However, we show in the next section that it is not necessary to calculate all entries in order to retrieve the exact values of the relevant blocks.

Recovering the exact values for all required entries without calculating the complete covariance matrix is not straightforward, but can be done efficiently by again exploiting the sparsity structure of the square root information matrix R. In general, the covariance matrix is obtained as the inverse of the information matrix

$$\Sigma := (A^{\mathrm{T}}A)^{-1} = (R^{\mathrm{T}}R)^{-1} \tag{3}$$

based on the factor matrix R by noting that

$$R^{\mathrm{T}}R\Sigma = I \tag{4}$$

and performing a forward substitution, followed by a back substitution

$$R^{\mathrm{T}}Y = I, \qquad R\Sigma = Y. \tag{5}$$

Because the information matrix is not band-diagonal in general, this would seem to require calculating all n^2 entries of the fully dense covariance matrix, which is infeasible for any non-trivial problem. This is where we exploit the sparsity of the square root information matrix R. Both Golub and Plemmons [13] and Triggs et al. [14] present an efficient method for recovering only the entries σ_{ij} of the covariance matrix Σ that coincide with non-zero