

---

# Incremental Sparse GP Regression for Continuous-time Trajectory Estimation & Mapping

---

**Xinyan Yan**  
College of Computing  
Georgia Institute of Technology  
Atlanta, GA 30332, USA  
xinyan.yan@cc.gatech.edu

**Vadim Indelman**  
Faculty of Aerospace Engineering  
Technion - Israel Institute of Technology  
Haifa 32000, Israel  
vadim.indelman@technion.ac.il

**Byron Boots**  
College of Computing  
Georgia Institute of Technology  
Atlanta, GA 30332, USA  
bboots@cc.gatech.edu

## Abstract

Recent work has investigated the problem of continuous-time trajectory estimation and mapping for mobile robots by formulating the problem as sparse Gaussian process regression. Gaussian processes provide a continuous-time representation of the robot trajectory, which elegantly handles asynchronous and sparse measurements, and allows the robot to query the trajectory to recover its estimated position at any time of interest. One of the major drawbacks of this approach is that Gaussian process regression formulates continuous-time trajectory estimation as a *batch* estimation problem. In this work, we provide the critical extensions necessary to transform this existing batch approach into an extremely efficient incremental approach. 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.

## 1 Introduction & Related Work

The problem of simultaneously recovering the location of a robot and a map of its environment from sensor readings is a fundamental challenge in robotics, persisting as a core research topic for several decades. Well-known approaches to this problem, such as smoothing and mapping (SAM) [1], 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. [2, 3] developed a critical extension to the batch SAM algorithm, called incremental smoothing and mapping (iSAM 2.0), that overcomes this problem by *incrementally* computing a solution. The approach employs an efficient data structure called a *Bayes tree* [4] to perform incremental variable reordering and just-in-time relinearization. iSAM 2.0 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 have formulated the problem in discrete time. However, in the real world, the trajectory of a robot is continuous and sensor measurements are often sampled asynchronously. A continuous-time formulation of the SAM problem where the robot trajectory is a *function*  $x(t)$  that maps any time  $t$  to a robot pose is more appropriate for this problem. The problem of estimating this function along with landmark loca-

tions has been dubbed *simultaneous trajectory estimation and mapping* (STEAM). Tong et al. [5, 6] 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 resolving the problem.

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 and handle asynchronous measurements, 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 extremely large datasets.

## 2 Batch Trajectory Estimation & Mapping as Gaussian Process Regression

Following Tong et al. [5] and Barfoot et al. [6], we take a Gaussian process regression approach to state estimation, where we represent robot trajectories  $\mathbf{x}$  as functions of time  $t$ .

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

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

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  $\mathcal{K}(t, t')$ . Measurements  $\mathbf{y}_i$  are nonlinear functions of the 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  $\mathbf{x}(t)$  at the measurement time and the associated landmark location  $\ell_j$ . The total number of measurements is  $N$ . The landmarks  $\boldsymbol{\ell} = [\ell_1 \ \ell_2 \ \dots \ \ell_M]$  are assumed to be sampled from a joint Gaussian distribution  $\boldsymbol{\ell} \sim \mathcal{N}(\mathbf{d}, \mathbf{W})$  with mean  $\mathbf{d}$  and covariance  $\mathbf{W}$ . The prior distribution of the combined state  $\boldsymbol{\theta}$  that consists of robot poses at measurement times and landmarks is, therefore, a joint Gaussian distribution:

$$\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\eta}, \mathcal{P}), \quad \boldsymbol{\eta} = [\boldsymbol{\mu} \ \mathbf{d}]^\top, \quad \mathcal{P} = \begin{bmatrix} \mathcal{K} & \\ & \mathbf{W} \end{bmatrix} \quad (3)$$

To solve the STEAM problem, we compute the *maximum a posteriori* (MAP) estimate of the combined state *conditioned* on measurements:

$$\begin{aligned} \boldsymbol{\theta}^* &= \underset{\boldsymbol{\theta}}{\operatorname{argmax}} p(\boldsymbol{\theta}|\mathbf{y}) = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} p(\boldsymbol{\theta}, \mathbf{y}) \\ &= \underset{\boldsymbol{\theta}}{\operatorname{argmax}} p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta}) = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} (-\log p(\boldsymbol{\theta}) - \log p(\mathbf{y}|\boldsymbol{\theta})) \\ &= \underset{\boldsymbol{\theta}}{\operatorname{argmin}} ((\boldsymbol{\theta} - \boldsymbol{\eta})^\top \mathcal{P}^{-1}(\boldsymbol{\theta} - \boldsymbol{\eta}) + (\mathbf{y} - \mathbf{h}(\boldsymbol{\theta}))^\top \mathbf{R}^{-1}(\mathbf{y} - \mathbf{h}(\boldsymbol{\theta}))) \end{aligned} \quad (4)$$

where  $\mathbf{h}(\boldsymbol{\theta})$  and  $\mathbf{R}$  are the mean and covariance of the measurements, respectively:

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

If measurement functions  $\mathbf{h}_i(\cdot)$  are nonlinear, nonlinear optimization methods are required. Examples include Gauss-Newton or Levenberg-Marquardt [7], which solve a sequence of approximated linearized problems to approach the minimum iteratively. A linearization of a measurement function at current state estimate  $\boldsymbol{\theta}_i$  can be accomplished by a first-order Taylor expansion:

$$\mathbf{h}_i(\bar{\boldsymbol{\theta}}_i + \delta\boldsymbol{\theta}_i) \approx \mathbf{h}_i(\bar{\boldsymbol{\theta}}_i) + \left. \frac{\partial \mathbf{h}_i}{\partial \boldsymbol{\theta}_i} \right|_{\bar{\boldsymbol{\theta}}_i} \delta\boldsymbol{\theta}_i \quad (5)$$

Combining with Eq. 4, the optimal increment  $\delta\boldsymbol{\theta}^*$  at the current combine state estimate  $\bar{\boldsymbol{\theta}}$  is

$$\delta\boldsymbol{\theta}^* = \underset{\delta\boldsymbol{\theta}}{\operatorname{argmin}} ((\bar{\boldsymbol{\theta}} + \delta\boldsymbol{\theta} - \boldsymbol{\eta})^\top \mathcal{P}^{-1}(\bar{\boldsymbol{\theta}} + \delta\boldsymbol{\theta} - \boldsymbol{\eta}) + (\mathbf{y} - \mathbf{h}(\bar{\boldsymbol{\theta}}) - \mathbf{H}\delta\boldsymbol{\theta})^\top \mathbf{R}^{-1}(\mathbf{y} - \mathbf{h}(\bar{\boldsymbol{\theta}}) - \mathbf{H}\delta\boldsymbol{\theta})) \quad (6)$$

Where  $\mathbf{H}$  is the measurement Jacobian:

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

To solve Eq. 6, we take the derivative with respect to  $\delta\boldsymbol{\theta}$ , and set it to zero, which gives us  $\delta\boldsymbol{\theta}^*$  embedded in a set of linear equations

$$\underbrace{(\mathcal{P}^{-1} + \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H})}_{\mathcal{I}} \delta\boldsymbol{\theta}^* = \underbrace{\mathbf{H}^\top \mathbf{R}^{-1}(\mathbf{y} - \bar{\mathbf{h}}) - \mathcal{P}^{-1}(\bar{\boldsymbol{\theta}} - \boldsymbol{\eta})}_{\mathbf{b}} \quad (7)$$

The positive definite matrix  $\mathcal{P}^{-1} + \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H}$  is the *a posteriori* information matrix, which we label  $\mathcal{I}$ . To solve this set of linear equations for  $\delta\boldsymbol{\theta}^*$ , we don't actually have to calculate the inverse  $\mathcal{I}^{-1}$ . Instead, factorization-based methods can be leveraged for a fast, numerically stable solution. For example,  $\delta\boldsymbol{\theta}^*$  can be found by first performing a Cholesky factorization  $\mathcal{L}\mathcal{L}^\top = \mathcal{I}$ , and then solving  $\mathcal{L}\mathbf{d} = \mathbf{b}$  and  $\mathcal{L}^\top \delta\boldsymbol{\theta}^* = \mathbf{d}$  by back substitution. If  $\mathcal{I}$  is dense, the time complexity of a Cholesky factorization and back substitution are  $O(N^3)$  and  $O(N^2)$  respectively. However, if  $\mathcal{I}$  has sparse structure, then the solution can be found much faster. For example, for a banded matrix, the computation time is  $O(N)$  instead of  $O(N^3)$ . Fortunately, we can guarantee sparsity for the STEAM problem (see Section 2.1 below). At each iteration we update the state  $\bar{\boldsymbol{\theta}} \leftarrow \bar{\boldsymbol{\theta}} + \delta\boldsymbol{\theta}^*$  and repeat the process until convergence.

## 2.1 Sparse Gaussian Process Regression

The efficiency of the Gaussian Process Gauss-Newton algorithm presented in Section 2 is heavily dependent on the choice of kernel. It is well-known that if the information matrix  $\mathcal{I}$  is sparse, then it is possible to very efficiently compute the solution to Eq. 7 [2]. For the batch GP trajectory estimation and mapping problem, Barfoot et al. [6] suggest a kernel matrix with a sparse inverse that is well-suited to the simultaneous trajectory estimation and mapping problem. In particular, Barfoot et al. show that  $\mathcal{K}^{-1}$  is exactly block-tridiagonal when the GP is assumed to be generated by linear, time-varying (LTV) stochastic differential equation (SDE) which we describe here:

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{v}(t) + \mathbf{F}(t)\mathbf{w}(t), \quad \mathbf{w}(t) \sim \mathcal{GP}(\mathbf{0}, \mathbf{Q}_c \delta(t-t')) \quad t_0 < t, t' \quad (8)$$

where  $\mathbf{x}(t)$  is state,  $\mathbf{v}(t)$  is exogenous input,  $\mathbf{w}(t)$  is process noise, and  $\mathbf{F}(t)$  is time-varying system matrix. The process noise  $\mathbf{w}(t)$  is modeled by a Gaussian process, and  $\delta(\cdot)$  is the *Dirac delta function*. (See [6] for details). We consider a specific case of this model in the experimental results in Section 4.1. Assuming the GP is generated by Eq. 8, and the measurements are range and odometry measurements, the sparse information matrix becomes

$$\mathcal{I} = \begin{bmatrix} \mathcal{I}_{xx} & \mathcal{I}_{xl} \\ \mathcal{I}_{xl}^\top & \mathcal{I}_{ll} \end{bmatrix} \quad (9)$$

where  $\mathcal{I}_{xx}$  is block-tridiagonal and  $\mathcal{I}_{ll}$  is block-diagonal.  $\mathcal{I}_{xl}$ 's density depends on the frequency of landmark measurements and how they are carried out.

## 3 Bayes Tree for Fast Incremental Updates to Sparse GP Regression

The iSAM 2.0 algorithm proposed by Kaess et al. [3] was designed to efficiently solve a nonlinear discrete-time trajectory estimation problem. It is an incremental and real-time approach that works by directly operating on the factor graph representation of the SAM problem. The core technology behind this approach is the *Bayes tree* data structure which allows for incremental variable reordering and fluid relinearization [4]. We apply the same data structure to sparse Gaussian process regression in the context of the STEAM problem, thereby eliminating the need for periodic batch computation.

To understand how the Bayes tree can be incorporated into the problem, it is helpful to understand the how the GP estimation problem can be represented as a factor graph. A factor graph is a bipartite graph  $G = (\mathcal{F}, \boldsymbol{\theta}, \mathcal{E})$ , where  $\mathcal{F}$  is the set of factor nodes that encodes all probabilistic constraints on variables, including range measurements, odometry measurements or smoothing priors,  $\boldsymbol{\theta}$  is the set of variable nodes to estimate, and  $\mathcal{E}$  is the set of edges in the graph. Based on the variable dependence captured in the factor graph, the joint probability of variables to estimate is factored as

$$f(\boldsymbol{\theta}) = \prod_i f_i(\boldsymbol{\theta}_i) \quad (10)$$

where  $f_i \in \mathcal{F}$  is one of the factors, and  $\theta_i$  is the set of variables directly connected to  $f_i$ . Then the estimation problem is to find  $\theta^*$  that maximizes Eq. 10. As stated in Section 2.1, Barfoot et al. prove that when GPs are generated by LTV SDE,  $\mathcal{K}^{-1}$  is block-tridiagonal [6]. This leads to an interesting interpretation of the states as Markovian, even though we are using a continuous-time prior. In other words, the Gaussian process smoothing factors only connect consecutive pairs of states, and they result from the underlying continuous-time process model.

The factor graph can be converted to a Bayes net by a *bipartite elimination game* [8]. This procedure is equivalent to converting Eq. 6 to least-squares form and computing the square-root information matrix by Householder reflections or Gram-Schmidt orthogonalization. In order to facilitate marginalization and optimization, a Bayes-tree is constructed from the Bayes net [4]. The Bayes tree groups several variables together based on their dependence, with each node corresponding to a clique in the net. From a linear algebra perspective, the Bayes tree captures the structure of the Cholesky factor  $\mathcal{L}$  of  $\mathcal{I}$ , and the sequence of back substitutions that can be performed. When we add a new measurement, add a prior for new variables, or relinearize a previous measurement,  $\mathcal{L}$  will change accordingly. Importantly, all these modifications to the factor graph only have *local* effects on  $\mathcal{L}$ . Exploiting this observation is the foundation for efficient incremental updates.

Since the nodes of Bayes tree encode conditional probability distributions which directly correspond to rows in  $\mathcal{L}$ , the structure of the tree can be leveraged to efficiently update the factor  $\mathcal{L}$ . The nodes containing the variables that are involved in new factors or whose linear step is larger than a predetermined threshold are identified.<sup>1</sup> Only these nodes and their ascendants in the Bayes tree are then updated. Note that when a sub-tree is updated, variables in the sub-tree are reordered by constrained COLAMD [9] to heuristically maximize the locality of future updates. Last but not least,  $\delta\theta^*$  is computed from tree root to leaves because of the information about the structure of  $\mathcal{L}$  encoded in the Bayes tree. This propagation stops when updates to the conditioning variables are below a predetermined threshold. Algorithm 1 summarizes the incremental Gaussian process regression leveraging the Bayes tree data structure in detail.

---

**Algorithm 1** Updating sparse GP Regression by Bayes tree

---

**while** collecting data **do**

1. Get new measurement results, and identify new factors  $\mathcal{F}_{new}$  and related variables  $\theta_{rel} = \bigcup \theta_i, f_i \in \mathcal{F}_{new}$
2. For each affected variable in  $\theta_{aff} = \theta_{lin} \cup \theta_{rel}$ , remove the corresponding clique and ascendants up to the root of Bayes tree
3. Relinearize all factors required to recreate the removed part of the tree
4. Add cached marginal factors from orphaned sub-trees of removed cliques and create a factor graph
5. Eliminate the factor graph by a new variable ordering, create a Bayes tree, and attach back orphaned sub-trees
6. Partially update estimate from root to leaves and stop walking down a branch when the updates to variables that the child clique is conditioned on are not significant enough
7. Collect variables involved in the measurement factors  $\mathcal{F}_{lin}$  where previous linearization point is far from current estimate,  $\theta_{lin} = \bigcup \theta_i, f_i \in \mathcal{F}_{lin}$

**end while**

---

## 4 Experimental Results

### 4.1 Synthetic SLAM Exploration Task

The data consists of an exploration task with 1,500 process states  $\mathbf{x}(t_i) = [ \mathbf{p}(t_i) \ \dot{\mathbf{p}}(t_i) ]^\top$ ,  $\mathbf{p}(t_i) = [ x(t_i) \ y(t_i) \ \theta(t_i) ]^\top$ , with a constant time interval, 151 landmarks, and 1,500 range and odometry measurements (Fig 1a). Each process state has one odometry and one range measurement. The trajectory is generated from a white noise acceleration prior  $\ddot{\mathbf{p}}(t) = \mathbf{w}(t)$ , i.e. constant velocity prior.

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{F}\mathbf{w}(t) \tag{11}$$

where

$$\mathbf{x}(t) = \begin{bmatrix} \mathbf{p}(t) \\ \dot{\mathbf{p}}(t) \end{bmatrix}, \quad \mathbf{p}(t) = \begin{bmatrix} x(t) \\ y(t) \\ \theta(t) \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}, \quad \mathbf{w}(t) \sim \mathcal{GP}(\mathbf{0}, \mathbf{Q}_c \delta(t - t'))$$

---

<sup>1</sup>The reader is referred to [3] for additional details regarding this just-in-time relinearization.




Figure 1: Solving the STEAM problem on a synthetic dataset. Ground truth (a) and combined state estimate (b). In (a), the lines between trajectory points and landmarks indicate range measurements. At each time step, the robot receives a noisy measurement of the distance between itself and a random nearby landmark. (b) shows that the estimate is very close to the ground truth.




Figure 2: Comparison of time spent on each time step and accumulated time for each of the three approaches – periodic batch relinearization, periodic batch relinearization with variable reordering, and the incremental Bayes tree-based approach. Because the number of landmarks, which is 151, is not negligible compared to the number of process states, variable reordering dramatically improves the performance. In all cases, a new iteration of the algorithm, i.e. an update to the combined state estimate, is carried out when 5 new range measurements have accumulated. So in (a), the computation time in many time steps is close to zero. Notice that the Bayes tree approach has almost no increase in computation time when the size of problem grows.

Note that velocity  $\dot{\mathbf{p}}(t)$  has to be included in process state to represent the motion in LTV SDE form, which leads to Markovian states and block-tridiagonal inverse kernel matrix  $\mathcal{K}^{-1}$ . The odometry and range measurements are specified in Eq. 12 and Eq. 13 respectively.

$$\mathbf{y}_{io} = \begin{bmatrix} \cos \theta(t_{i-1}) \cdot (x(t_i) - x(t_{i-1})) + \sin \theta(t_{i-1}) \cdot (y(t_i) - y(t_{i-1})) \\ \theta(t_i) - \theta(t_{i-1}) \end{bmatrix} + \mathbf{n}_o \quad (12)$$

$$y_{ir} = \left\| \begin{bmatrix} x(t_i) \\ y(t_i) \end{bmatrix} - \ell_{ij} \right\|_2 + n_r \quad (13)$$

where  $\mathbf{y}_{io}$  consists of robot-oriented distance and heading difference between  $t_i$  and  $t_{i-1}$ , and  $y_{ir}$  is the distance between the robot and a specific landmark  $\ell_{ij}$  at  $t_i$ . The results are shown in Fig. 1b. In Fig. 2, we compare the performance of the periodic relinearization approach, a naive periodic relinearization with variable reordering approach, and the Bayes tree approach from Section 3.

## 4.2 Autonomous Lawnmower

In our second experiment, we applied all three approaches to a freely available range-only SLAM dataset collected from an autonomous lawn-mowing robot [10]. The environment, including the




Figure 3: Solving the STEAM problem on a real dataset. Ground truth (a) and estimate (b) of “Plaza” dataset. In (a), the lines between trajectory points and landmarks represent range measurements. The range measurements are sparse, with approximately 11 time steps (and up to 500 steps) between range readings for the worst landmark. (b) shows that ground truth and estimate are very close.




Figure 4: Comparison of time spent on each time step and accumulated time. There are 4 landmarks, which is extremely small relative to the number of process states, so no performance gain is observed from variable reordering. However, in the Bayes tree-based approach, which is fully incremental, the update of each iteration has very limited local effects and the computation time almost remains the same even with a large number of process states. The gap around time step 5,000 in (a) is caused by the vacancy of range measurements in that time period.

locations of the landmarks and the ground truth paths, are shown in Fig. 3a. While collecting the dataset the robot travelled 1.9km, occupied 9,658 poses, and received 3,529 range measurements. The path is a commonly used path pattern in lawn mowing applications. This dataset is very sparse, with approximately 11 time steps (and up to 500 steps) between range readings for the worst landmark. Besides range measurements, the dataset also contains odometry measurements at each time step. And all the measurements are considered in the same manner as in Section 4.1. The results of our algorithm are shown in Fig. 3b and demonstrate that we are able to accurately estimate the robot’s trajectory and map.

Performance of three approaches – periodic batch relinearization, periodic batch relinearization with variable reordering, and incremental Bayes tree are compared in Fig. 4. In all cases, as in Section 4.1, a new iteration, i.e. an update to the combined state estimate, is carried out when 5 new range measurements have accumulated. In this dataset, the number of landmarks are 4, which is extremely small relative to the number of process states, so there is no performance gain from reordering. However, the Bayes tree-based approach dramatically outperforms the other two approaches. As the problem size increases, there is negligible increase in computation time, even for close to 10,000 process states.

## References

- [1] Frank Dellaert and Michael Kaess. Square root sam: Simultaneous localization and mapping via square root information smoothing. *International Journal of Robotics Research*, 25:2006, 2006.
- [2] M. Kaess, A. Ranganathan, and F. Dellaert. isam: Incremental smoothing and mapping. *Robotics, IEEE Transactions on*, 24(6):1365–1378, Dec 2008.
- [3] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J.J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research, IJRR*, 31(2):217–236, Feb 2012.
- [4] Michael Kaess, Viorela Ila, Richard Roberts, and Frank Dellaert. The bayes tree: An algorithmic foundation for probabilistic robot mapping. In *Algorithmic Foundations of Robotics IX*, pages 157–173. Springer, 2011.
- [5] Chi Hay Tong, Paul Furgale, and Timothy D Barfoot. Gaussian process gauss–newton for non-parametric simultaneous localization and mapping. *The International Journal of Robotics Research*, 32(5):507–525, 2013.
- [6] Tim Barfoot, Chi Hay Tong, and Simo Sarkka. Batch continuous-time trajectory estimation as exactly sparse gaussian process regression. In *Proceedings of Robotics: Science and Systems*, Berkeley, USA, July 2014.
- [7] J. E. Dennis, Jr. and Robert B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations (Classics in Applied Mathematics, 16)*. Soc for Industrial & Applied Math, 1996.
- [8] Pinar Heggernes and Pontus Matstoms. Finding good column orderings for sparse qr factorization. Technical report, In Second SIAM Conference on Sparse Matrices, 1996.
- [9] Timothy A. Davis, John R. Gilbert, Stefan I. Larimore, and Esmond G. Ng. Algorithm 836: Colamd, a column approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw.*, 30(3):377–380, September 2004.
- [10] Joseph Djugash. *Geolocation with Range: Robustness, Efficiency and Scalability*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, November 2010.