



Incremental sparse GP regression for continuous-time trajectory estimation and mapping

Xinyan Yan^{a,*}, Vadim Indelman^b, Byron Boots^a

^a Institute for Robotics and Intelligent Machines, College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, USA

^b Faculty of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel

HIGHLIGHTS

- An incremental sparse GP regression algorithm for STEAM problems is proposed.
- The benefits of GP-based approaches and incremental smoothing are combined.
- The approach elegantly handles asynchronous and sparse measurements.
- Results indicate significant speed-up in performance with little loss in accuracy.

ARTICLE INFO

Article history:

Received 26 January 2016

Accepted 4 October 2016

Available online 14 October 2016

Keywords:

State estimation

Localization

Continuous time

SLAM

Gaussian process regression

ABSTRACT

Recent work on simultaneous trajectory estimation and mapping (STEAM) for mobile robots has used Gaussian processes (GPs) to efficiently represent the robot's trajectory through its environment. GPs have several advantages over discrete-time trajectory representations: they can represent a continuous-time trajectory, elegantly handle asynchronous and sparse measurements, and allow the robot to query the trajectory to recover its estimated position at any time of interest. A major drawback of the GP approach to STEAM is that it is formulated as a *batch* trajectory estimation problem. In this paper we provide the critical extensions necessary to transform the existing GP-based batch algorithm for STEAM into an extremely efficient incremental algorithm. In particular, we are able to vastly speed up the solution time through efficient variable reordering and incremental sparse updates, which we believe will greatly increase the practicality of Gaussian process methods for robot mapping and localization. Finally, we demonstrate the approach and its advantages on both synthetic and real datasets.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction & related work

Simultaneously recovering the location of a robot and a map of its environment from sensor readings is a fundamental challenge in robotics [1–3]. Well-known approaches to this problem, such as square root smoothing and mapping (SAM) [4], have focused on regression-based methods that exploit the sparse structure of the problem to efficiently compute a solution. The main weakness of the original SAM algorithm was that it was a *batch* method: all of the data must be collected before a solution can be found. For a robot traversing an environment, the inability to update an estimate of its trajectory online is a significant drawback. In response to this weakness, Kaess et al. [5] developed a critical extension to the batch SAM algorithm, iSAM, that overcomes this problem by *incrementally* computing a solution. The main drawback of iSAM,

was that the approach required costly periodic batch steps for variable reordering to maintain sparsity and relinearization. This approach was extended in iSAM 2.0 [6], which employs an efficient data structure called the *Bayes tree* [7] to perform incremental variable reordering and just-in-time relinearization, thereby eliminating the bottleneck caused by batch variable reordering and relinearization. The iSAM 2.0 algorithm and its extensions are widely considered to be state-of-the-art in robot trajectory estimation and mapping.

The majority of previous approaches to trajectory estimation and mapping, including the smoothing-based SAM family of algorithms, have formulated the problem in discrete time [1–4,6,8,9]. However, discrete-time representations are restrictive: they are not easily extended to trajectories with irregularly spaced way-points or asynchronously sampled measurements. A continuous-time formulation of the SAM problem where measurements constrain the trajectory at any point in time, would elegantly contend with these difficulties. Viewed from this perspective, the robot trajectory is a *function* $\mathbf{x}(t)$, that maps any time t to a robot

* Corresponding author.

E-mail address: xinyan.yan@cc.gatech.edu (X. Yan).

state. The problem of estimating this function along with landmark locations has been dubbed *simultaneous trajectory estimation and mapping* (STEAM) [10].

Tong et al. [11] proposed a Gaussian process (GP) regression approach to solving the STEAM problem. While their approach was able to accurately model and interpolate asynchronous data to recover a trajectory and landmark estimate, it suffered from significant computational challenges: naive Gaussian process approaches to regression have notoriously high space and time complexity. Additionally, Tong et al.'s approach is a *batch* method, so updating the solution necessitates saving all of the data and completely re-solving the problem. In order to combat the computational burden, Tong et al.'s approach was extended in Barfoot et al. [10] to take advantage of the sparse structure inherent in the STEAM problem. The resulting algorithm significantly speeds up solution time and can be viewed as a continuous-time analog of Dellaert's original square-root SAM algorithm [4]. Unfortunately, like SAM, Barfoot et al.'s GP-based algorithm remains a batch algorithm, which is a disadvantage for robots that need to continually update the estimate of their trajectory and environment.

In this work, we provide the critical extensions necessary to transform the existing Gaussian process-based approach to solving the STEAM problem into an extremely efficient incremental approach. Our algorithm elegantly combines the benefits of Gaussian processes and iSAM 2.0. Like the GP regression approaches to STEAM, our approach can model continuous trajectories, handle asynchronous measurements, and naturally interpolate states to speed up computation and reduce storage requirements, and, like iSAM 2.0, our approach uses a Bayes tree to efficiently calculate a *maximum a posteriori* (MAP) estimate of the GP trajectory while performing incremental factorization, variable reordering, and just-in-time relinearization. The result is an online GP-based solution to the STEAM problem that remains computationally efficient while scaling up to large datasets.

The present paper is an extension of the work presented in [12,13]. As a further contribution, in this manuscript we elaborate more on variable re-ordering that is key to making both batch and incremental GP-regression computationally more efficient, and we study the performance of the proposed approach in an additional real-world dataset. Furthermore, we release an efficient implementation of the approach developed herein as open source code.¹

2. Batch trajectory estimation & mapping as Gaussian process regression

We begin by describing how the simultaneous trajectory estimation and mapping (STEAM) problem can be formulated in terms of Gaussian process regression. Following Tong et al. [11] and Barfoot et al. [10], we represent robot trajectories as functions of time t sampled from a Gaussian process:

$$\mathbf{x}(t) \sim \mathcal{GP}(\boldsymbol{\mu}(t), \mathcal{K}(t, t')), \quad t_0 < t, t'. \quad (1)$$

Here, $\mathbf{x}(t)$ is the continuous-time trajectory of the robot through state-space, represented by a Gaussian process with mean $\boldsymbol{\mu}(t)$ and covariance $\mathcal{K}(t, t')$ functions.

We next define a finite set of measurements:

$$\mathbf{y}_i = \mathbf{h}_i(\boldsymbol{\theta}_i) + \mathbf{n}_i, \quad \mathbf{n}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_i), \quad i = 1, 2, \dots, N. \quad (2)$$

The measurement $\mathbf{y}_i$ can be any linear or nonlinear functions of a set of related variables $\boldsymbol{\theta}_i$ plus some Gaussian noise $\mathbf{n}_i$. The related variables for a range measurement are the robot state at the corresponding measurement time $\mathbf{x}(t_i)$ and the associated landmark

location ℓ_j . We assume the total number of measurements is N , and the number of trajectory states at measurement times is M .

Based on the definition of Gaussian processes, any finite collection of robot states has a joint Gaussian distribution [14]. So the robot states at measurement times are normally distributed with mean $\boldsymbol{\mu}$ and covariance $\mathcal{K}$.

$$\begin{aligned} \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\mu}, \mathcal{K}), \quad \mathbf{x} = [\mathbf{x}_1^\top \dots \mathbf{x}_M^\top]^\top, \quad \mathbf{x}_i = \mathbf{x}(t_i) \\ \boldsymbol{\mu} &= [\boldsymbol{\mu}_1^\top \dots \boldsymbol{\mu}_M^\top]^\top, \quad \boldsymbol{\mu}_i = \boldsymbol{\mu}(t_i), \quad \mathcal{K}_{ij} = \mathcal{K}(t_i, t_j). \end{aligned} \quad (3)$$

Note that any point along the continuous-time trajectory can be estimated from the Gaussian process model. Therefore, the trajectory does not need to be discretized and robot trajectory states do not need to be evenly spaced in time, which is an advantage of the Gaussian process approach over discrete-time approaches (e.g. Dellaert's square-root SAM [4]).

The landmarks ℓ which represent the map are assumed to conform to a joint Gaussian distribution with mean $\mathbf{d}$ and covariance $\mathbf{W}$ (Eq. (4)). The prior distribution of the combined state $\boldsymbol{\theta}$ that consists of robot trajectory states at measurement times and landmarks is, therefore, a joint Gaussian distribution (Eq. (5)).

$$\ell \sim \mathcal{N}(\mathbf{d}, \mathbf{W}), \quad \ell = [\ell_1^\top \ell_2^\top \dots \ell_N^\top]^\top \quad (4)$$

$$\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\eta}, \mathcal{P}), \quad \boldsymbol{\eta} = [\boldsymbol{\mu}^\top \mathbf{d}^\top]^\top, \quad \mathcal{P} = \begin{bmatrix} \mathcal{K} & \\ & \mathbf{W} \end{bmatrix}. \quad (5)$$

To solve the STEAM problem, given the prior distribution of the combined state and the likelihood of measurements, we compute the *maximum a posteriori* (MAP) estimate of the combined state conditioned on measurements via Bayes' rule:

$$\begin{aligned} \boldsymbol{\theta}^* &\triangleq \boldsymbol{\theta}_{\text{MAP}} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} p(\boldsymbol{\theta}|\mathbf{y}) = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \frac{p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})}{p(\mathbf{y})} \\ &= \underset{\boldsymbol{\theta}}{\operatorname{argmax}} p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}) = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} (-\log p(\boldsymbol{\theta}) - \log p(\mathbf{y}|\boldsymbol{\theta})) \\ &= \underset{\boldsymbol{\theta}}{\operatorname{argmin}} (\|\boldsymbol{\theta} - \boldsymbol{\eta}\|_{\mathcal{P}}^2 + \|\mathbf{h}(\boldsymbol{\theta}) - \mathbf{y}\|_{\mathbf{R}}^2) \end{aligned} \quad (6)$$

where the norms are Mahalanobis norms defined as: $\|\mathbf{e}\|_{\Sigma}^2 = \mathbf{e}^\top \Sigma^{-1} \mathbf{e}$, and $\mathbf{h}(\boldsymbol{\theta})$ and $\mathbf{R}$ are the mean and covariance of the measurements collected, respectively:

$$\mathbf{h}(\boldsymbol{\theta}) = [\mathbf{h}_1(\boldsymbol{\theta}_1)^\top \mathbf{h}_2(\boldsymbol{\theta}_2)^\top \dots \mathbf{h}_N(\boldsymbol{\theta}_N)^\top]^\top \quad (7)$$

$$\mathbf{R} = \operatorname{diag}(\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N). \quad (8)$$

Because both covariance matrices $\mathcal{P}$ and $\mathbf{R}$ are positive definite, the objective in Eq. (6) corresponds to a least squares problem. Consequently, if some of the measurement functions $\mathbf{h}_i(\cdot)$ are nonlinear, this becomes a nonlinear least squares problem, in which case iterative methods including Gauss-Newton and Levenberg-Marquardt [15] can be utilized; in each iteration, an optimal update is computed given a linearized problem at the current estimate. A linearization of a measurement function at current state estimate $\boldsymbol{\theta}_i$ can be accomplished by a first-order Taylor expansion:

$$\mathbf{h}_i(\bar{\boldsymbol{\theta}}_i + \delta\boldsymbol{\theta}_i) \approx \mathbf{h}_i(\bar{\boldsymbol{\theta}}_i) + \left. \frac{\partial \mathbf{h}_i}{\partial \boldsymbol{\theta}_i} \right|_{\bar{\boldsymbol{\theta}}_i} \delta\boldsymbol{\theta}_i. \quad (9)$$

Combining Eq. (9) with Eq. (6), the optimal increment $\delta\boldsymbol{\theta}^*$ at the current combined state estimate $\bar{\boldsymbol{\theta}}$ is

$$\delta\boldsymbol{\theta}^* = \underset{\delta\boldsymbol{\theta}}{\operatorname{argmin}} \left(\|\bar{\boldsymbol{\theta}} + \delta\boldsymbol{\theta} - \boldsymbol{\eta}\|_{\mathcal{P}}^2 + \|\mathbf{h}(\bar{\boldsymbol{\theta}}) + \mathbf{H}\delta\boldsymbol{\theta} - \mathbf{y}\|_{\mathbf{R}}^2 \right) \quad (10)$$

$$\mathbf{H} = \operatorname{diag}(\mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_N), \quad \mathbf{H}_i = \left. \frac{\partial \mathbf{h}_i}{\partial \boldsymbol{\theta}_i} \right|_{\bar{\boldsymbol{\theta}}_i} \quad (11)$$

where $\mathbf{H}$ is the measurement Jacobian matrix. To solve the linear least squares problem in Eq. (10), we take the derivative with

¹ Please check out the code at <https://github.com/XinyanGT/online-gpslam-code>.

Download English Version:

<https://daneshyari.com/en/article/4948835>

Download Persian Version:

<https://daneshyari.com/article/4948835>

[Daneshyari.com](https://daneshyari.com)