Incremental Sparse Gaussian Process Regression for Continuous-time Trajectory Estimation & Mapping

Xinyan Yan [1], Vadim Indelman [2], Byron Boots [1]

[1] College of Computing
Georgia Institute of Technology

[2] Faculty of Aerospace Engineering
Technion - Israel Institute of Technology

1 Motivation

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. We want:

- a *continuous-time representation* of the robot trajectory, which elegantly handles asynchronous and sparse measurements, and
- an efficient *incremental approach* to updating the estimate.

2 Batch Trajectory Estimation & Mapping as GP Regression

- Representation[1]
  - robot trajectory: $x(t) \sim \mathcal{GP}\big(\mu(t), K(t,t')\big)$
  - mm. model: $y_i = h_i(\theta_i) + n_i \quad n_i \sim \mathcal{N}(0, R_i)$
  - robot states at mm.: $x \sim \mathcal{N}(\mu, K)$ landmarks (map): $\ell \sim \mathcal{N}(d, W)$
  - combined states: $\theta \sim \mathcal{N}(\eta, P) \quad \theta = [\,x^T \;\ell^T]^T \quad \eta = [\,\mu^T \; d^T]^T \quad P = \begin{bmatrix} K & \\ & W \end{bmatrix}$
- MAP estimate of the combined states:
  - nonlinear ls: $\theta^* = \underset{\theta}{\operatorname{argmax}}\, p(\theta|y) = \operatorname{argmin}(\|\theta-\eta\|_P^2 + \|y - h(\theta)\|_R^2)$
- Gauss-Newton method, *batch* update: $\bar\theta \leftarrow \bar\theta + \delta\theta^*$. At each iteration:
  - linearized mm.: $h_i(\bar\theta_i + \delta\theta_i) \approx h_i(\bar\theta_i) + H_i\delta\theta_i \qquad H_i = \frac{\partial h_i}{\partial \theta_i}\Big|_{\bar\theta_i}$
  - linearized ls: $\delta\theta^* = \underset{\delta\theta}{\operatorname{argmin}}\left(\|\bar\theta + \delta\theta - \eta\|_P^2 + \|y - h(\bar\theta) - H\delta\theta\|_R^2\right)$
  - solution embedded: $\underbrace{(P^{-1} + H^T R^{-1} H)}_{I}\,\delta\theta^* = \underbrace{H^T R^{-1}(y - \bar h) - P^{-1}(\bar\theta - \eta)}_{b}$
  - solved by Cholesky decomposition ($I = L^T L$) and back substitution

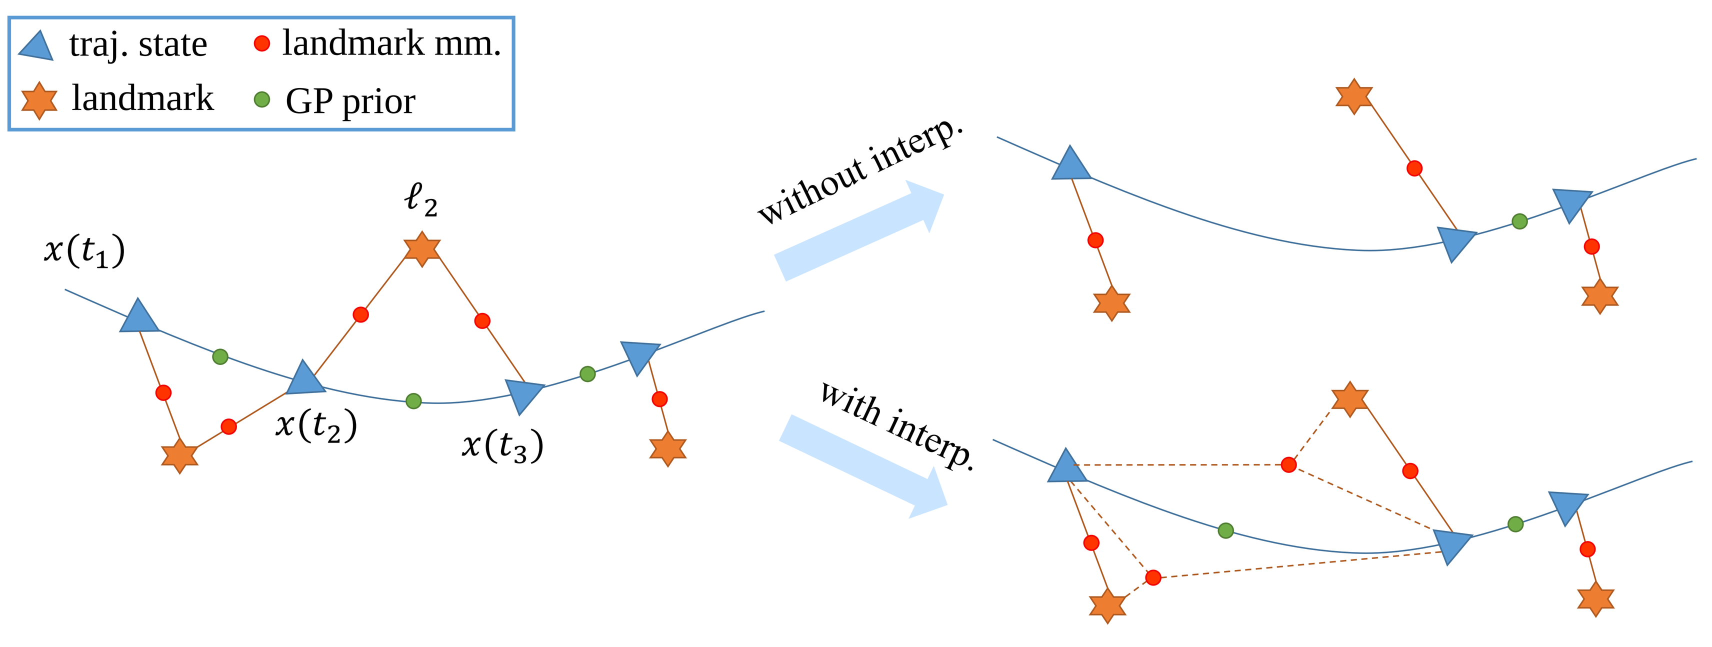

Figure 1 A simple factor graph, with $x(t_2)$ being a missing state.

2A State Interpolation

Any state can be interpolated from other states by computing the posterior mean:

interpolation : $\bar x(t) = \mu(t) + K(t)K^{-1}(\bar x - \mu)$

linearized mm. $h_i(\bar\theta_i + \delta\theta_i) = h_i\big(\bar x(\tau) + \delta x(\tau)\big)$
$$\approx h_i\big(\mu(\tau) + K(\tau)K^{-1}(\bar x - \mu)\big) + H_i K(\tau) K^{-1}\delta x$$

We can utilize measurement i without explicitly estimating the states it relates to.

2B Sparse Gaussian Process Regression

K^{-1} is exactly *block-tridiagonal* when the GP is generated by linear, time-varying (LTV) stochastic differential equation (SDE)[2]:

LTV-SDE: $\dot x(t) = A(t)x(t) + v(t) + F(t)w(t)$

$K(\tau)K^{-1}$ has a specific pattern: only two column blocks are non-zero. $t_{i-1} < \tau < t_i$.

sparse pattern: $K(\tau)K^{-1} = [0 \dots 0 \;\; \Lambda(\tau) \;\; \Psi(\tau) \;\; 0 \dots 0]$

interpolation : $\bar x(\tau) = \mu(\tau) + [\Lambda(\tau) \;\; \Psi(\tau)]\left(\begin{bmatrix}\bar x(t_{i-1}) \\ \bar x(t_i)\end{bmatrix} - \begin{bmatrix}\mu(t_{i-1}) \\ \mu(t_i)\end{bmatrix}\right)$

linearized mm. $h_k(\bar\theta_k + \delta\theta_k) = h_k\big(\bar x(\tau) + \delta x(\tau)\big)$
$$\approx h_k\big(\bar x(\tau)\big) + H_k[\Lambda(\tau) \;\; \Psi(\tau)]\begin{bmatrix}\delta x(t_{i-1}) \\ \delta x(t_i)\end{bmatrix}$$

3 The Bayes Tree for Incremental Updates to Sparse GPR

In order to incrementally update the combined state, we utilize a Bayes tree[3]. It leverages a *factor graph* interpretation of the problem to directly update the Cholesky factor L with just-in-time relinearization while maintaining sparsity.

factorization: $f(\theta) = \prod_i f_i(\theta_i)$

GP factor: $f_j(\theta_j) = f_j(x(t_{i-1}, t_i)) \propto \exp\{-\frac{1}{2}\|\Phi(t_i, t_{i-1})x(t_{i-1}) + v_i - x(t_i)\|_{Q_i}^2\}$

mm factor: $f_j(\theta_j) \propto \exp\{-\frac{1}{2}\|h_k(\bar x(\tau) + \delta x(\tau)) - y_k\|\} \qquad \theta_j \triangleq \delta x(\tau)$

interp. $f_j(\theta_j) \propto \exp\{-\frac{1}{2}\|h_k(\bar x(\tau)) + H_k K(\tau)K^{-1}\delta x - y_k\|_{R_k}^2 \qquad \theta_j \triangleq \delta x$

sparse interp. $f_j(\theta_j) \propto \exp\{-\frac{1}{2}\|h_k(\bar x(\tau)) + H_k[\Lambda(\tau) \;\; \Psi(\tau)]\theta_j - y_k\|_{R_k}^2 \qquad \theta_j \triangleq \begin{bmatrix}\delta x(t_{i-1}) \\ \delta x(t_i)\end{bmatrix}$

4 Experimental Results

We evaluate the utilization of Bayes tree and interpolation for incremental updates on both synthetic and real datasets.

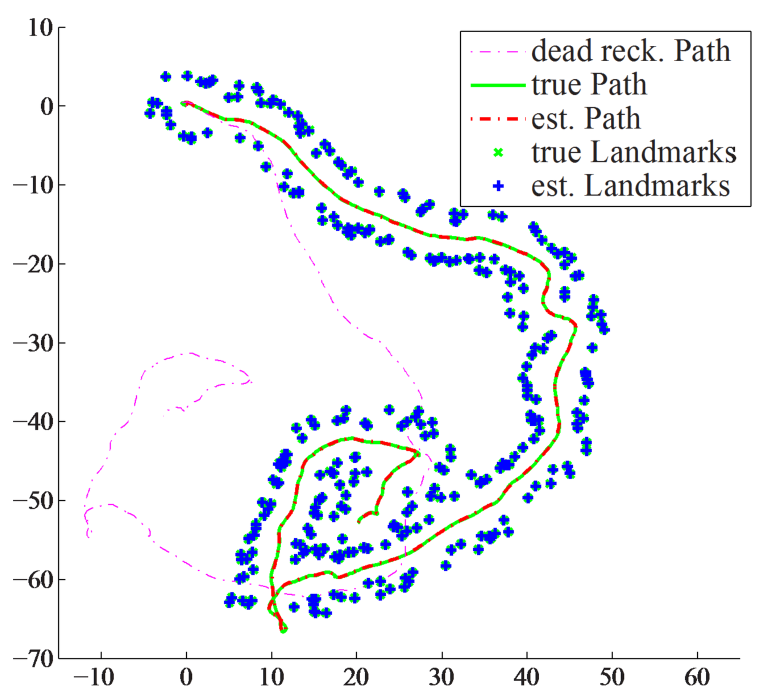

Figure 2 Synthetic dataset, 1,500 ts

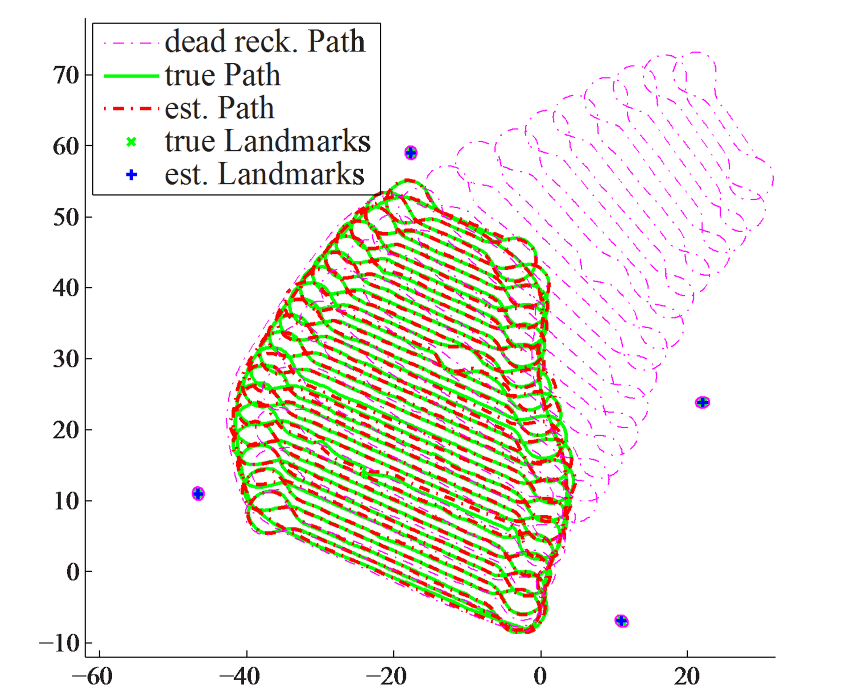

Figure 3 Autonomous mower dataset, 9,658 ts

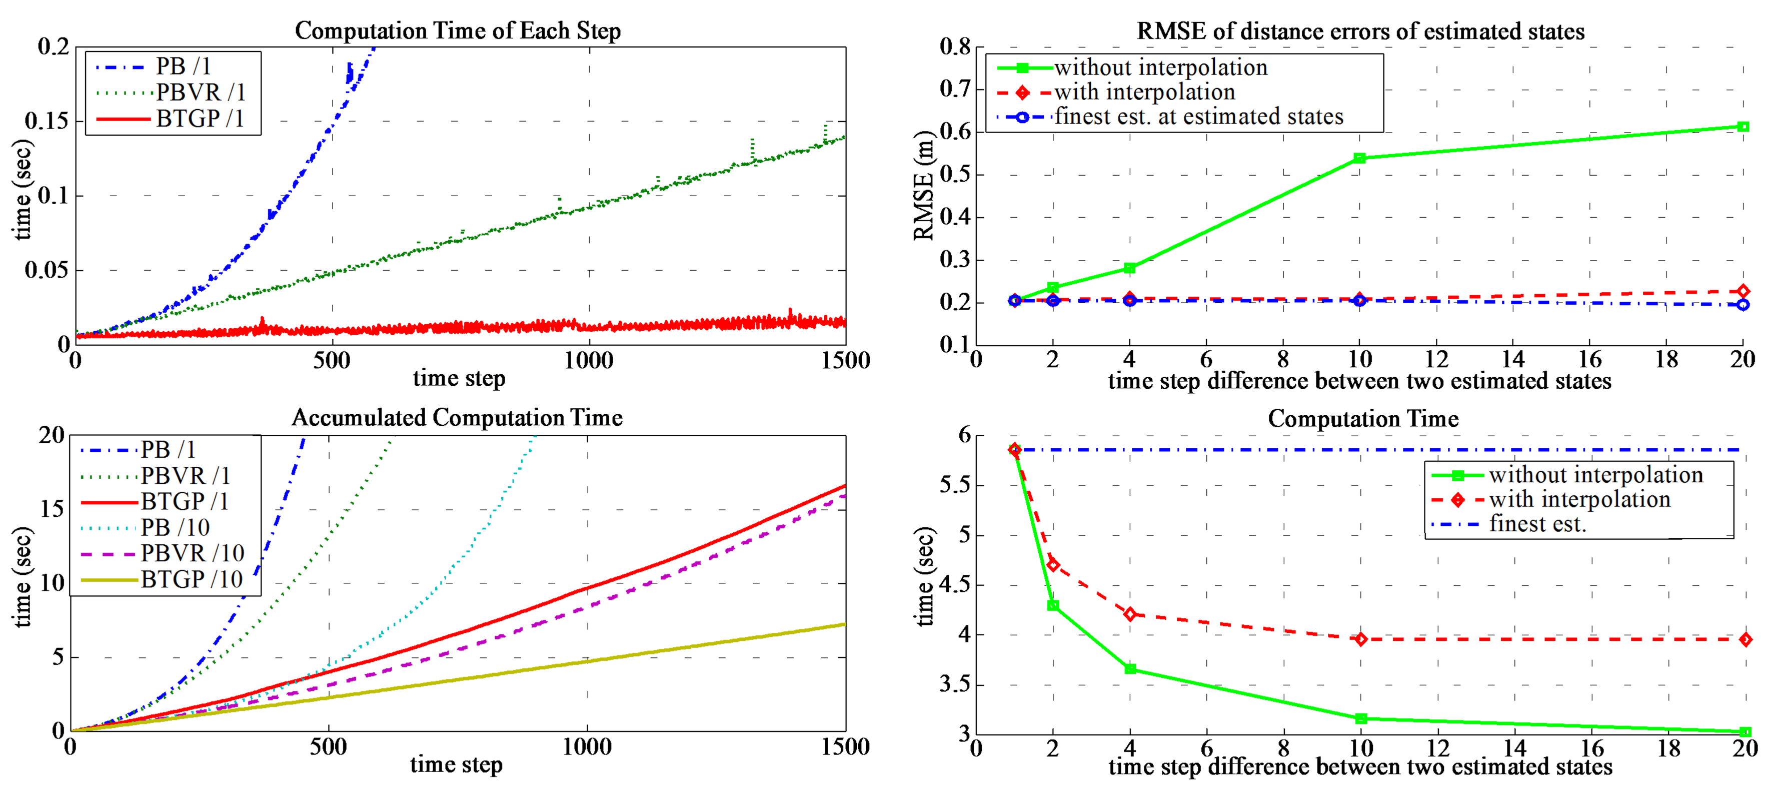

Figure 4 Performance on the synthetic dataset

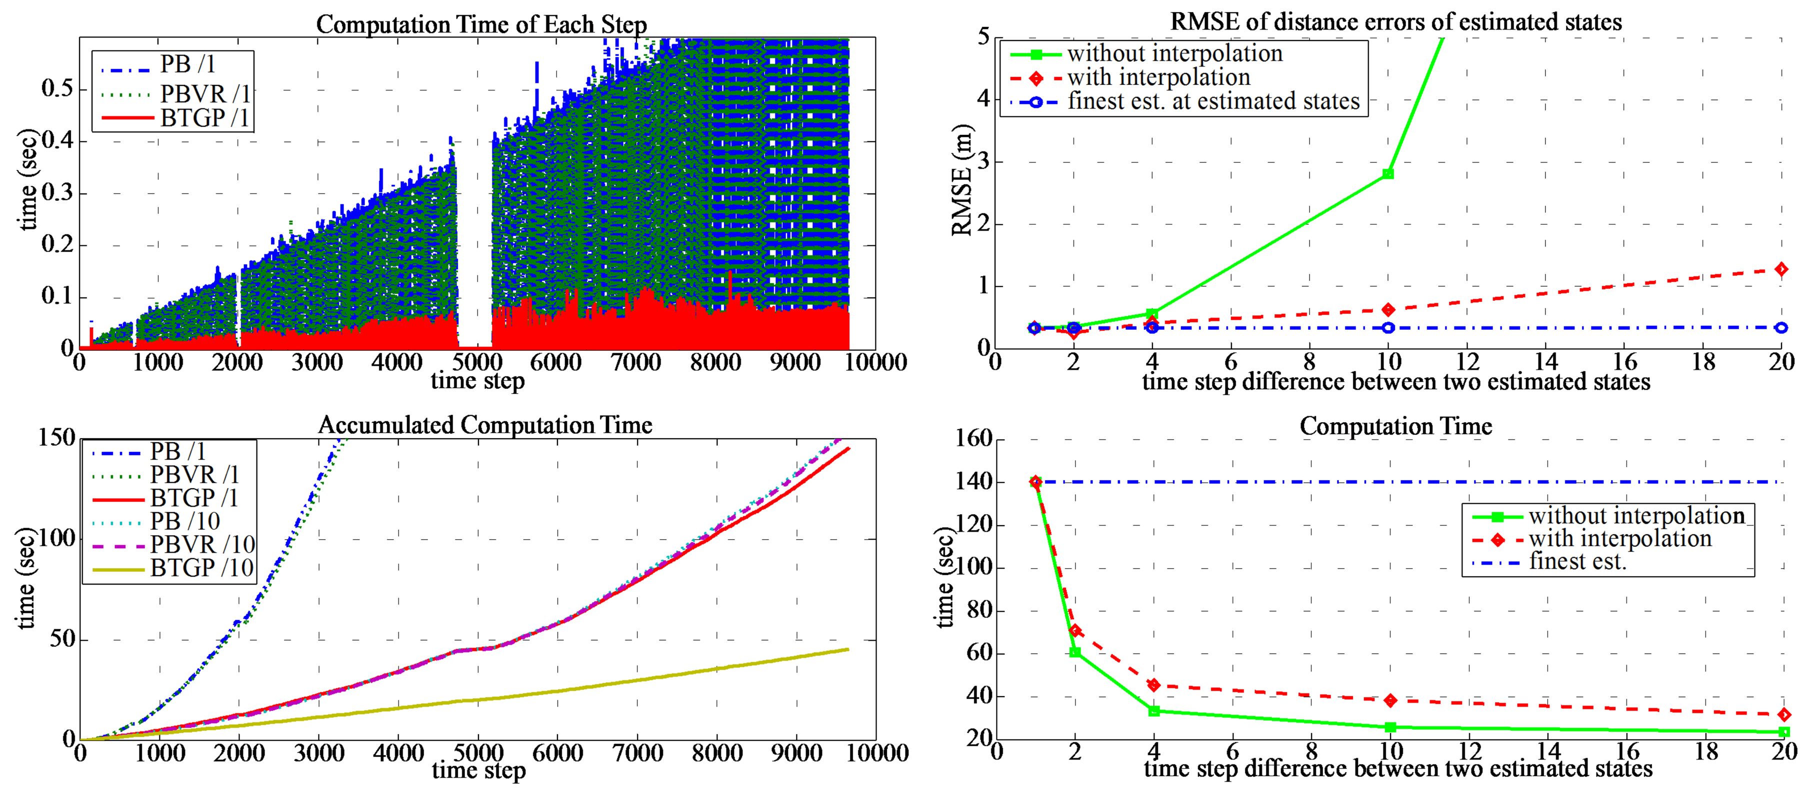

Figure 5 Performance on the autonomous mower dataset

5 Conclusion

We have introduced an incremental sparse GPR algorithm that elegantly combines

- the benefits of continuous-time GP-based approaches, and
- the state-of-the-art innovations from incremental discrete-time algorithms for smoothing and mapping.

References

[1] Chi Hay Tong, Paul Furgale, and Timothy D Barfoot. Gaussian process gauss–newton for non-parametric simultaneous localization and mapping. IJRR, 32(5):507–525, 2013.

[2] Tim Barfoot, Chi Hay Tong, and Simo Sarkka. Batch continuous-time trajectory estimation as exactly sparse gaussian process regression. RSS, Berkeley, USA, July 2014.

[3] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J.J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. IJRR, 31(2):217–236, Feb 2012.