Efficient Spatial Prediction Using Gaussian Markov Random Fields Under Uncertain Localization

Mahdi Jadaliha[1], Yunfei Xu[1], and Jongeun Choi[1,2]

[1]Department of Mechanical Engineering
Michigan State University
[2]Department of Electrical and Computer Engineering
Michigan State University

2012 ASME Dynamic Systems and Control Conference
October 17 - 19, 2012
Fort Lauderdale, FL

Outline

Background

- mobile robotic sensors in environmental monitoring
- statistically model physical phenomena
- cheap sensor networks are prone to localization uncertainty
- The significant computational complexity due to the growing number of observations

Objectives

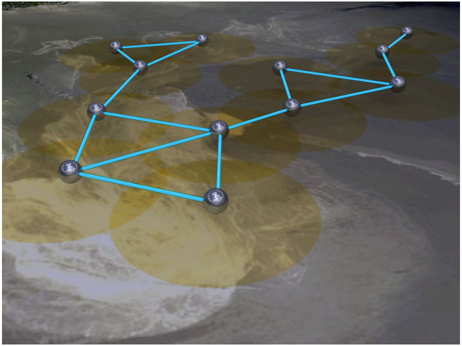

- predicting a spatio-temporal random field
- using sequential noisy observation
- incorporating the effects of localization uncertainty in the prediction

Uncertain localization

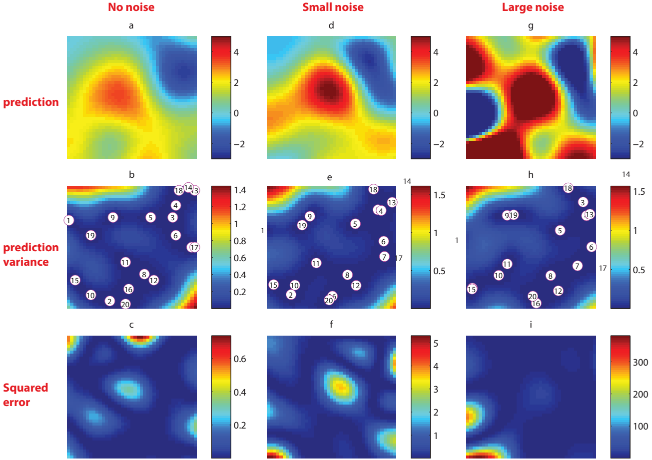

The first column is under the true sampling positions. Second and third columns are under the noisy sampling positions with $\Sigma_1 = 0.1I$ and $\Sigma_2 = 0.4I$ noise covariance matrices, respectively.

Sensor network

The measurement model is given by

$$y^{[i]} := y(q_c^{[i]}) = z(q_c^{[i]}) + \epsilon^{[i]}, \forall i = 1, \cdots, N$$

q_c is uncertain and determined by a prior probability distribution $\pi(q_c)$.

$\pi(q_c)$ could be the output of a common localization algorithm such as Kalman filter, SLAM, and etc.

Gaussian process regression

The posterior distribution for $z \in \mathbb{R}^n$ given true positions, is

$$z|q_c, y \sim \mathcal{N}(\mu, \Sigma).$$

The predictive mean $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n\times n}$ can by obtained by

$$\mu = \lambda + K^T C^{-1}(y - \lambda), \quad \Sigma = \Sigma_0 - K^T C^{-1} K,$$

where the covariance matrices are defined as :
$K := \mathrm{Cov}(y, z) \in \mathbb{R}^{N\times n}$, $C := \mathrm{Cov}(y, y) \in \mathbb{R}^{N\times N}$, $\Sigma_0 := \mathrm{Cov}(z, z) \in \mathbb{R}^{n\times n}$.

Then the predictive distribution of z given the measured locations is

$$\pi(z|\tilde{q}_c, y) = \int_{q\in\mathcal{S}_c} \pi(z|q, y)\pi(q|\tilde{q}_c, y)dq,$$

Discretization

We discretize the compact domain $\mathcal{S}_c := [0 \; x_{max}] \times [0 \; y_{max}]$ into n spatial sites, where $n = hx_{max} \times hy_{max}$.

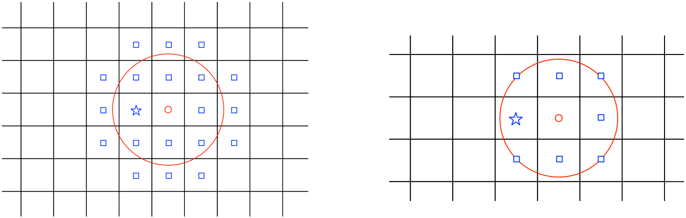

(a) h_1 (b) h_2

Spatio-temporal field

Video

The value of the scalar field is modeled by

$$z_t{}^{[i]} = \lambda_t^{[i]} + \eta_t^{[i]}, \forall i \in \{1, \cdots, n\}, t \in \mathbb{Z}_{>0}.$$

The prior distribution of η_t is given by $\eta_t \sim \mathcal{N}(0, \Sigma_0)$, and so we have

$$z_t \sim \mathcal{N}\left(\lambda_t, \Sigma_0^{-1}\right),$$

where $\Sigma_0 \in \mathbb{R}^{n \times n}$ is the covariance matrix, or $Q_\theta = \Sigma_0^{-1}$ is the precision matrix.

Mean function

Here the mean function $\lambda_t^{[i]} : \mathcal{S} \times \mathbb{Z}_{>0} \to \mathbb{R}$ is defined as

$$\lambda_t^{[i]} = f(s^{[i]})^T \beta_t,$$

where $f(s^{[i]})$ is a known regression function and β_t is an unknown vector of regression coefficients.

The time evolution of $\beta_t \in \mathbb{R}^p$ is modeled by

$$\beta_{t+1} = A_t \beta_t + B_t \omega_t,$$

With precise localization

Assumptions:

A.1 The spatio-temporal random field is generated by the proposed model in the previous slides.

A.2 The precision matrix Q_θ is a given function of an uncertain hyperparameter vector θ.

A.3 The noisy measurements $\{y_t\}$ are continuously collected by robotic sensors in time t.

A.4 The sample positions $\{q_t\}$ are measured precisely by robotic sensors in time t.

A.5 The prior distribution of the hyperparameter vector θ is discrete with a support $\Theta = \{\theta^{(1)}, \cdots, \theta^{(L)}\}$.

Problem 1: Consider the assumptions A.1-A.5. Our problem is to find the predictive mean, and variance of the spatio-temporal field, using successive noisy measurements, precise localization and uncertain hyperparameters.

Solution to problem 1 (Algorithm 1)

Prediction:

$$\mu_{x_t|\theta,\mathcal{D}_{1:t-1}} = \begin{pmatrix} F_s \mu_{\beta_t|\theta,\mathcal{D}_{1:t-1}} \\ \mu_{\beta_t|\theta,\mathcal{D}_{1:t-1}} \end{pmatrix},$$

$$Q_{x_t|\theta,\mathcal{D}_{1:t-1}} = \begin{pmatrix} Q_\theta & -Q_\theta F_s \\ -F_s^T Q_\theta & F_s^T Q_\theta F_s + \Sigma^{-1}_{\beta_t|\theta,\mathcal{D}_{1:t-1}} \end{pmatrix},$$

- $\mu_{\beta_t|\theta,\mathcal{D}_{1:t-1}} = A_t \mu_{\beta_{t-1}|\theta,\mathcal{D}_{1:t-1}}$ denotes the expectation of β_t
- $\Sigma_{\beta_t|\theta,\mathcal{D}_{1:t-1}} = A_t \Sigma_{\beta_{t-1}|\theta,\mathcal{D}_{1:t-1}} A_t^T + B_t W B_t^T$ denotes the associated estimation error covariance matrix.

Correction:

$$\begin{aligned} Q_{x_t|\theta,\mathcal{D}_{1:t}} &= Q_{x_t|\theta,\mathcal{D}_{1:t-1}} + \sigma_\epsilon^{-2} \Gamma_{q_t} \Gamma_{q_t}^T, \\ \mu_{x_t|\theta,\mathcal{D}_{1:t}} &= \mu_{x_t|\theta,\mathcal{D}_{1:t-1}} + \sigma_\epsilon^{-2} Q^{-1}_{x_t|\theta,\mathcal{D}_{1:t}} \Gamma_{q_t} (y_t - \Gamma_{q_t}^T \mu_{x_t|\theta,\mathcal{D}_{1:t-1}}). \end{aligned}$$

Uncertain hyperparameters

The posterior distribution of the hyperparameter vector θ:

$$\pi(\theta|\mathcal{D}_{1:t}) \propto \pi(y_t|\theta, \mathcal{D}_{1:t-1}, q_t)\pi(\theta|\mathcal{D}_{1:t-1}),$$

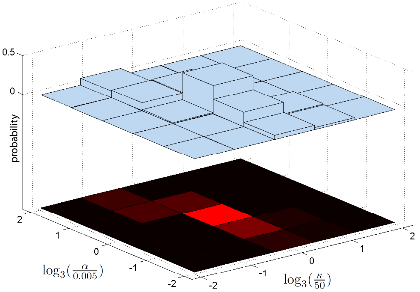

The predictive mean and variance:

$$\mu_{x_t|\mathcal{D}_{1:t}} = \sum_{\theta\in\Theta} \mu_{x_t|\theta,\mathcal{D}_{1:t}} \pi(\theta|\mathcal{D}_{1:t}),$$

$$\Sigma_{x_t|\mathcal{D}_{1:t}} = \sum_{\theta\in\Theta} \left[\Sigma_{x_t|\theta,\mathcal{D}_{1:t}} + (\mu_{x_t|\theta,\mathcal{D}_{1:t}} - \mu_{x_t|\mathcal{D}_{1:t}})(\mu_{x_t|\theta,\mathcal{D}_{1:t}} - \mu_{x_t|\mathcal{D}_{1:t}})^T\right] \pi(\theta|\mathcal{D}_{1:t}).$$

With uncertain localization

Assumptions:

A.4 ~~The sample positions $\{q_t\}$ are measured precisely by robotic sensors in time t.~~

A.6 The prior distribution $\pi(q_t)$ is discrete with a support $\Omega(t) = \{q_t^{(k)} | k \in \mathcal{I}(t)\}$, which is given at time t along with the corresponding measurement y_t.

Problem 2: Consider the assumptions A.1-A.3 and A.5-A.6. Our problem is to find the predictive mean, and variance of the spatio-temporal field, using successive noisy measurements, uncertain localization and uncertain hyperparameters.

Solution to problem 2 (Algorithm 2)

The posterior distribution of q_t:

$$\pi\left(q_{1:t-1}^{(n)}, q_t^{(k)} | \mathcal{R}_{1:t}\right) \propto \pi(q_{1:t-1}^{(n)} | \mathcal{R}_{1:t-1}) \pi(y_t | \mathcal{D}_{1:t-1}^{(n)}, q_t^{(k)}) \pi(q_t^{(k)}).$$

The predictive mean and variance:

$$\begin{aligned} \mu_{x_t|\mathcal{R}_{1:t}} &= \sum_{i \in \mathcal{I}(1:t)} \mu_{x_t|\mathcal{D}_{1:t}^{(i)}} \pi\left(q_{1:t}^{(i)} | \mathcal{R}_{1:t}\right), \\ \Sigma_{x_t|\mathcal{R}_{1:t}} &= \sum_{i \in \mathcal{I}(1:t)} \left[\Sigma_{x_t|\mathcal{D}_{1:t}^{(i)}} + \left(\mu_{x_t|\mathcal{D}_{1:t}^{(i)}} - \mu_{x_t|\mathcal{R}_{1:t}}\right)\right. \\ &\quad \left.\left(\mu_{x_t|\mathcal{D}_{1:t}^{(i)}} - \mu_{x_t|\mathcal{R}_{1:t}}\right)^T\right] \pi\left(q_{1:t}^{(i)} | \mathcal{R}_{1:t}\right), \end{aligned}$$

Simulation results

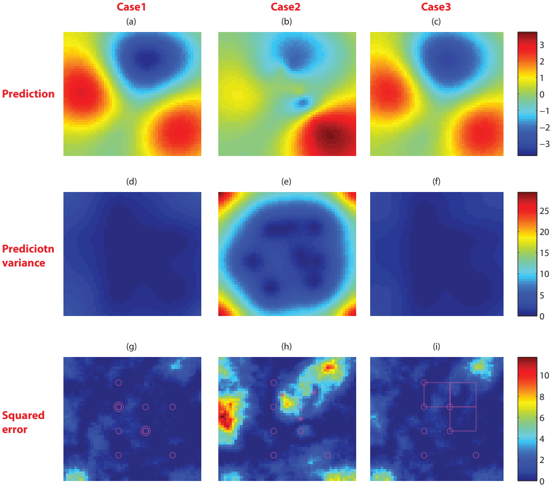

- Case 1: using Algorithm 1 with exact sampling positions
- Case 2: applying Algorithm 1 naively to the measured noisy sampling positions
- Case 3: applying Algorithm 2 to the uncertain sampling positions

Simulation results

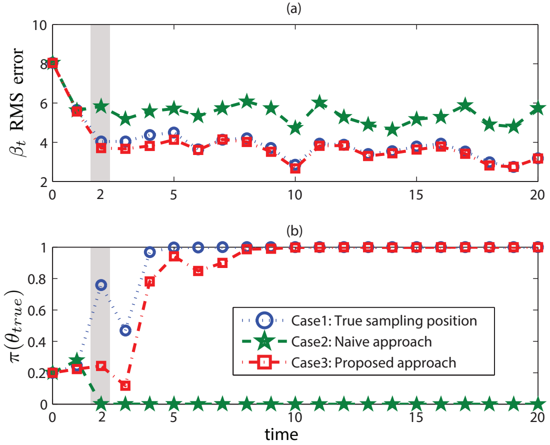

(a) The RMS estimation error of β_t v.s. time and (b) the posterior probability of the true hyperparameter vector v.s. time.

Implementation

- Environmental sensors
- Computer
- Micro-controller
- GPS/INS modules
- Communication modules
- Batteries

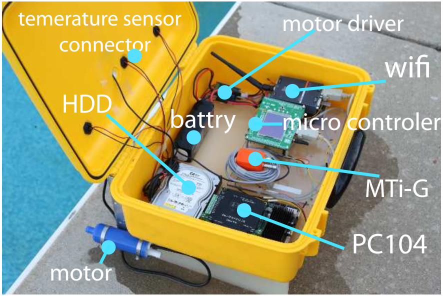

Experimental results

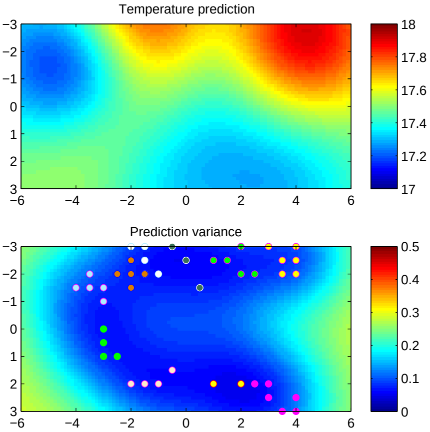

Video

- The experimental environment is a 12 x 6 meters outdoor swimming pool.
- All possible sampling positions for each observation are represented with the same color.

Conclusion

- We have tackled a problem of predicting a spatio-temporal field using successive noisy measurements, uncertain hyperparameters, and uncertain localization.
- We developed the spatio-temporal field of interest using a GMRF and designed sequential prediction algorithms for computing the exact and approximated predictive inference from a Bayesian point of view.
- The most important contribution is that the computation times for Algorithm 1 and Algorithm 2 do not grow as the number of measurements increases.

Acknowledgment

- This work has been supported, in part, by the National Science Foundation through CAREER Award CMMI-0846547. This support is gratefully acknowledged.

Thank you!