

LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

# Value Function Approximation for Multiscale Dynamical Systems

Pingping Zhu, Julian Morelli

Advisor : Silvia Ferrari

Laboratory for Intelligent Systems and Controls

Conference for Decision and Control

Las Vegas

December 14, 2016

# Adaptive Planning for Intelligent Collaborative Systems

- **Overview:**

- Multi-agent collection of autonomous robotic vehicles,  $O(10^2)$ 
  - Large temporal and spatial scales
- Uncertainty in environmental conditions

- **Applications:**

- Coastal & Environmental monitoring
- Surveillance
- Search and Rescue
- Recognizance



United States Coast Guard Rescue Team [1]



REMUS Autonomous submarine [2]

# Background

- **Multi-Agent Planning Methods:**

- Prioritized Path Planning
- Multi-Agent Potential Field Methods
- **Distributed Optimal Control**

- **Limitations:**

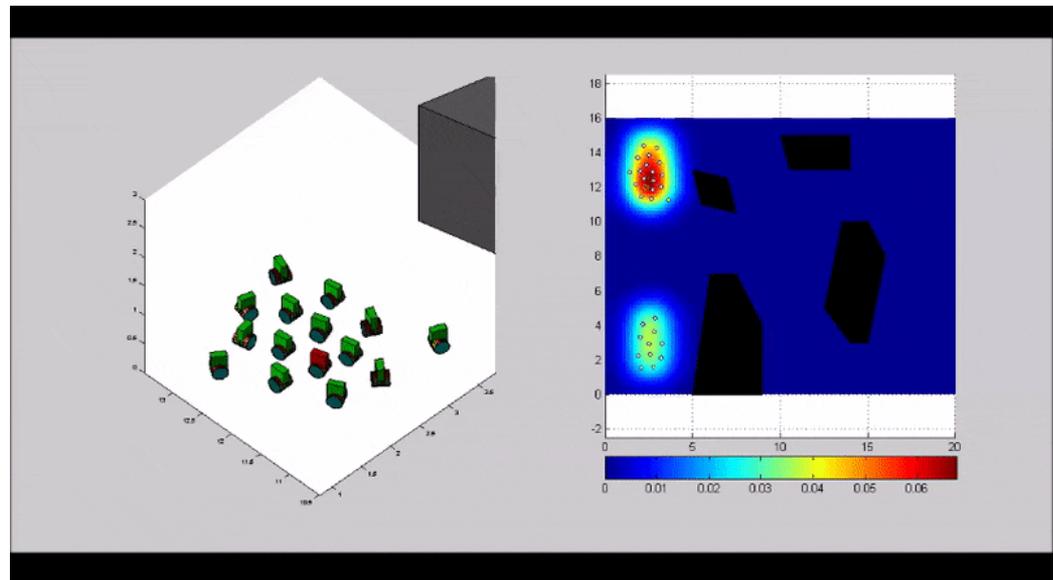
- High computational complexity
- Sub- or non-optimal
- Non-adaptive/Offline methods

- **Distributed Optimal Control:**

- Currently best method for a problem this size
- Numerical solution via Generalized Reduced Gradient Method

- **Technical Challenges:**

- Hundreds of ODEs
  - Nonlinear dynamics
- Functional approximation
- Rapidly changing environment





# Distributed Optimal Control

- Consider  $N$  autonomous robotic agents with dynamics described by the following system of ordinary differential equations:

$$\dot{\mathbf{x}}_i(t) = \mathbf{f}[\mathbf{x}_i(t), \mathbf{u}_i(t), t] + \mathbf{G}\mathbf{w}_i(t)$$

- Macroscopic state, *restriction operator*:  $\wp(\mathbf{x}, t) \Phi : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{P}$
- Macroscopic dynamics represented by the *advection-diffusion* equation:
  - $\mathbf{G}$  is a constant matrix,  $\mathbf{w}_i(t)$  is an additive Gaussian noise
  - Where  $\mathbf{v} \equiv \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t)$  and  $\nu = \nabla(\mathbf{G}\mathbf{G}^T)$

$$\begin{aligned} \frac{\partial \wp}{\partial t} &= -\nabla \cdot [\wp(\mathbf{x}, t)\mathbf{v}] + \frac{1}{2} \nabla \cdot [(\mathbf{G}\mathbf{G}^T)\nabla \wp(\mathbf{x}, t)] \\ &= -\nabla \cdot [\wp(\mathbf{x}, t)\mathbf{f}(\mathbf{x}, \mathbf{u}, t)] + \nu \nabla^2 \wp(\mathbf{x}, t) \end{aligned}$$

## Distributed Optimal Control

- Cost function in terms of restriction operator

$$J = \phi[\wp(\cdot, T_f)] + \int_{T_0}^{T_f} \int_{\mathcal{X}} \mathcal{L}[\wp(\mathbf{x}, t), \mathbf{u}(\mathbf{x}, t), t] d\mathbf{x}_i dt$$

- PDF restriction operator must also satisfy:  $P(\mathbf{x} \in B, t) = \int_B \wp(\mathbf{x}, t) d\mathbf{x}$

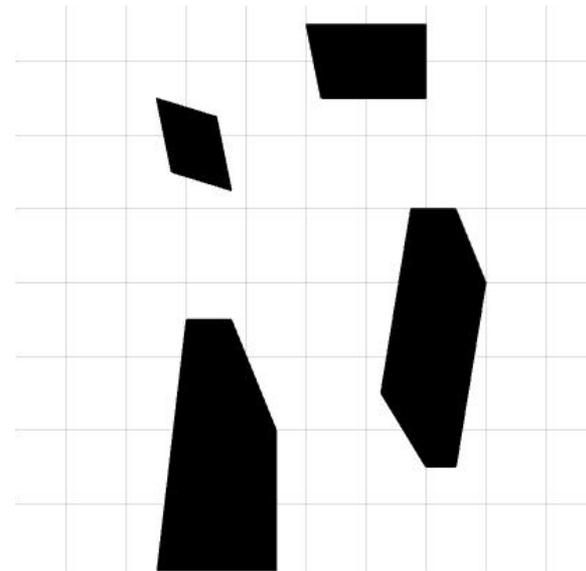
and these other conditions:



$$\begin{aligned} \text{I.C.:} \quad & \wp(\mathbf{x}, T_0) = \wp_0(\mathbf{x}); \\ \text{B.C.:} \quad & \wp(\mathbf{x} \in \partial\mathcal{X}, t) = 0, \quad \forall t \in [T_0, T_f]; \\ \text{A.C.:} \quad & \int_{\mathcal{X}} \wp(\mathbf{x}, t) d\mathbf{x}_i = 1; \\ & \wp(\mathbf{x} \notin \partial\mathcal{X}, t) = 0, \quad \forall t \in [T_0, T_f]. \end{aligned}$$

# Distributed Optimal Control

- Currently the most effective way of solving optimal control of large multiscale systems
- **Limitations:**
  - Offline Method / Non-adaptive
  - Must know *a priori*:
    - Microscopic agent dynamics
    - Macroscopic evolution equation
    - Definition of the restriction operator
    - Environmental conditions ← 



Obstacle Configuration

# Research Problem and Motivation

- **Workspace:**  $\mathcal{X} \subset \mathbb{R}^2$
- **Agent Dynamics:**  $\dot{\mathbf{x}}_i(t) = \mathbf{f}[\mathbf{x}_i(t), \mathbf{u}_i(t), t] + \mathbf{G}\mathbf{w}_i(t)$
- **Restriction Operator:**  $\Phi(\mathbf{x}, t) : \mathbb{R} \rightarrow \mathcal{P}$ 

Gaussian Mixed Model  
 Expectation Maximization algorithm
- **Cost Function:**  $J = \phi[\varphi(\cdot, T_f)] + \int_{T_0}^{T_f} \int_{\mathcal{X}} \mathcal{L}[\varphi(\mathbf{x}, t), \mathbf{u}(\mathbf{x}, t), t] d\mathbf{x}_i dt$
- **Optimal Value Function:**  $\mathcal{V}^* = \min_{\mathbf{u}_k} \left\{ \int_{\mathcal{X}} \mathcal{L}(\varphi_k^*, \mathbf{u}_k) d\mathbf{x}_i + \mathcal{V}^*[\varphi_{k+1}^*, \mathcal{C}^*(\cdot)] \right\}$



# Adaptive Critics

- Dynamic Programming optimizes a *Value function*:

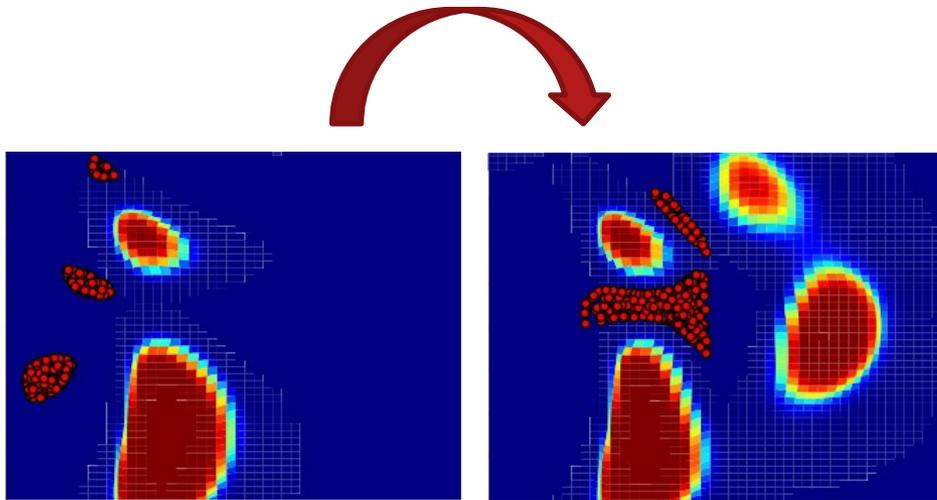
$$\mathcal{V}^* = \min_{\mathbf{u}_k} \left\{ \int_{\mathcal{X}} \mathcal{L}(\varrho_k^*, \mathbf{u}_k) d\mathbf{x}_i + \mathcal{V}^*[\varrho_{k+1}^*, \mathcal{C}^*(\cdot)] \right\}$$

- Hamilton-Jacobi-Bellman Equation

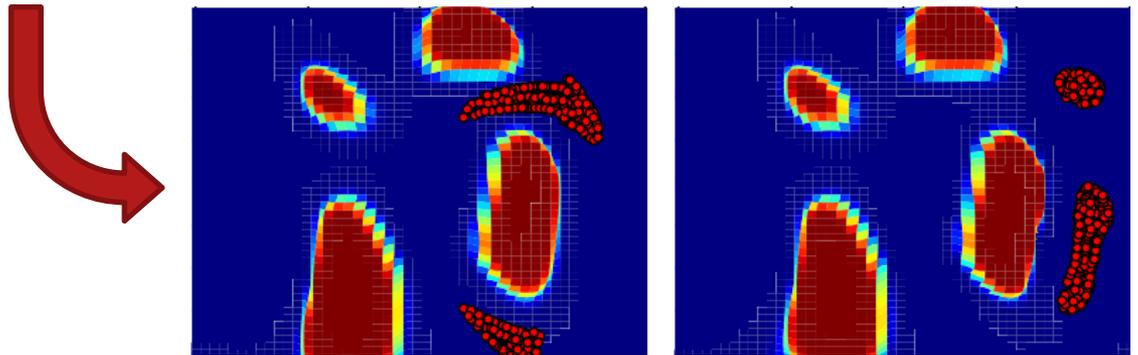
- Associated with a cost function (at time k):

$$\mathcal{L}(\varrho_k) = \int_{\mathbf{x} \in \mathcal{X}} \underbrace{\varrho_k(\mathbf{x}) \log \frac{\varrho_k(\mathbf{x})}{g(\mathbf{x})}}_{\text{Kullback-Leibler Divergence}} + \underbrace{\varrho_k(\mathbf{x}) U_{\text{rep}}}_{\text{Obstacle Repulsion}} + \underbrace{\varrho_k(\mathbf{x}) e^{\{w_u[u_1(x,k)^2 + u_2(x,k)^2]/2\}}}_{\text{Energy Consumption}} d\mathbf{x}$$

# Adaptive Critics



$$\mathcal{L}(\rho_k) = \int_{\mathbf{x} \in \mathcal{X}} \underbrace{\rho_k(\mathbf{x}) \log \frac{\rho_k(\mathbf{x})}{g(\mathbf{x})}}_{\text{KL-Divergence}} + \underbrace{\rho_k(\mathbf{x}) U_{\text{rep}}}_{\text{Obstacle Repulsion}} + \underbrace{\rho_k(\mathbf{x}) e^{\{w_u [u_1(x,k)^2 + u_2(x,k)^2] / 2\}}}_{\text{Energy Consumption}} d\mathbf{x}$$





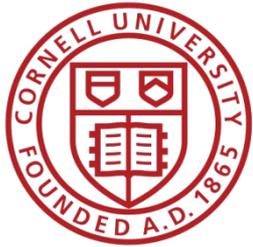
## Adaptive Critics

- Given a value function corresponding to a control law, and improved control law can be obtained as follows:

$$\mathcal{C}_{\ell+1}(\varrho_k) = \arg \min_{\mathbf{u}_k} \left\{ \int_{\mathcal{X}} \mathcal{L}(\varrho_k, \mathbf{u}_k) d\mathbf{x} + \mathcal{V}[\varrho_{k+1}, \mathcal{C}_{\ell}(\cdot)] \right\}$$

- Given a control law, the value function can be updated according to the following rule:

$$\mathcal{V}_{\ell+1}[\varrho_k, \mathcal{C}(\cdot)] = \int_{\mathcal{X}} \mathcal{L}(\varrho_k, \mathbf{u}_k) d\mathbf{x} + \mathcal{V}_{\ell}[\varrho_{k+1}, \mathcal{C}(\cdot)]$$



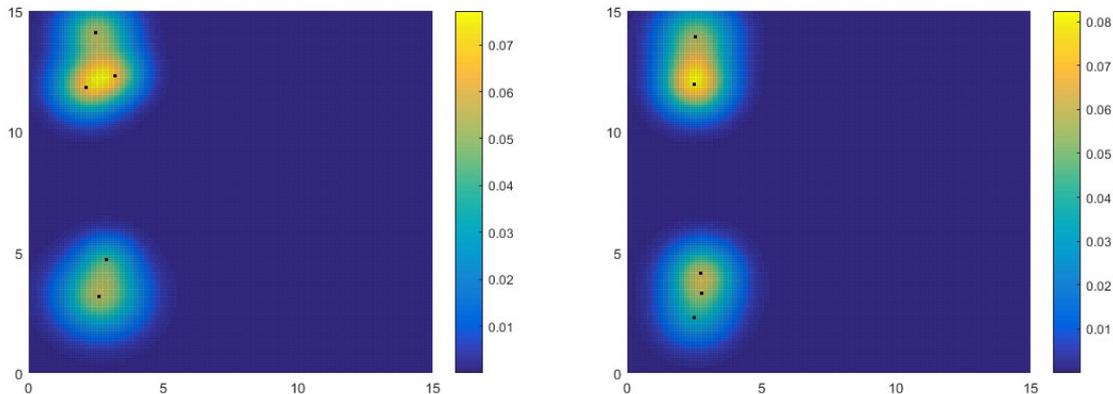
---

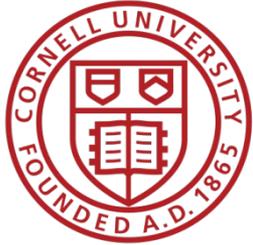
LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

# Gaussian Mixed Model

# Gaussian Mixed Model

- GMM used to approximate macroscopic state PDF
- Form: 
$$\tilde{\varphi}_k(\mathbf{x}) = \sum_{\tau=1}^n \omega_{\tau} \mathcal{N}(\mathbf{x} | \mu_{k,\tau}, \Sigma_{k,\tau})$$
- Parameters cannot be ordered!
  - Parameter tube is discontinuous with respect to the time step  $(\omega_{k,\tau}, \mu_{k,\tau}, \Sigma_{k,\tau})$





---

LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

# Projection onto Fixed Basis

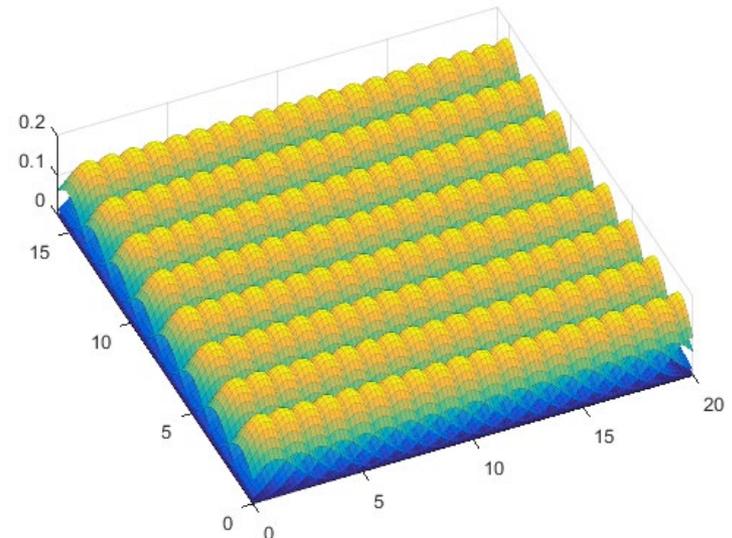
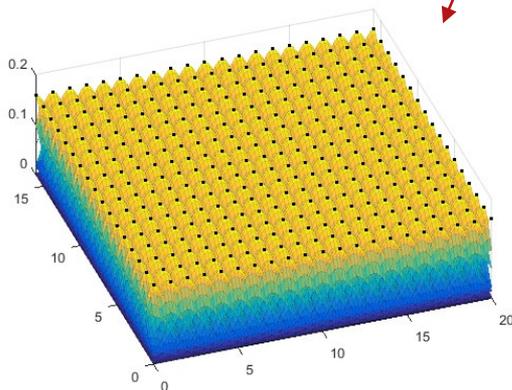
# Projection onto Fixed Basis

- Project onto fixed basis of Gaussians using inner product:

$$\mathbf{p}_k = I_j(\tilde{\varphi}_k) = \int_{\mathbf{x} \in \mathcal{X}} \tilde{\varphi}_k(\mathbf{x}) b_j(\mathbf{x}) d\mathbf{x}, \quad \text{for } j = 1, \dots, m$$

- $\mathbf{p}_k$  is a coefficient vector for the fixed basis set  $\{b_j\}_{j=1\dots m}$

- Gaussian basis functions





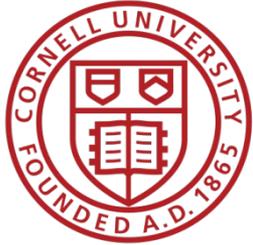
## Projection onto Fixed Basis

- Inner product can be computed in closed form:
  - *Gaussian Identity Property:*

$$\sum_{\tau=1}^n \int_{\mathbf{x} \in \mathcal{X}} \omega_{\tau} \mathcal{N}(\mathbf{x} | \mu_{k,\tau}, \Sigma_{k,\tau}) \mathcal{N}(\mathbf{x} | \mu_j, \Sigma_j) d\mathbf{x} = \sum_{\tau=1}^n \omega_{\tau} \mathcal{N}(\mu_{k,\tau} | \mu_j, \Sigma_{k,\tau} + \Sigma_j)$$

- Using the normalized basis previously defined:

$$\begin{aligned} \mathbf{p}_k &= I_j(\tilde{\phi}_k) = \left\langle \tilde{\phi}_k, \frac{b_j}{\|b_j\|} \right\rangle_{\mathcal{P}} \\ &= \frac{\sum_{\tau=1}^n \omega_{\tau} \mathcal{N}(\mu_{k,\tau} | \mu_j, \Sigma_{k,\tau} + \Sigma_j)}{\mathcal{N}(\mu_j | \mu_j, 2\Sigma_j)}, \end{aligned}$$



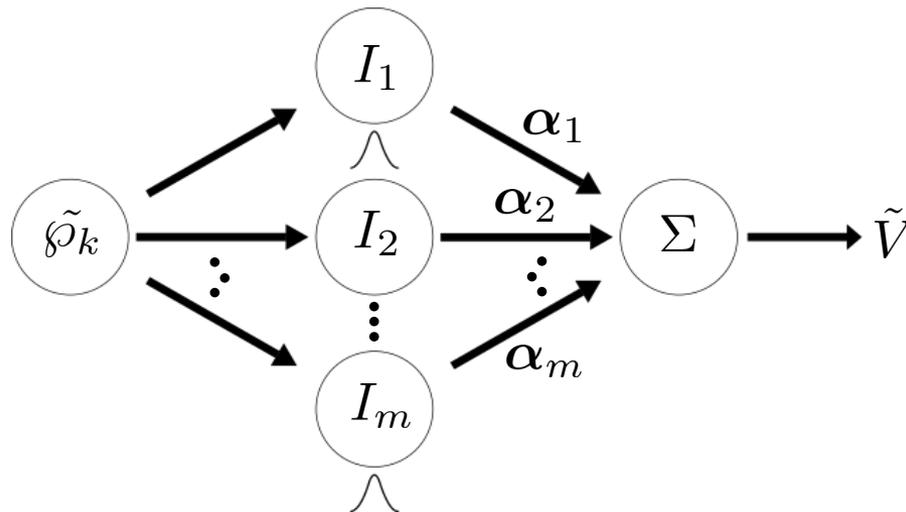
---

LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

# Functional Learning by Temporal Difference

# Value Function Learning

- Uses existing Temporal Difference (TD) and Recursive Least Squares Temporal Difference (RLSTD):
  - Learn parameters  $\alpha$  such that:  $\tilde{V}(\tilde{\phi}_k) = \alpha^T \mathbf{p}_k$





# Value Function Learning

- Batch Learning (Offline)

- Matrix  $\mathbf{A}$  is created with the outer product
- Vector  $\mathbf{b}$  is created with the inner product between  $\mathbf{z}$  and  $R$

$$\mathbf{A} = \sum_{i=0}^t \mathbf{z}_i (\mathbf{p}_i - \mathbf{p}_{i+1}) (\mathbf{p}_i - \mathbf{p}_{i+1})^T, \quad \mathbf{b} = \sum_{i=0}^t \mathbf{z}_i R_i$$

- where

$$\mathbf{z}_{t+1} = \lambda \mathbf{z}_t + \mathbf{p}_{t+1}, \quad \mathbf{z}_{t_0} \triangleq \mathbf{p}_{t_0}$$

- Solve with matrix inverse:

$$\boldsymbol{\alpha} = \mathbf{A}^{-1} \mathbf{b}$$



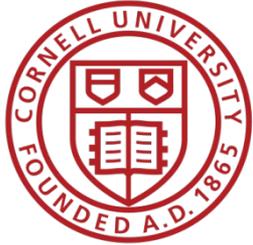
# Value Function Learning

- Recursive Learning (Online)
  - $e$  is an error term
  - $C$  is an intermediate term to compute parameters

$$e_t = R_t - (\mathbf{p}_t - \gamma \mathbf{p}_{t+1})^T \boldsymbol{\alpha}_{t-1}$$

$$\mathbf{C}_t = \mathbf{C}_{t-1} - \frac{\mathbf{C}_{t-1} \mathbf{p}_t (\mathbf{p}_t - \gamma \mathbf{p}_{t+1})^T \mathbf{C}_{t-1}}{1 + (\mathbf{p}_t - \gamma \mathbf{p}_{t+1})^T \mathbf{C}_{t-1} \mathbf{p}_t}$$

$$\boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} + \frac{\mathbf{C}_{t-1}}{1 + (\mathbf{p}_t - \gamma \mathbf{p}_{t+1})^T \mathbf{C}_{t-1} \mathbf{p}_t} e_t \mathbf{p}_t$$



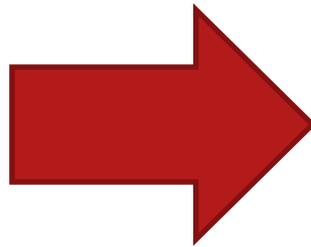
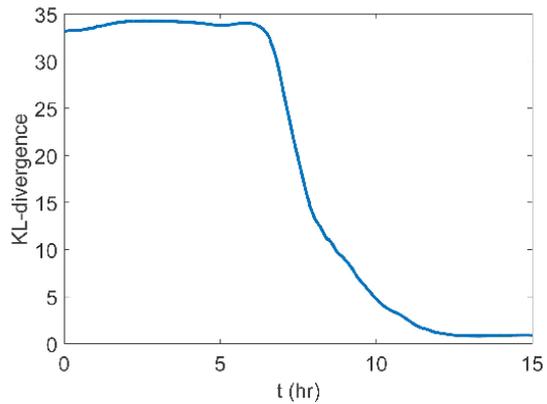
---

LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

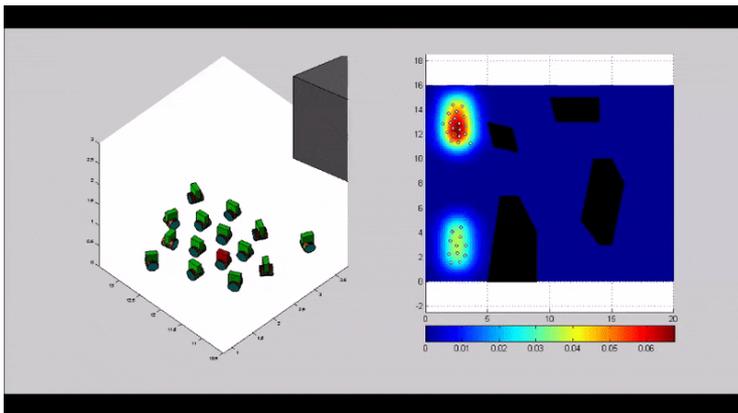
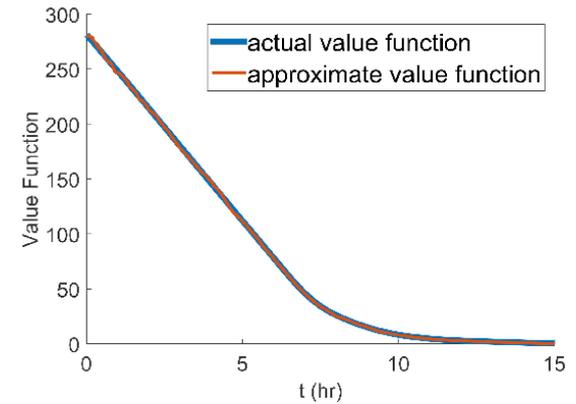
# Results and Conclusions

# Results

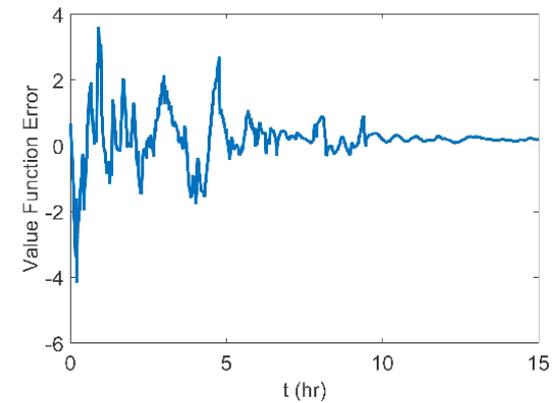
## Actual Cost Function



## Actual and Approximate Value function



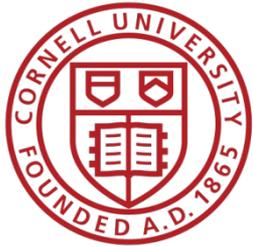
## Error:





## Conclusions

- Describe a collective of agents with a restriction operator
- Approximate the optimal value function on-line from the parameters of the PDF
  - Project GMM onto fixed basis
  - Using existing temporal difference algorithms
- **Future Work:**
  - Learn the optimal control law on-line by iteratively improving it



LABORATORY FOR INTELLIGENT  
SYSTEMS AND CONTROLS

# Value Function Approximation for Multiscale Dynamical Systems

Pingping Zhu, Julian Morelli

Advisor : Silvia Ferrari

Laboratory for Intelligent Systems and Controls

Conference for Decision and Control

Las Vegas

December 14, 2016

# Gaussian Identity Property

$$\begin{aligned}
 \mathbf{p}_k &= I_j(\tilde{\wp}_k) = \int_{\mathbf{x} \in \mathcal{X}} \tilde{\wp}_k(\mathbf{x}) b_j(\mathbf{x}) d\mathbf{x} \\
 &= \int_{\mathbf{x} \in \mathcal{X}} \left[ \sum_{\tau=1}^n \omega_\tau \mathcal{N}(\mathbf{x} | \mu_{k,\tau}, \Sigma_{k,\tau}) \right] \mathcal{N}(\mathbf{x} | \mu_j, \Sigma_j) d\mathbf{x} \\
 &= \sum_{\tau=1}^n \int_{\mathbf{x} \in \mathcal{X}} \omega_\tau \mathcal{N}(\mathbf{x} | \mu_{k,\tau}, \Sigma_{k,\tau}) \mathcal{N}(\mathbf{x} | \mu_j, \Sigma_j) d\mathbf{x} \\
 &= \sum_{\tau=1}^n \omega_\tau \mathcal{N}(\mu_{k,\tau} | \mu_j, \Sigma_{k,\tau} + \Sigma_j).
 \end{aligned}$$

# Gaussian Identity Property

$$\begin{aligned}
 I_j(\tilde{\varphi}_k) &= \left\langle \tilde{\varphi}_k, \frac{b_j}{\|b_j\|} \right\rangle_{\mathcal{P}} = \frac{\langle \tilde{\varphi}_k, b_j \rangle_{\mathcal{P}}}{\|b_j\|} = \frac{\langle \tilde{\varphi}_k, b_j \rangle_{\mathcal{P}}}{\langle b_j, b_j \rangle_{\mathcal{P}}} \\
 &= \frac{\int_{\mathbf{x} \in \mathcal{X}} \tilde{\varphi}_k(\mathbf{x}) b_j(\mathbf{x}) d\mathbf{x}}{\int_{\mathbf{x} \in \mathcal{X}} b_j(\mathbf{x}) b_j(\mathbf{x}) d\mathbf{x}} \\
 &= \frac{\sum_{j=1}^n \omega_{\tau} \mathcal{N}(\mu_{k,\tau} | \mu_j, \Sigma_{k,\tau} + \Sigma_j)}{\mathcal{N}(\mu_j | \mu_j, 2\Sigma_j)},
 \end{aligned}$$