Gaussian Processes Performance Bounds for Decentralized Control with Intermittent Communications

Silvia Ferrari
Professor of Engineering and Computer Science
Department of Mechanical Engineering and Materials Science
Duke University

ONR MURI: Nonparametric Bayesian Models to Represent Knowledge and Uncertainty for Decentralized Planning

Year 3 Review Meeting, MIT, Boston (MA)
September 29, 2014

Background and Motivation

Previous Work (Years 1-2): BNP-based sensor planning and control for modeling agent behaviors in the presence of significant uncertainties

- Learn BNP models of dynamic agents
- BNP model size and parameters are both learned from data

Motivation: BNP (e.g. DP-GP and DDP-GP) models can be used to learn trajectory and velocity fields from data

- Dirichlet process mixture (DP-GP) infers number of trajectory field classes
- Dependent Dirichlet process mixture (DDP-GP) extends to time-varying trajectory fields

Previous Technical Accomplishments (Years 1-2):

- Developed **information value functions** for DP-GP models of dynamic targets
- Developed **planning/control** algorithms for multiple dynamic sensors constructing DP-GP models of multiple dynamic targets
- Developed **decentralized planning/control** algorithms for DP-GP models of dynamic targets for multiple sensors with continuous state and control, tracking multiple moving targets

Technical Accomplishments – Year 3

- Analyzed **properties** of DP-GP **information value functions**
(1) Closed-form approximation and bounds for DP-GP expected KL divergence
(2) Performance analysis for different target motion models and prior information

- Developed **optimized visibility-based motion planning** algorithms
(1) Plan the motion of mobile robots based on objectives of exteroceptive sensors with bounded field-of-view
(2) Objectives: simultaneous localization and tracking of a moving target

- Developed **decentralized communication control** methods
(1) Decentralized GP learning with intermittent (controllable) communications
(2) Nominal performance of decentralized GP learning by a sensor network
(3) Plan the motion and communication times of decentralized sensor networks
(4) Guaranteed network performance bounds with intermittent communications

Benchmark Example Problem

Problem: Control sensor(s) mode and motion for modeling dynamic **target behaviors** in a workspace $\mathcal{W}$, via DP-GP mixtures

- Plan admissible **control inputs**, $\mathbf{u}$: - field-of-view (FOV) configuration, $\mathbf{q}$
  - pan-tilt-zoom (PTZ) variables, $\boldsymbol{l}$
  - communication time, t_c
- **Optimize sensor objectives**, φ : - DP-GP information value (KL)

Targets are non-cooperative, independent and obey a time-invariant velocity field:

$$\dot{\mathbf{x}}_j(t) = \mathbf{f}_i[\mathbf{x}_j(t)] \equiv \mathbf{v}_j(t), \quad j = 1, \ldots, N(t) \quad \rightarrow \{\mathcal{F}, \boldsymbol{\pi}\}$$ **"DP-GP target model"**

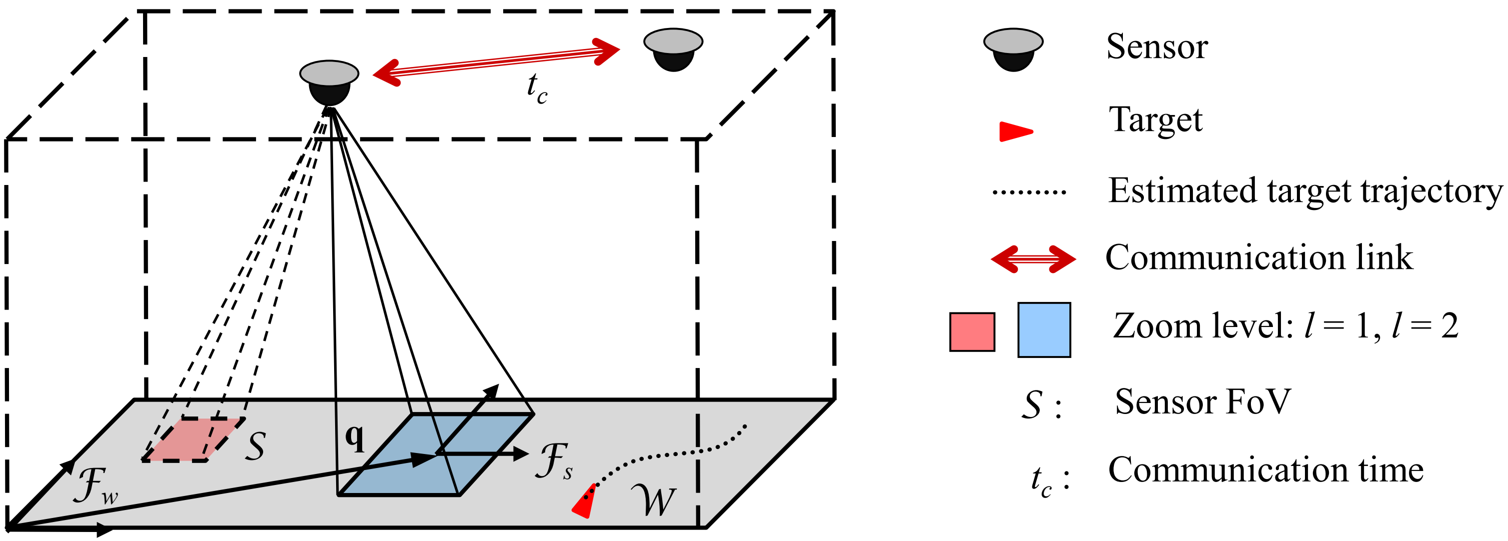

DP-GP Target Model

Sensor objectives: construct DP-GP model of target dynamics from data

- Noisy position and velocity measurements (perfect data-target association)
- Velocity field: 2D spatial phenomenon → Gaussian process

  $\dot{\mathbf{x}}_j(t) = \mathbf{f}_i[\mathbf{x}_j(t)] \equiv \mathbf{v}_j(t), \quad j = 1,\ldots,N(t)$
- Target – velocity field association unknown
- Clustering → Dirichlet process

DP-GP mixture model[1]:

$$\{\boldsymbol{\theta}_i, \boldsymbol{\pi}\} \sim \mathrm{DP}(\alpha, \mathrm{GP}_0), \; i = 1,\ldots,\infty$$
$$g_j \sim \mathrm{Cat}(\boldsymbol{\pi}), \; j = 1,\ldots,N$$
$$\mathbf{f}_{g_j}(\mathbf{x}) \sim \mathrm{GP}(\boldsymbol{\theta}_{g_j}, \psi), \; j = 1,\ldots,N$$

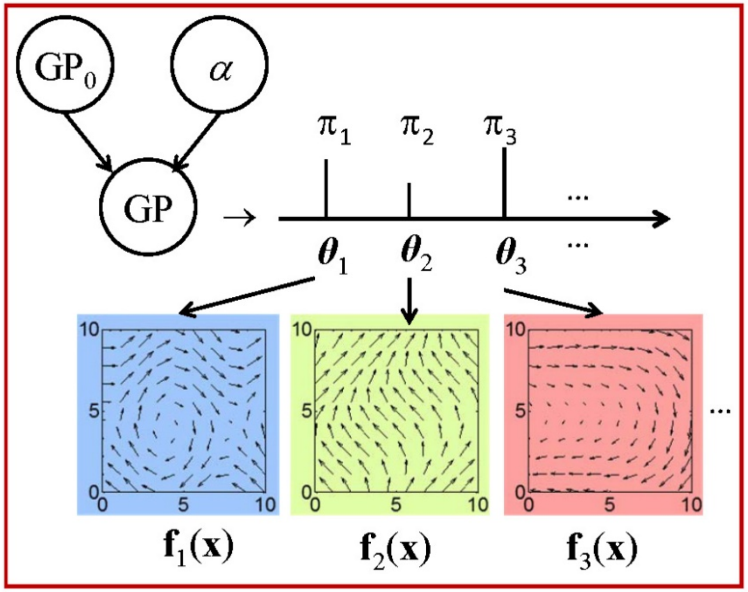

[1] Joseph, Joshua Mason, Finale Doshi-Velez, and Nicholas Roy. "A Bayesian Nonparametric Approach to Modeling Mobility Patterns." *AAAI*. 2010.

Decentralized communication control:

- Determine decentralized control policies
- Determine when to communicate and, thus, share target measurements
- **Competing objectives:** navigation, communication, sensing

Where:

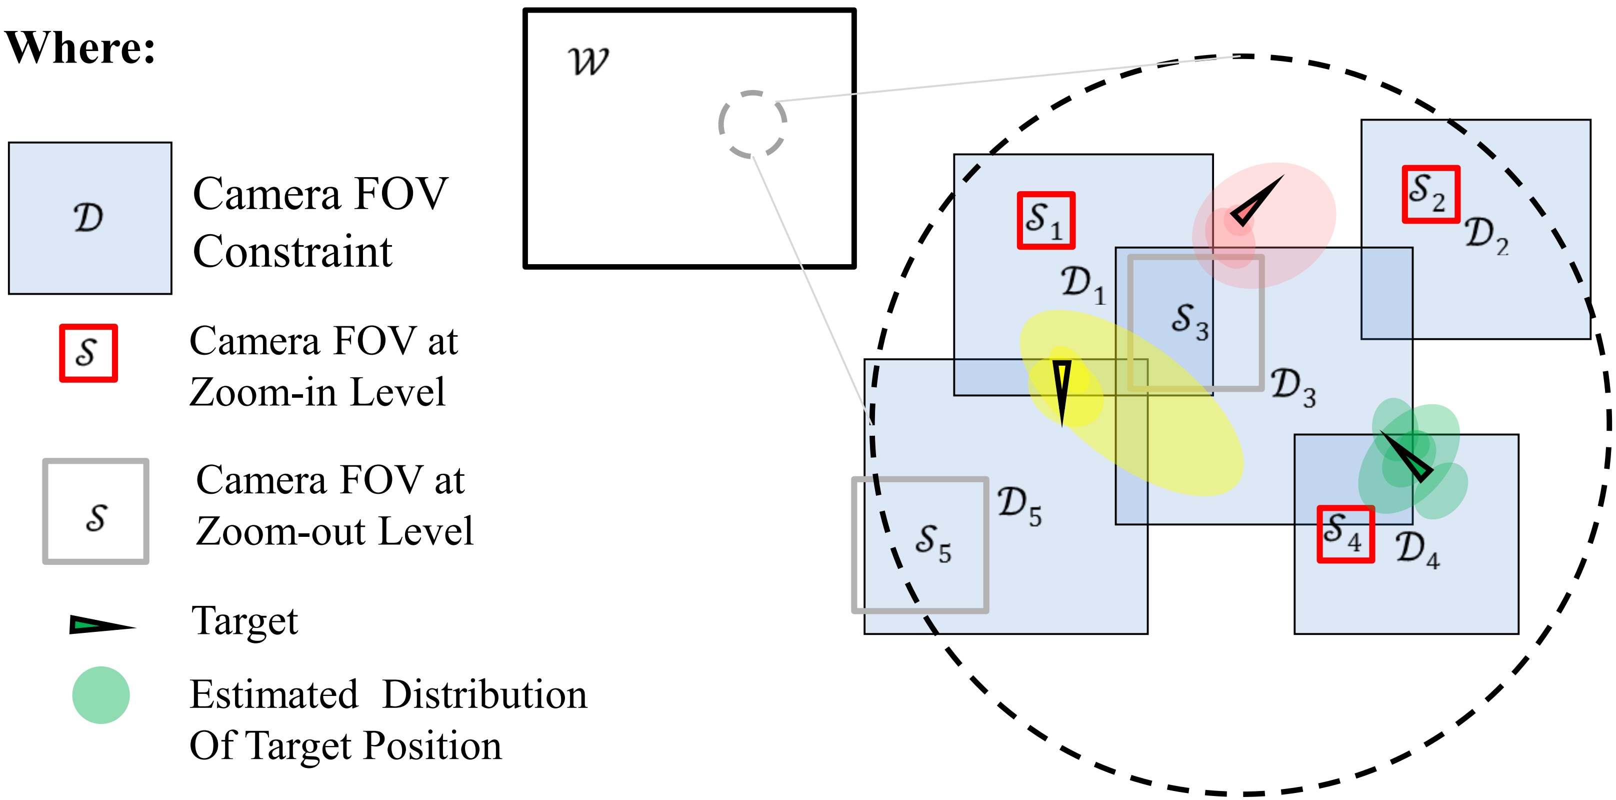

Analysis of BNP Information Value Functions

DP-GP Information Value

- Goal: difference between posterior and prior DP-GP model

$$\varphi[\mathbf{F}(\mathbf{X});\mathbf{m}(k+1) \mid \mathcal{M}(k), \mathcal{E}(k), \mathbf{u}(k)]$$
$$= D\{p[\mathbf{F}(\mathbf{X}) \mid \mathcal{M}(k+1), \mathcal{E}(k), \mathbf{u}(k)] \,||\, p[\mathbf{F}(\mathbf{X}) \mid \mathcal{M}(k), \mathcal{E}(k), \mathbf{u}(k)]\}$$

where

D: Kullback-Leibler (or other) divergence

$$D(P||Q) = \int_{-\infty}^{\infty} \ln\left(\frac{p(x)}{q(x)}\right) p(x)\,dx$$

$\mathbf{F}(\mathbf{X}) = [\mathbf{f}_1(\mathbf{X}) \;\ldots\; \mathbf{f}_M(\mathbf{X})]$, $\mathbf{X}$: points of interests

$\mathcal{M}(k) = \{\mathbf{m}(\ell) | \mathbf{m}(\ell) \neq \emptyset, k' < \ell \leq k\}$,
measurement history not used in updating DP-GP model

$\mathcal{E}(k) = \{\mathbf{m}(\ell) | \mathbf{m}(\ell) \neq \emptyset, 0 \leq \ell \leq k'\}$,
measurement history used in updating DP-GP model

Properties of KL Divergence

Assuming the noise in position measurements is small, the expected KL divergence,

$$\hat{\varphi}[\mathbf{u}(k)] = \mathrm{E}_{g_j}[\mathrm{E}_{\mathbf{m}(k+1)}[\varphi(\mathbf{u}(k)) \mid g_j = i]]$$

can be written as,

$$\hat{\varphi}[\mathbf{u}(k)] = \sum_{j=1}^{N}\sum_{i=1}^{M} w_{ij} \int h_i(\mathbf{x}_j) p(\mathbf{x}_j) d\mathbf{x}_j \approx \sum_{j=1}^{N}\sum_{i=1}^{M} \frac{w_{ij}}{S} \sum_{s=1}^{S} h_i(\mathbf{x}_{js})$$

such that the computation is reduced from an 8$^{\text{th}}$-order integral to a double integral, where $h(\bullet)$ is a known analytic function with the following properties:

- $h_i(\bullet)$ is bounded from below by 0
- $h_i(\bullet)$ is bounded from above by $4k\left[\dfrac{k+\sigma_v^2}{\sigma_v^2(1+k+\sigma_v^2)}\right]^2 \mathrm{tr}(\Sigma_1^{-1}) < \infty$

Theorem: The above approximation is an unbiased estimator of the DP-GP expected KL divergence, and the variance of the approximation error decreases linearly at a rate of $1/S$.

w_{ij} = i-j^{th} association probability, N = no. targets, M = no. velocity fields in $\mathcal{F}$, S = no. particles, $\mathbf{x}_{js}$ = sth particle of the j^{th} velocity field, $p(\bullet)$ = posterior probability on the previous slide, k = present time step, σ_v = std of velocity measurement noise, Σ_1 = covariance matrix at grid points.

Performance Analysis:
Example of Velocity Fields

First VF Second VF Third VF Fourth VF

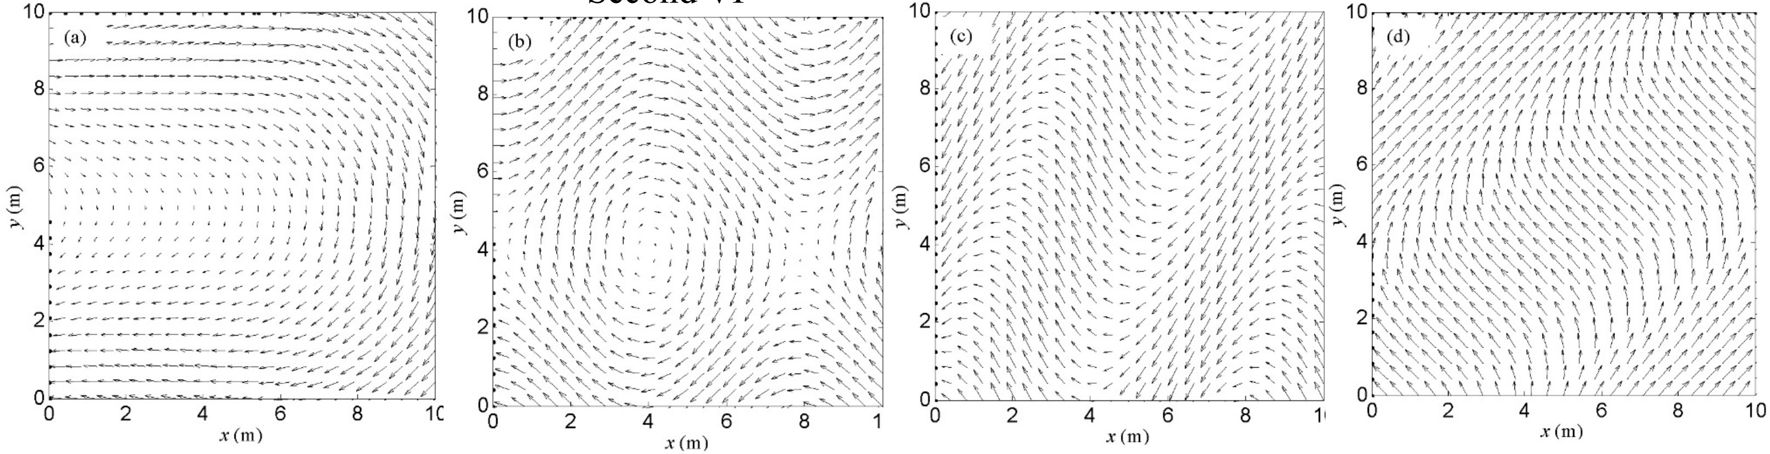

Target Trajectories:

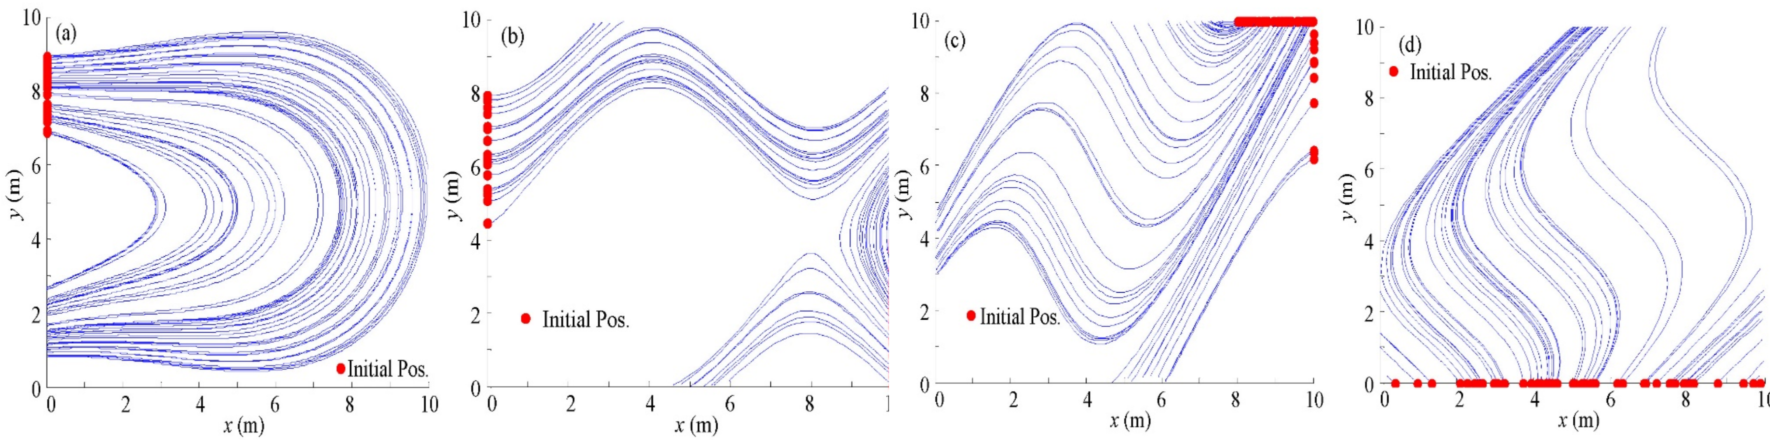

First VF Second VF Third VF Fourth VF

KL Divergence vs. Entropy

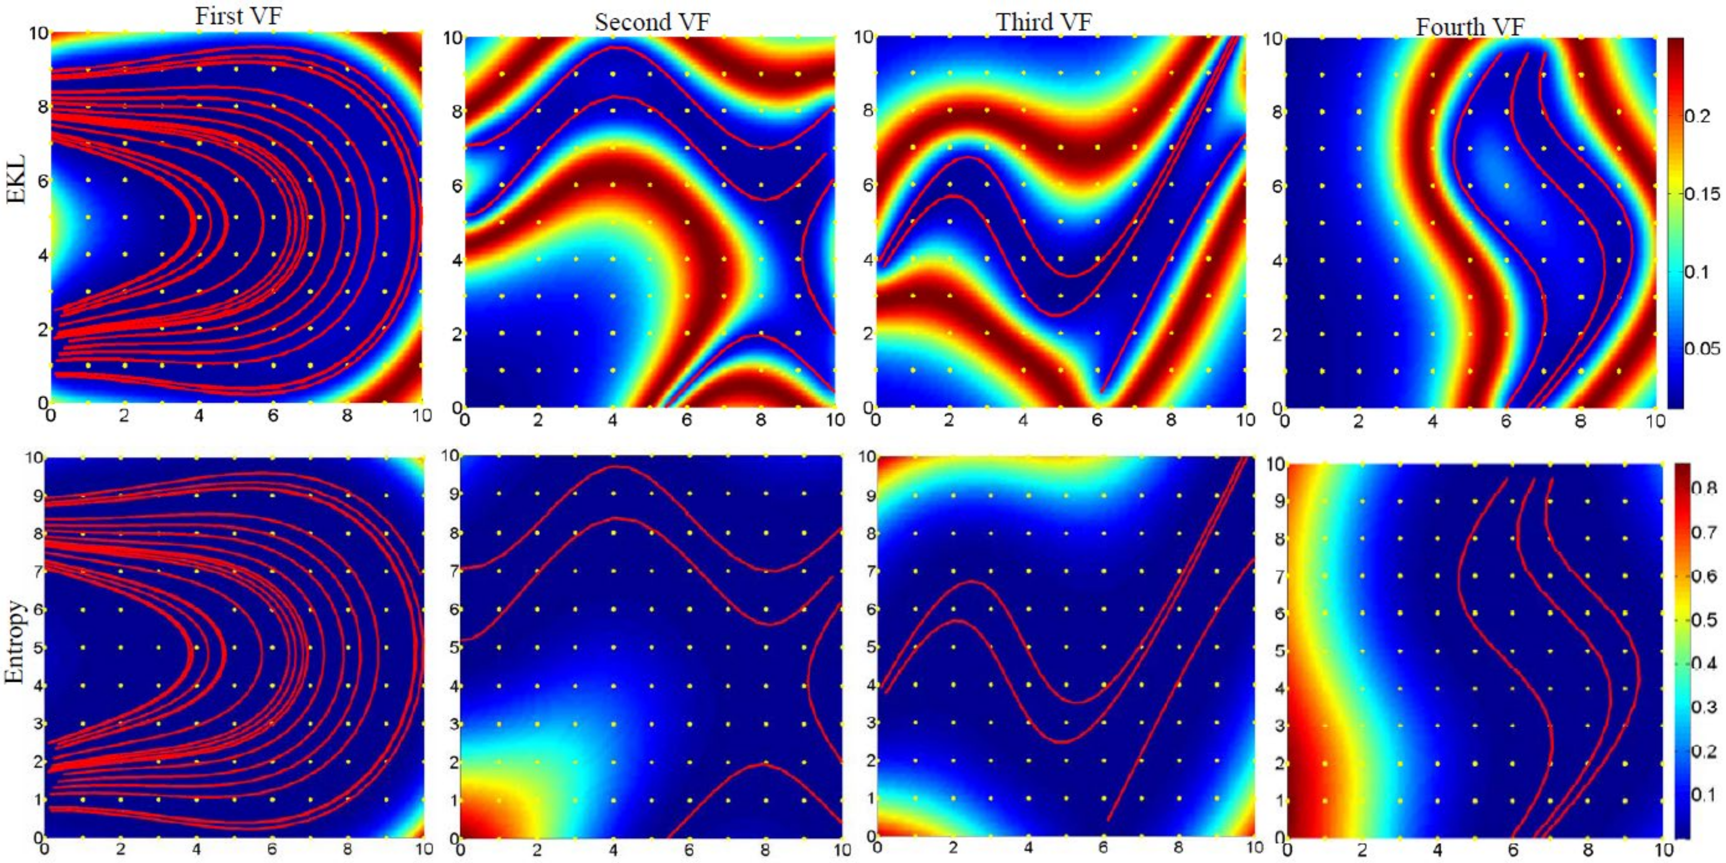

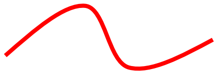

: Observed target trajectories up to *k*

• Grid Points for KL computation

Performance Comparison

More **Prior Information** Less

RMS of Velocity Error (m/s):

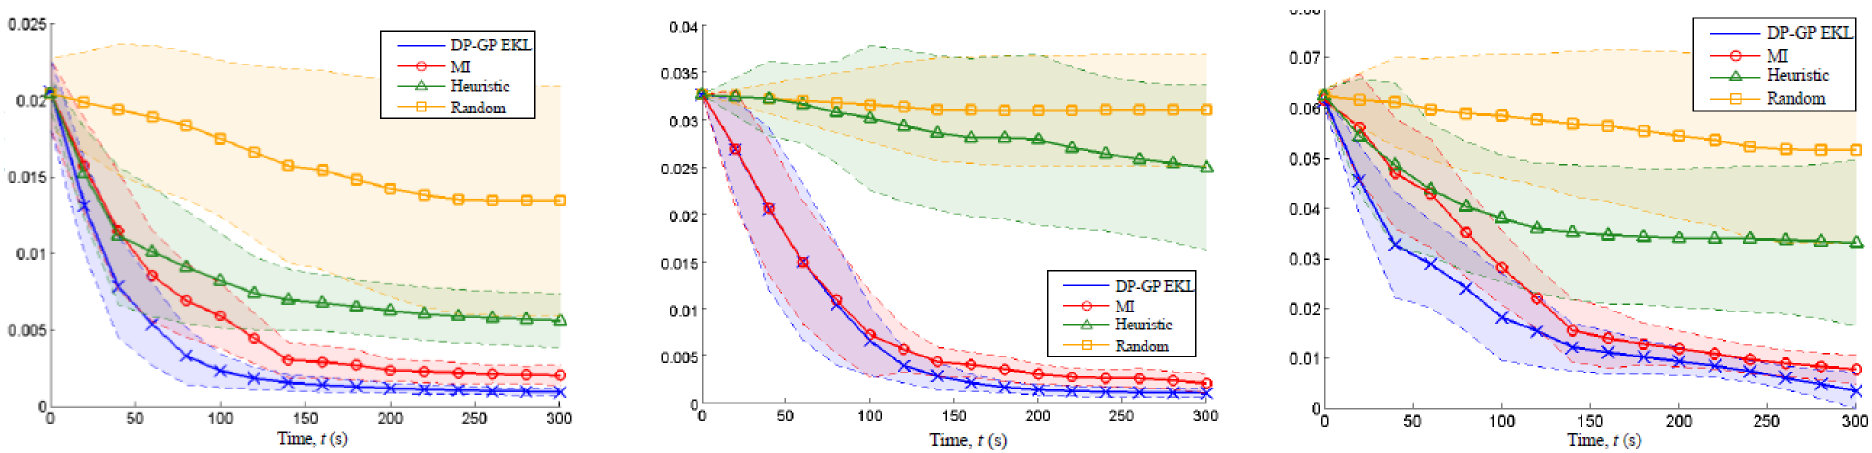

× **DP-GP EKL:** camera PTZ levels are controlled by optimizing the expected KL divergence.

○ **MI:** camera PTZ levels are controlled by optimizing the mutual information (MI).

△ **Heuristic Search:** camera PTZ levels are controlled such that the FoV centroid tracks the estimated position of the nearest target.

□ **Random Search:** camera PTZ levels are controlled by a randomized search algorithm.

Decentralized Visibility-based Control

Decentralized Simulation Results

- **Decentralized Active Camera Control** for mobile intruder BNP modeling and tracking, **without FOV constraints**.

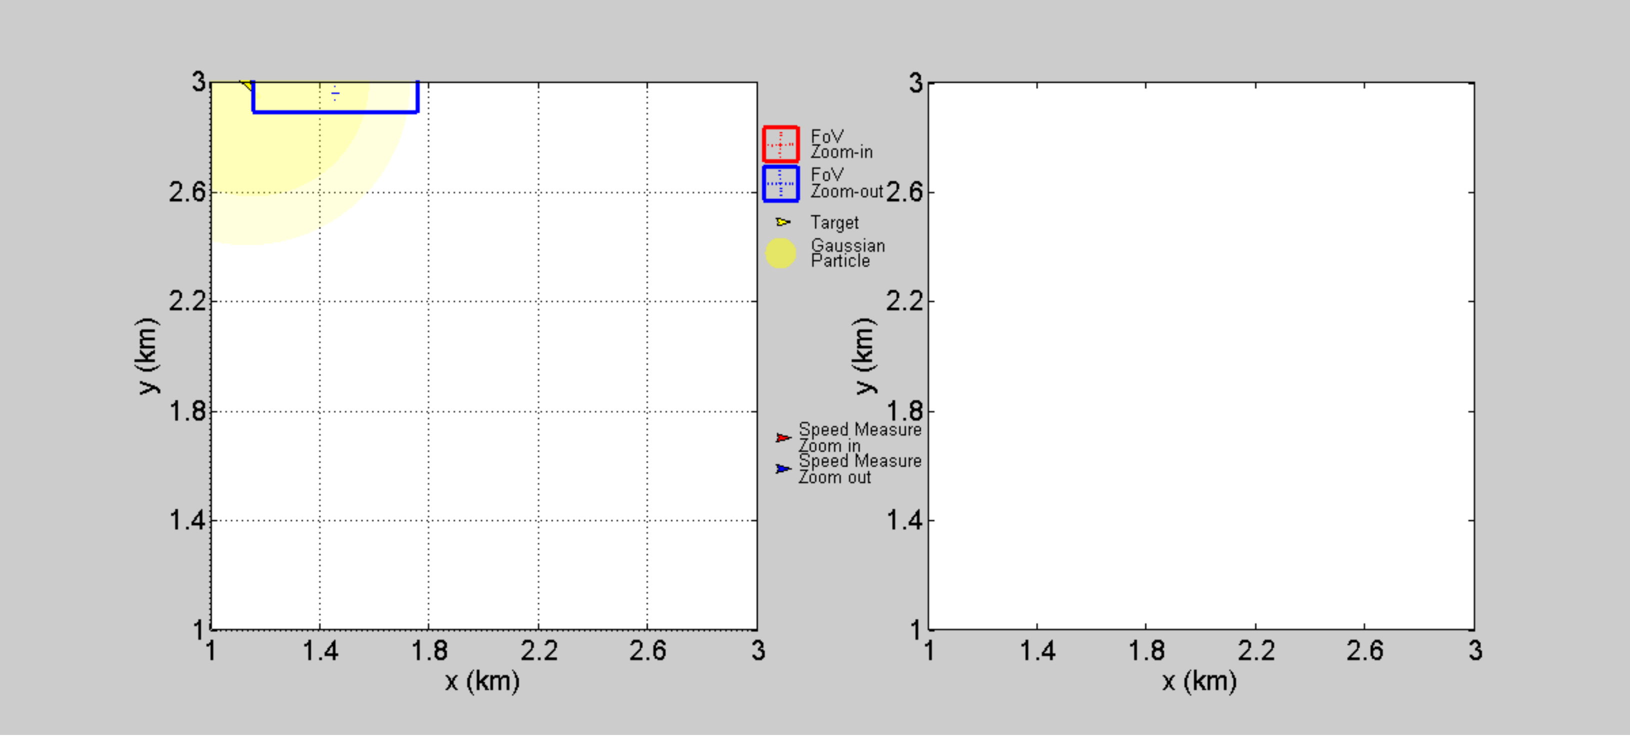

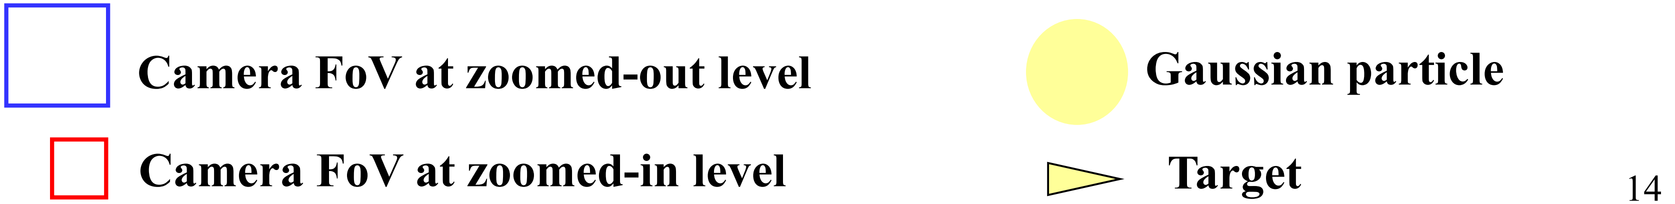

Decentralized Simulation Results

- **Decentralized Active Camera Control** for mobile intruder BNP modeling and tracking, **with FoV constraints**.

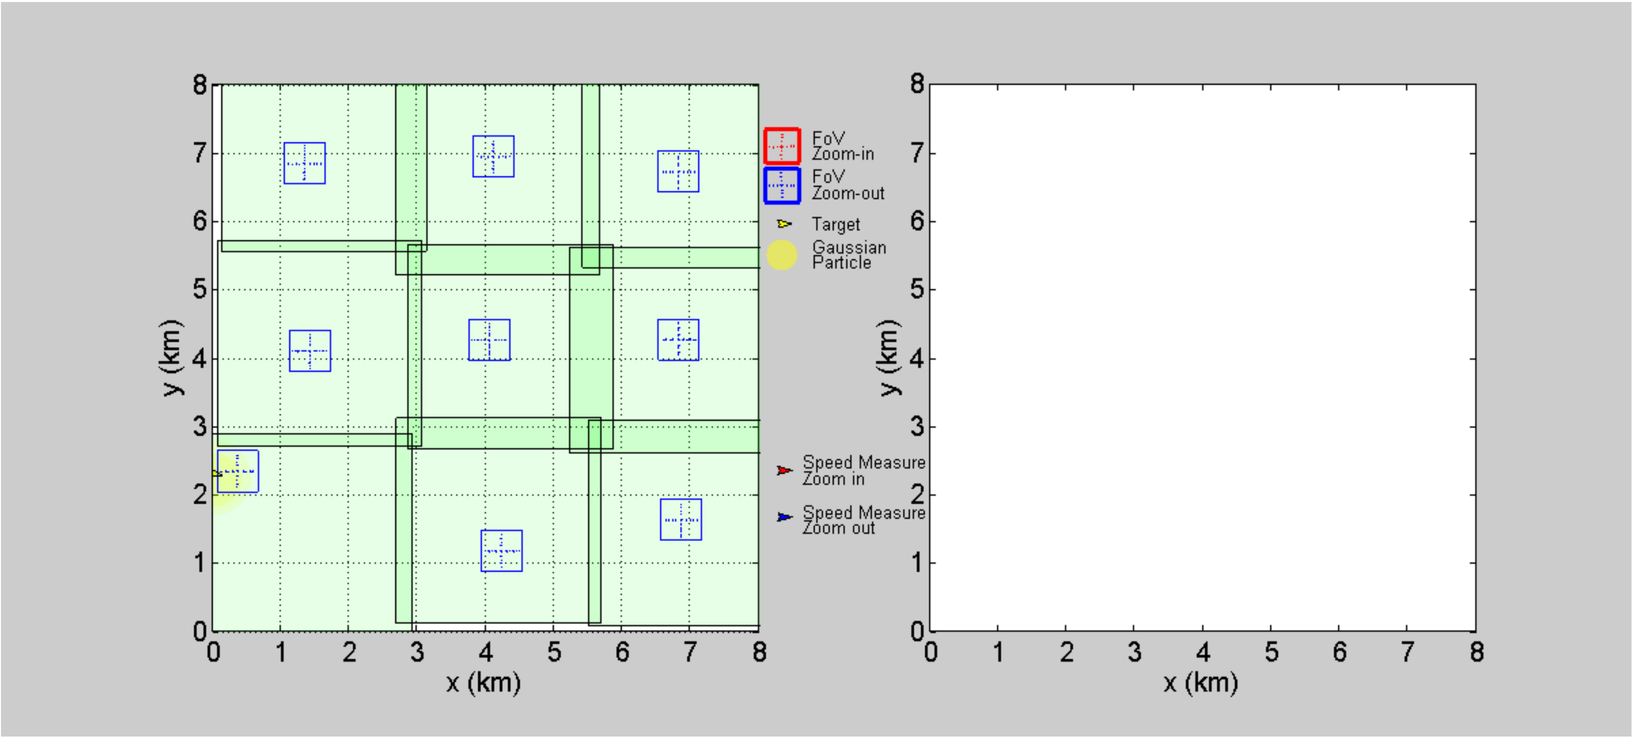

Camera FoV at zoomed-out level

Camera FoV at zoomed-in level

FoV constraint

Target

Decentralized Control with Intermittent Comm

Communication Control

Motivation: Develop planning and control algorithms for collaborative networks with intermittent communications

- Existing decentralized optimization methods assume constant communications (network is a connected graph) or detailed prior information (perfect models)
- Consider networks in which some or all nodes (agents) may be disconnected some of the time, and there is no or little prior information (high uncertainty)
- Agents aim to construct BNP model from data
- Disconnected agents can determine when their own local information is insufficient, and it is time to reestablish communications

Model a spatial phenomenon:

$g(\mathbf{x}), \mathbf{x} \in \mathcal{A}$

- Max temperature over a 2D ROI

  $\mathcal{A} \subset \mathbb{R}^2$

- Time invariant
- Observable at a set of target locations:

  $\mathcal{T} = \{\mathbf{t}_i | i = 1, \cdots, r\}, \mathbf{t}_i \in \mathcal{A}$

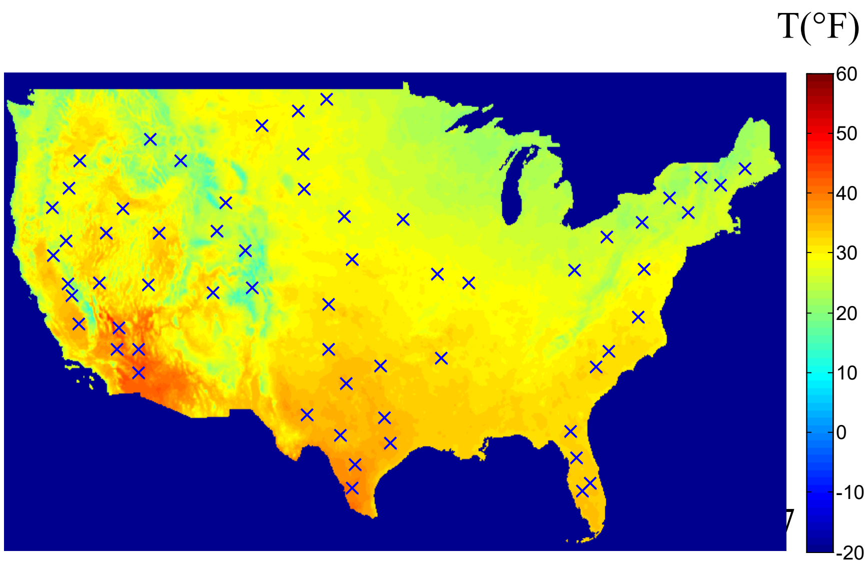

GP Model and Information Value

Estimation of spatial phenomenon:

- $g(\mathbf{x}) \xleftarrow{\text{estimation}} f(\mathbf{x}), \mathbf{x} \in \mathcal{A}$
- Measurements: $f(\mathbf{x}) \sim$ Gaussian process

Gaussian process:

$$f(\mathbf{x}) \sim \mathrm{GP}(\mu(\mathbf{x}),\ \psi(\mathbf{x}_1,\mathbf{x}_2));\ \mu(\mathbf{x}) = \mathrm{E}[f(\mathbf{x})]$$

$$\psi(\mathbf{x}_1,\ \mathbf{x}_2) = \mathrm{E}[(\,f(\mathbf{x}_1) - \mu(\mathbf{x}_1)\,)(\,f(\mathbf{x}_2) - \mu(\mathbf{x}_2)\,)]$$

Planning objective: at time k choose locations and measurements $\{\mathbf{y}_k, z_k\}$ to maximize,

$$D(p(\mathbf{f}(\mathbf{X}_T)|\mathbf{Y}_k, \mathbf{Z}_k)||p(\mathbf{f}(\mathbf{X}_T)|\mathbf{Y}_{k-1}, \mathbf{Z}_{k-1}))$$

where,

$\mathbf{X}_T = [\mathbf{x}_1 \,|\, \cdots \,|\, \mathbf{x}_r]$, $\mathbf{x}_i \in \mathcal{T}$, $i = 1, \ldots, r$; $\mathbf{Y}_k = [\mathbf{y}_1 \,|\, \cdots \,|\, \mathbf{y}_k]$; $\mathbf{Z}_k = [z_1 \,|\, \cdots \,|\, z_k]$

Since $\mathbf{z}_k$ is unknown, optimize **expected discrimination gain (EDG):**

$$\hat{\varphi}_D(\mathbf{f}(\mathbf{X}_T); z_k|\mathbf{y}_k, \mathbf{Y}_{k-1}, \mathbf{Z}_{k-1}) = \int D(p(\mathbf{f}(\mathbf{X}_T)|\mathbf{Y}_{k-1}, \mathbf{Z}_{k-1}, \mathbf{y}_k, z_k)||p(\mathbf{f}(\mathbf{X}_T)|\mathbf{Y}_{k-1}, \mathbf{Z}_{k-1})) \times p(z_k|\mathbf{Y}_{k-1}, \mathbf{Z}_{k-1}, \mathbf{x}_k)dz_k.$$

Disconnected Network

Modeling of a spatial phenomenon, $g(\mathbf{x})$, by four robots with disjoint workspaces:

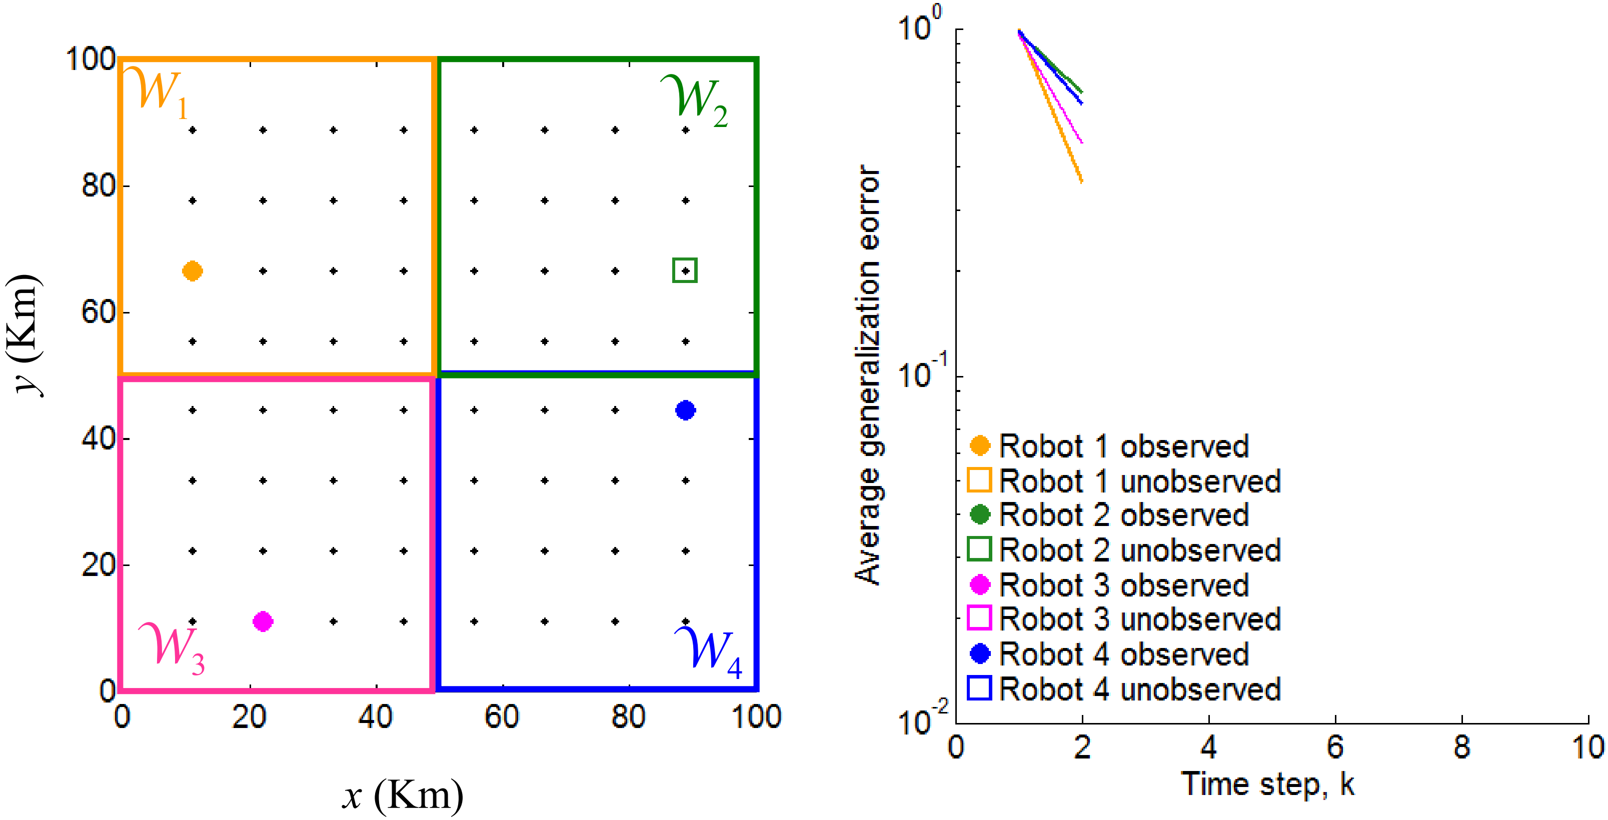

Nominal Network Performance

- Let $\mathbf{\Sigma}$ denote the covariance matrix and $\Psi(\mathbf{x}, \mathbf{y})$ denote the cross-covariance matrix, then the GP **average generalization error** (AGE),

$$\varepsilon(k) = \mathrm{E}_{\mathbf{x}}\left\{ \Sigma(\mathbf{x}) - \Psi(\mathbf{x}, \mathbf{Y}_k)\left[\Sigma(\mathbf{Y}_k) + \delta^2 \mathbf{I}\right]^{-1} \Psi^T(\mathbf{x}, \mathbf{Y}_k) \right\}$$

represent a measure of GP performance.

- From the latest GP, the posterior covariance, and the network nominal AGE can be estimated from an assumed probability distribution for future measurement locations, and an assumed probability of detection p_t

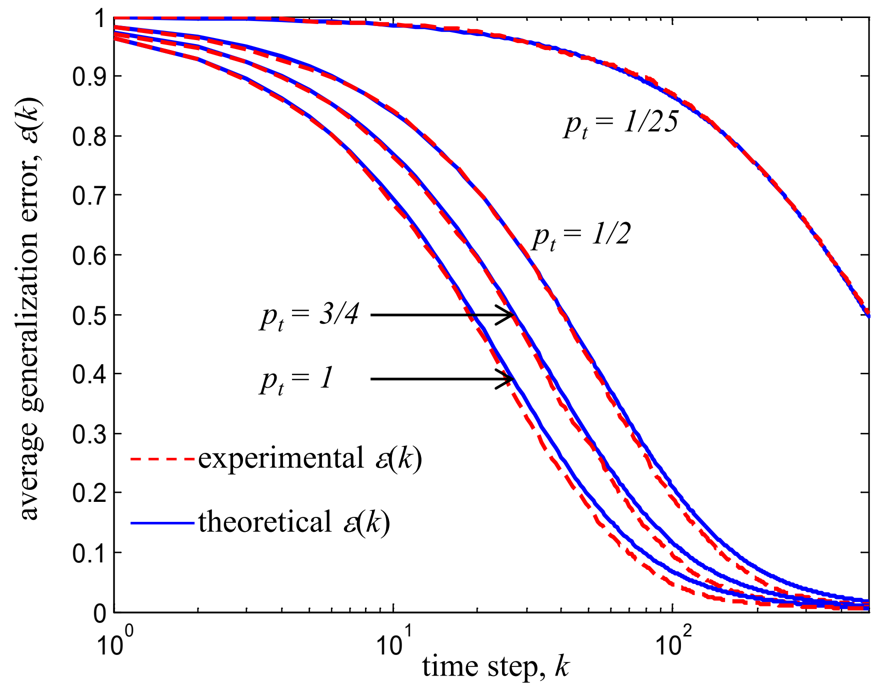

GP Communication Control

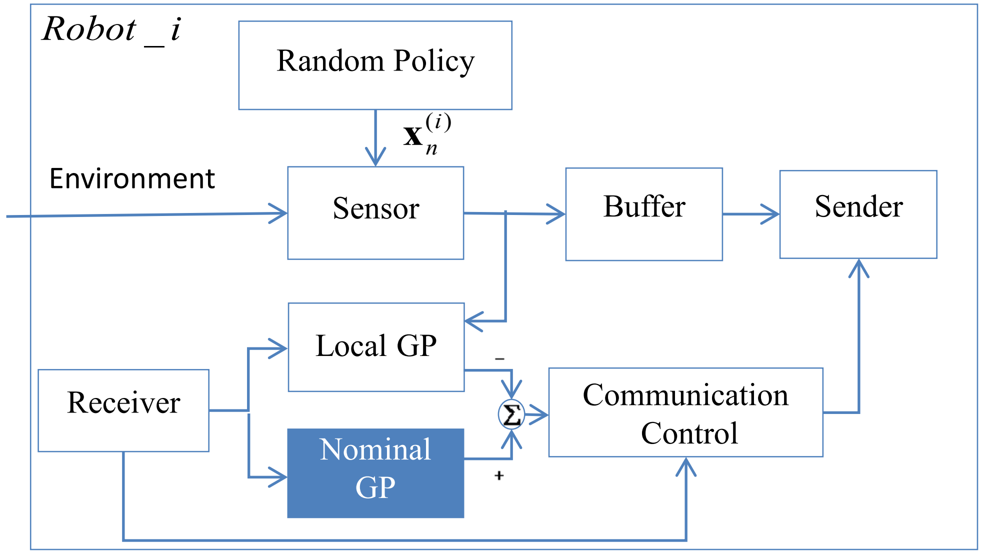

- Approximate nominal average generalization error (AGE) from latest GP

GP Communication Control

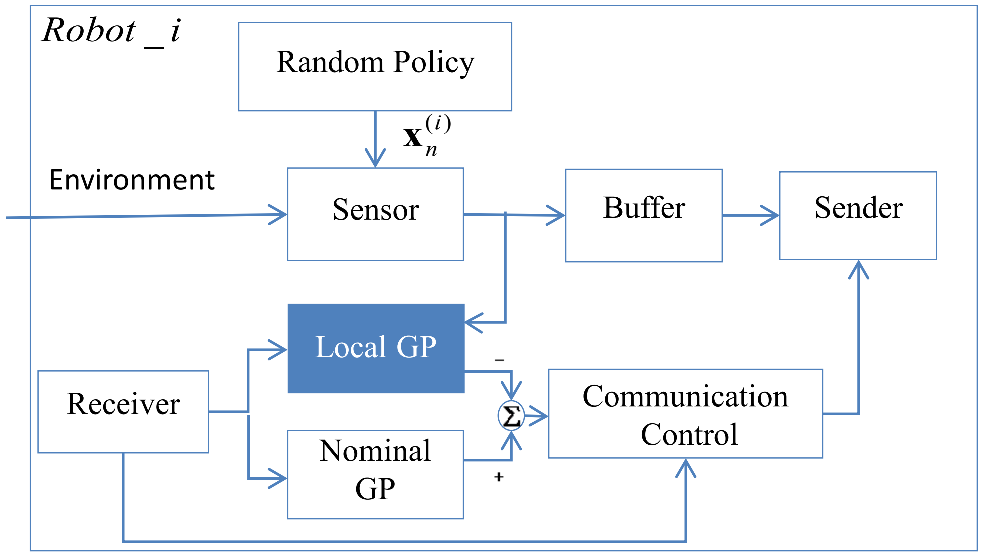

- Approximate nominal average generalization error (AGE) from latest GP
- Local GP (GPL) computation
  - GPL is updated using local measurements (obtained by robot *i*)
  - Actual AGE is calculated from the local covariance function.

GP Communication Control

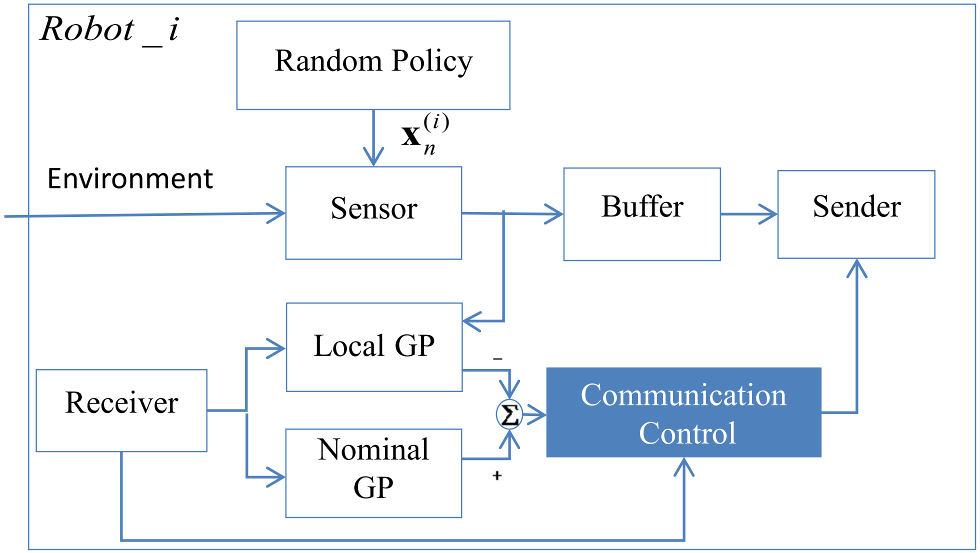

- Approximate nominal average generalization error (AGE) from latest GP
- Local GP (GPL) computation
- Communication time: at the n^{th} time step, the i^{th} robot communicates if

$$\max_i \left| \sum_{k=n_0}^{n} \varepsilon_i(k) - \varepsilon_{\text{nominal}}(k) \right| > \gamma$$

where, γ is predefined performance threshold.

GP Communication Control

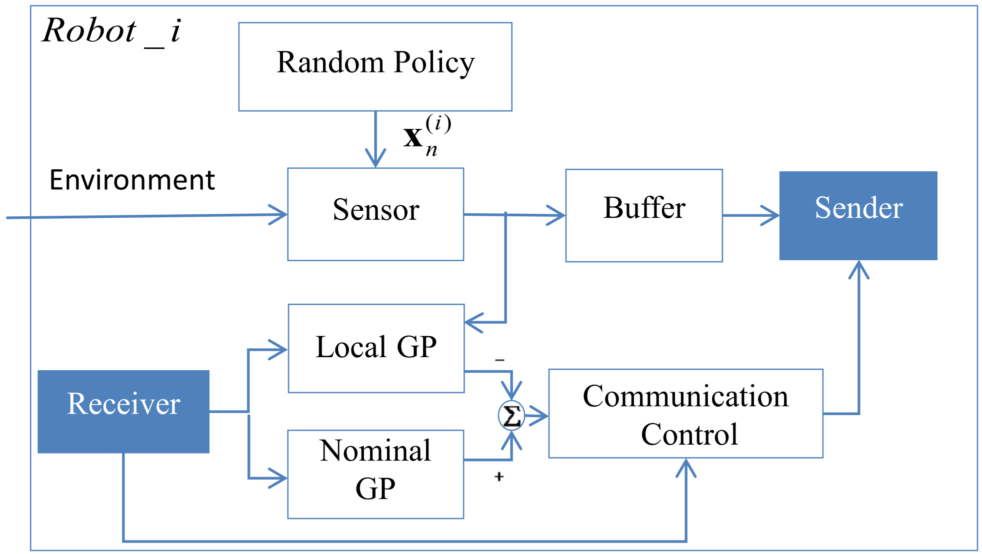

- Approximate nominal average generalization error (AGE) from latest GP
- Local GP (GPL) computation
- Communication time
- Information sharing: all new measurements are communicated and used to update the robot GP

Two-robot GP Comm Control

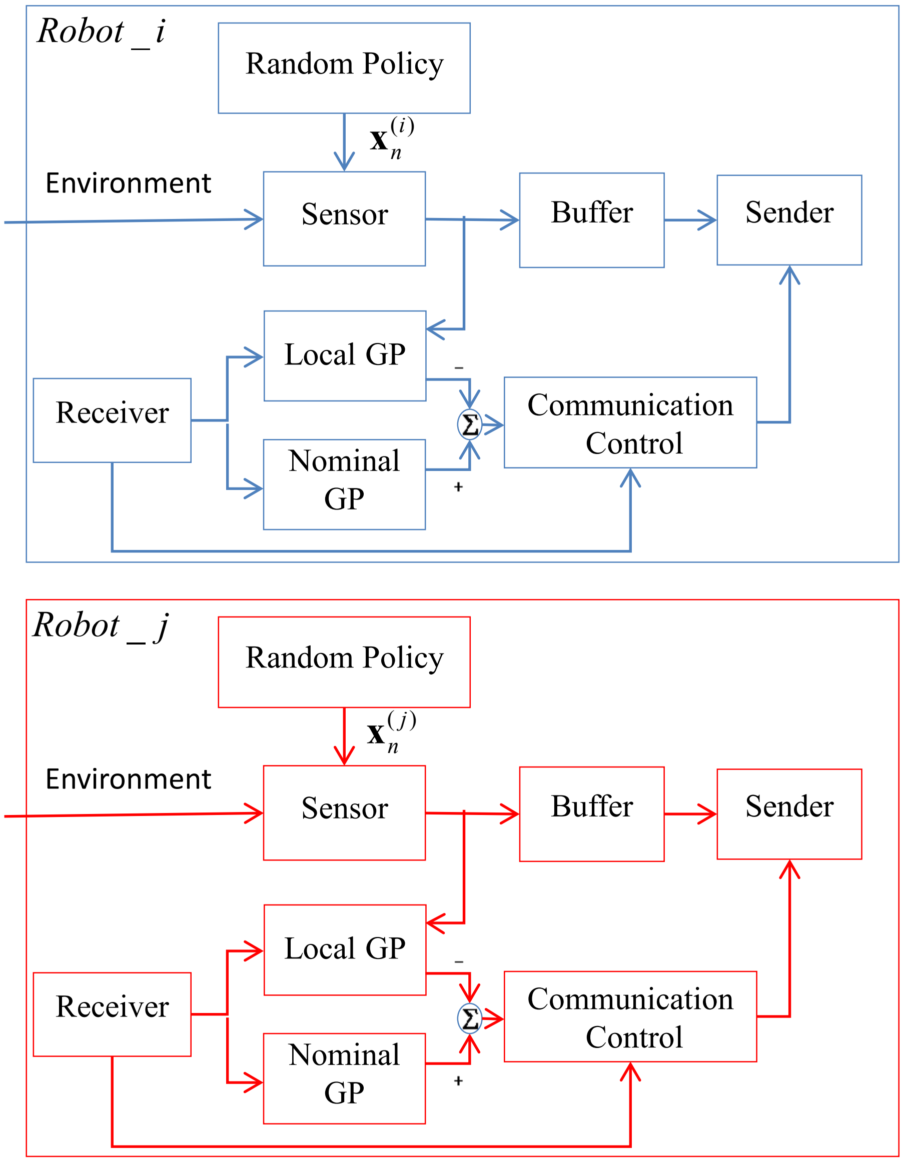

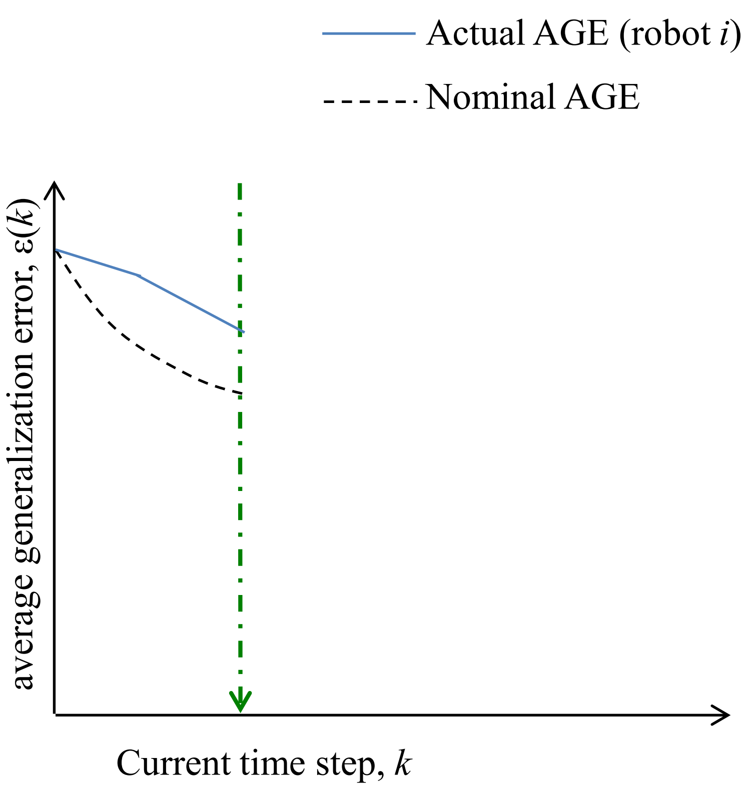

Two-robot GP Comm Control

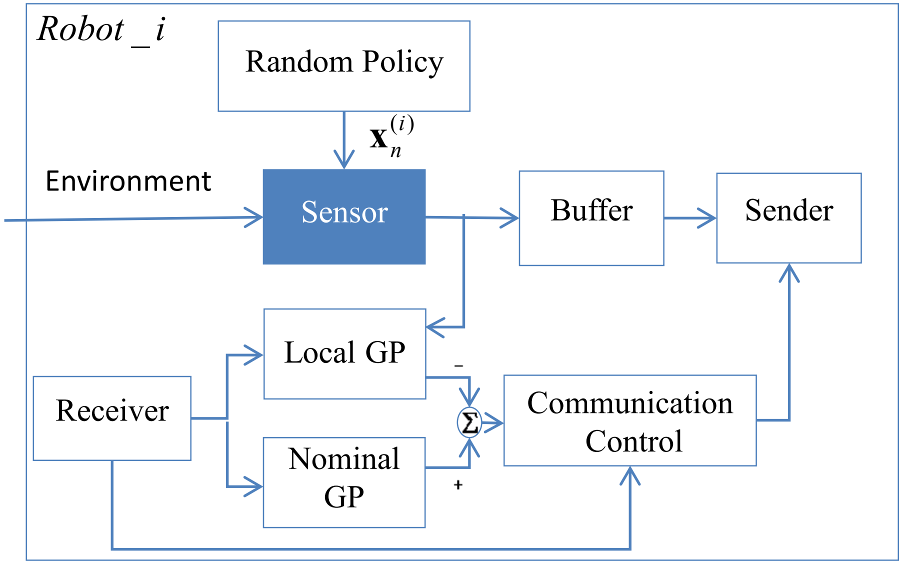

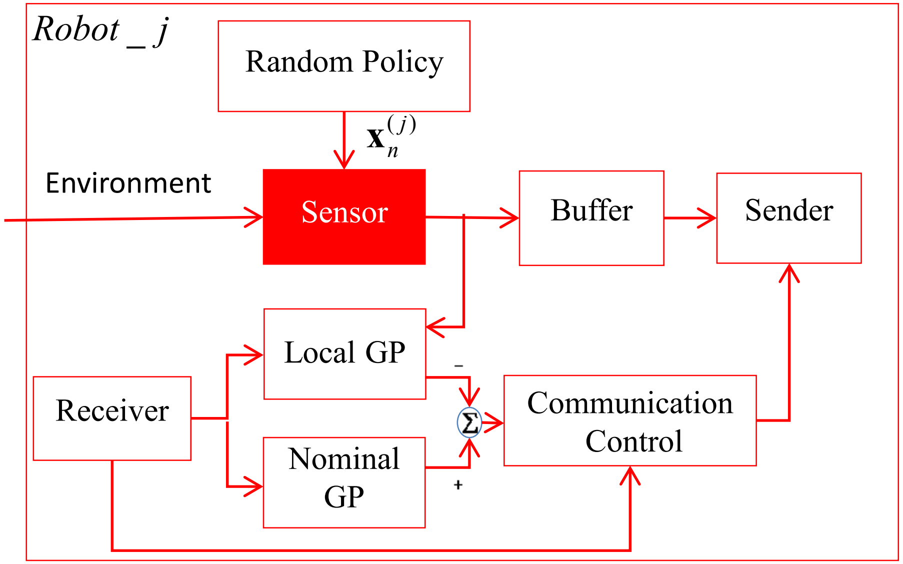

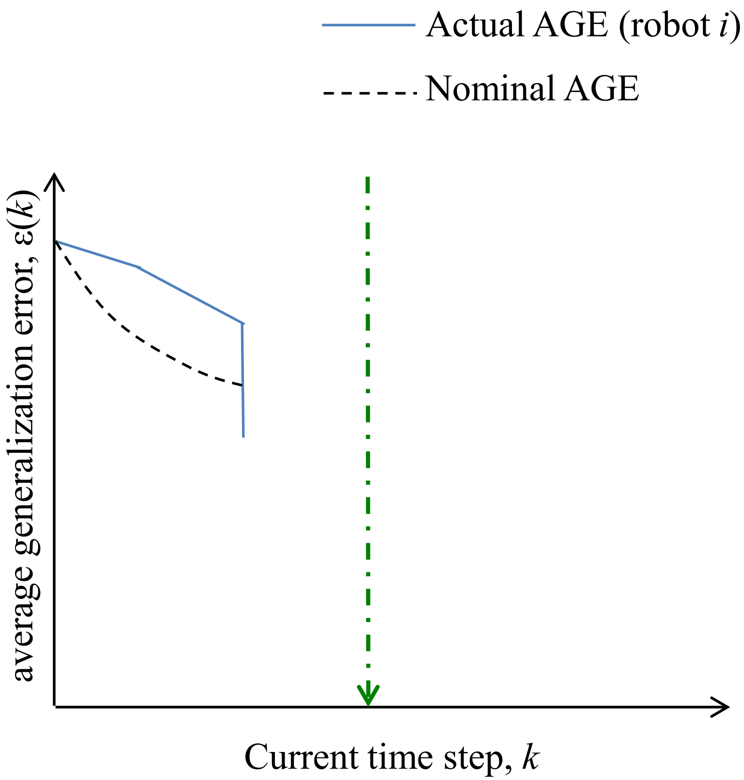

Two-robot GP Comm Control

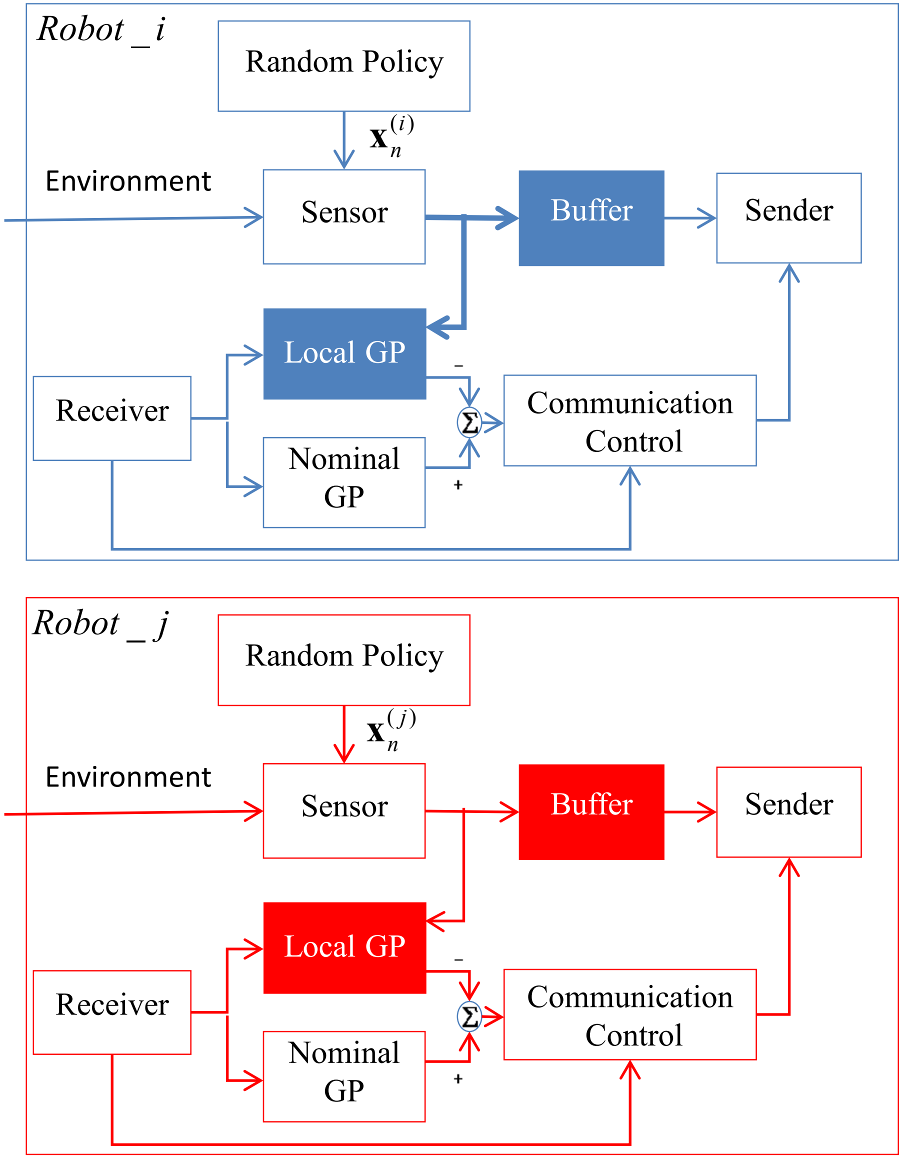

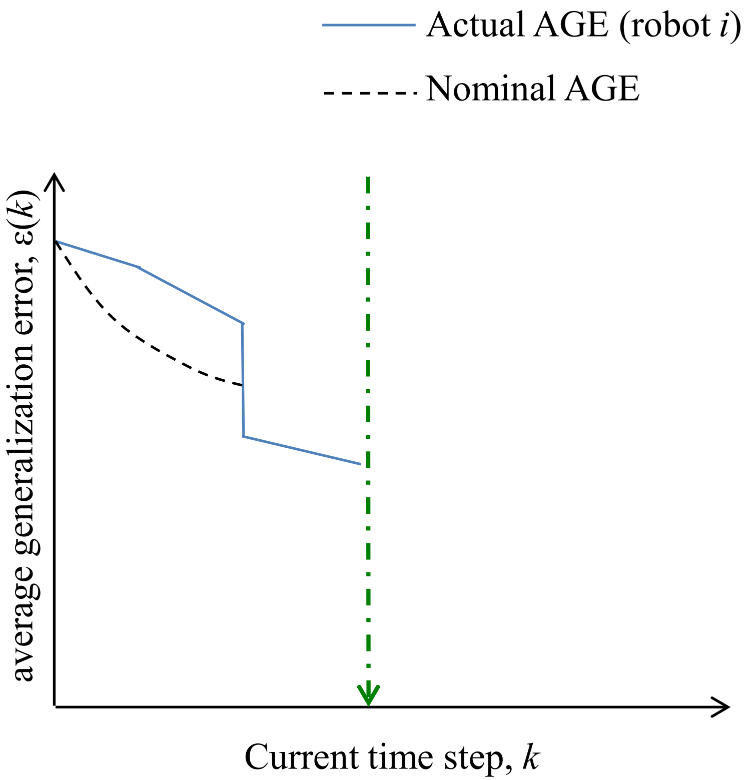

Two-robot GP Comm Control

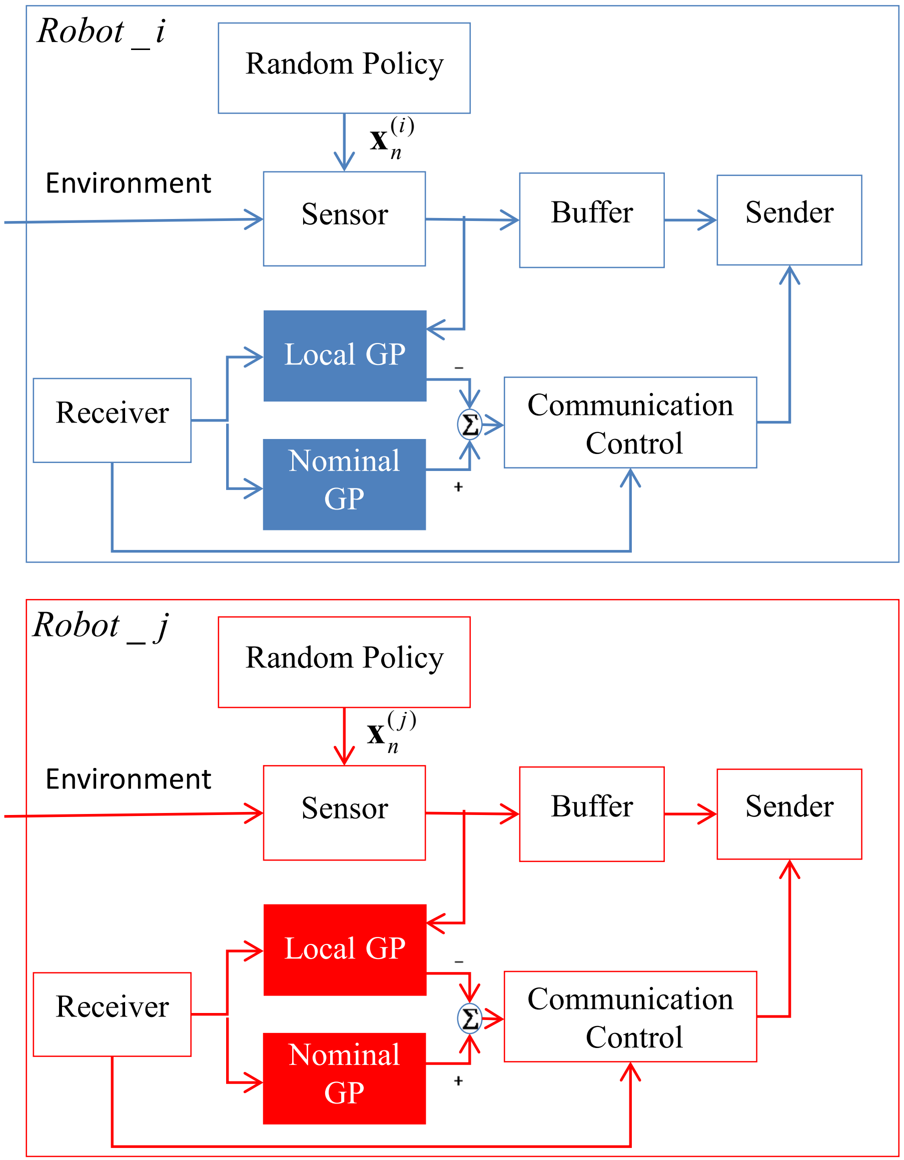

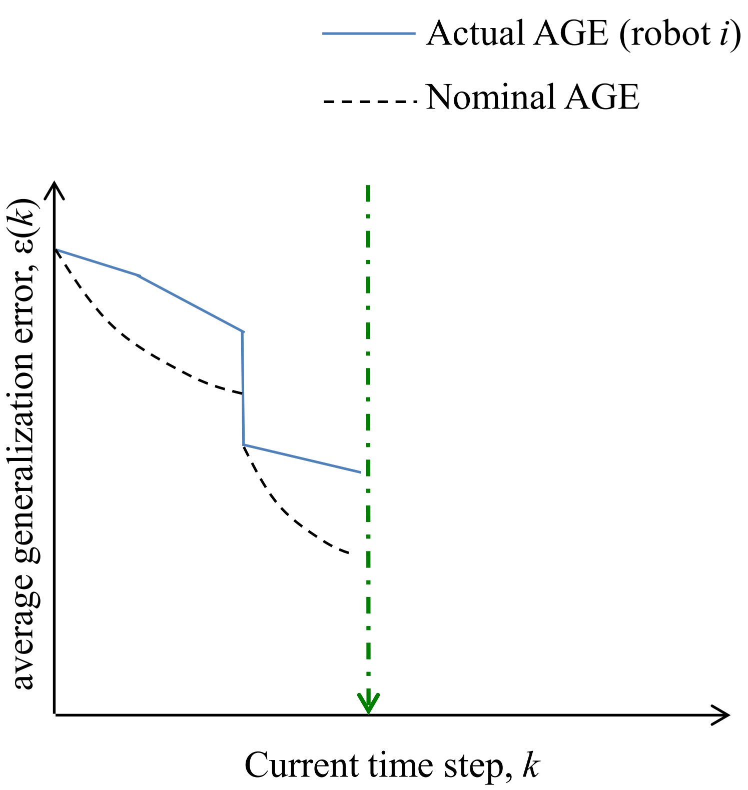

Simulation Results

Modeling of a spatial phenomenon, $g(\mathbf{x})$, by four robots with disjoint workspaces:

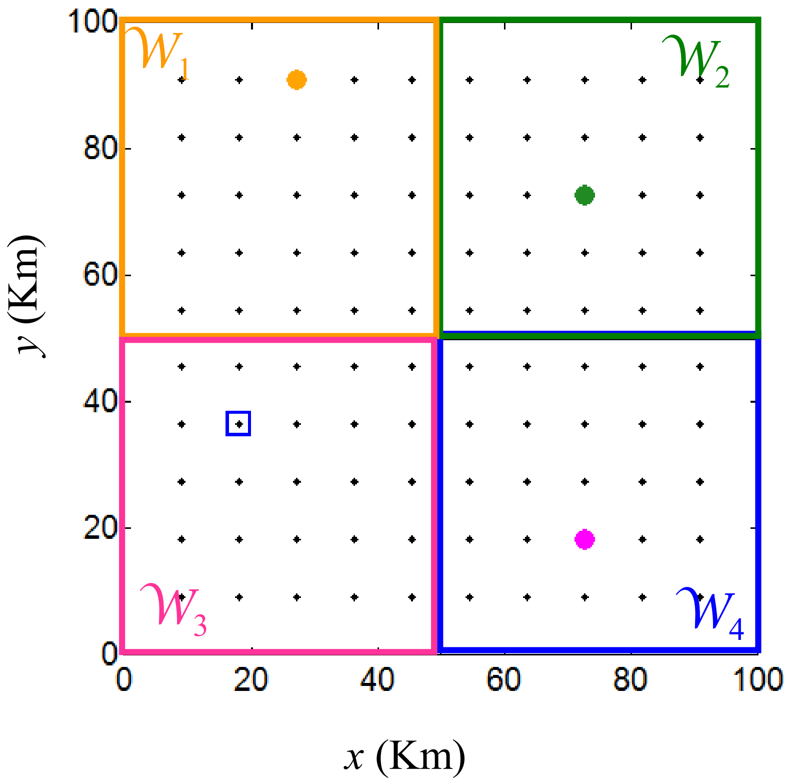

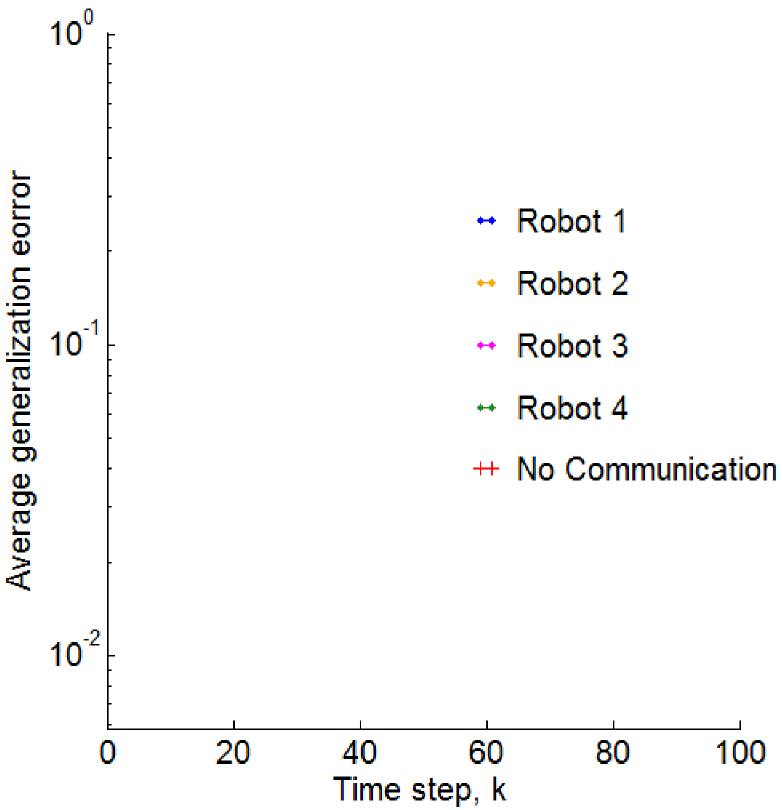

Summary and Future Work

Technical Accomplishments – Year 3:

- Derived properties of **BNP information value functions**
- Conducted performance analysis for **BNP planning/control** algorithms
- Developed and demonstrated **BNP decentralized planning/control** algorithms for connected sensor networks, modeling multiple moving targets collaboratively
- Developed approach for **BNP communication control** in decentralized disconnected sensor networks, modeling a spatial process collaboratively

Future Work:

- Complete BNP communication planning/control theory
- Develop convergence and performance guarantees
- Develop BNP simultaneous motion and communication planning algorithms
- Demonstrate motion and communication planning algorithms on collaborative mobile sensors modeling multiple moving targets, in comm-denied environments

Acknowledgements:
This research was funded by the ONR MURI Grant N00014-11-1-0688.

?

Questions?