On the Design of LQR Kernels for Efficient Controller Learning

Alonso Marco[1], Philipp Hennig[1], Stefan Schaal[1,2] and Sebastian Trimpe[1]

Learning-based control
Invited session
Chair: Sebastian Trimpe
Co-chair: Angela P. Schoellig

56th IEEE Conference on Decision and Control
Melbourne, Australia
12-15 Dec, 2017

[1]Autonomous Motion Department, Max Planck Institute for Intelligent Systems, Tübingen, Germany
[2]Computational Learning and Motor Control Lab, University of Southern California, Los Angeles, USA

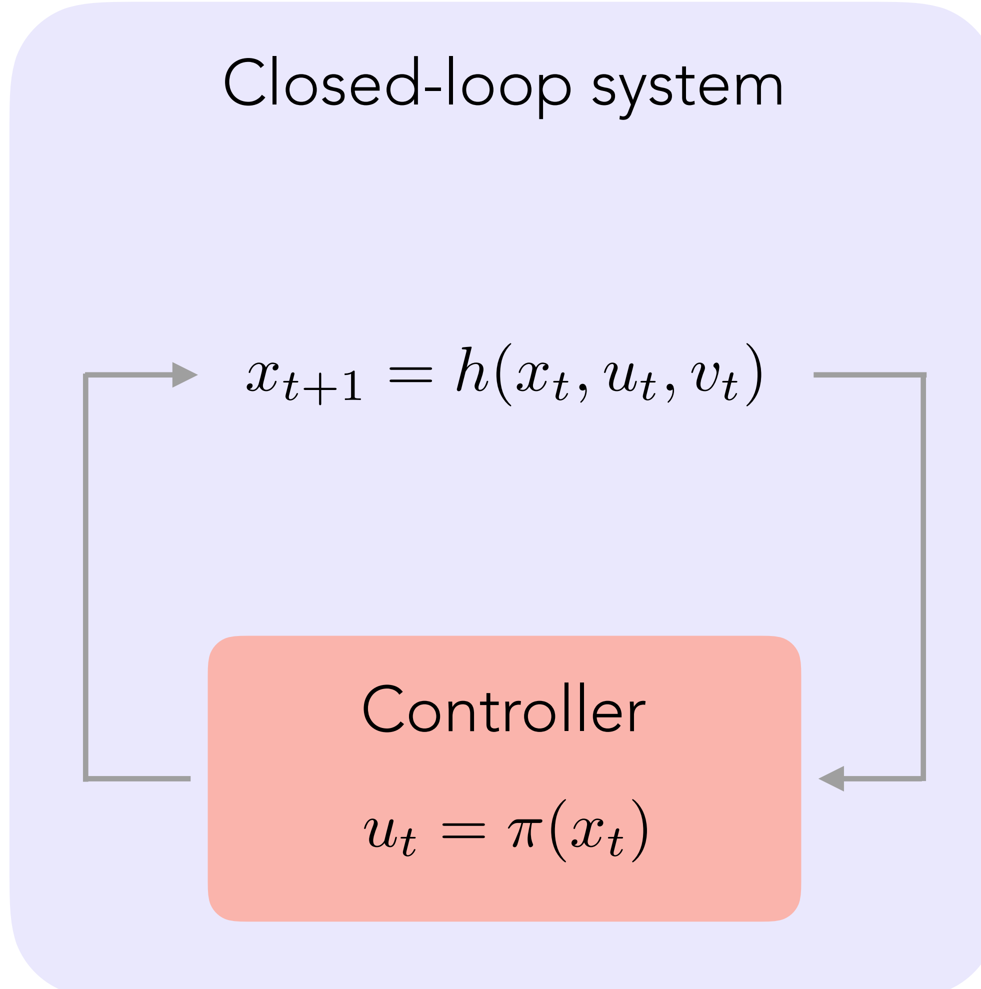

x_t : states

u_t : control input

v_t : process noise

π : feedback controller

Alonso Marco | Max Planck Institute for Intelligent Systems

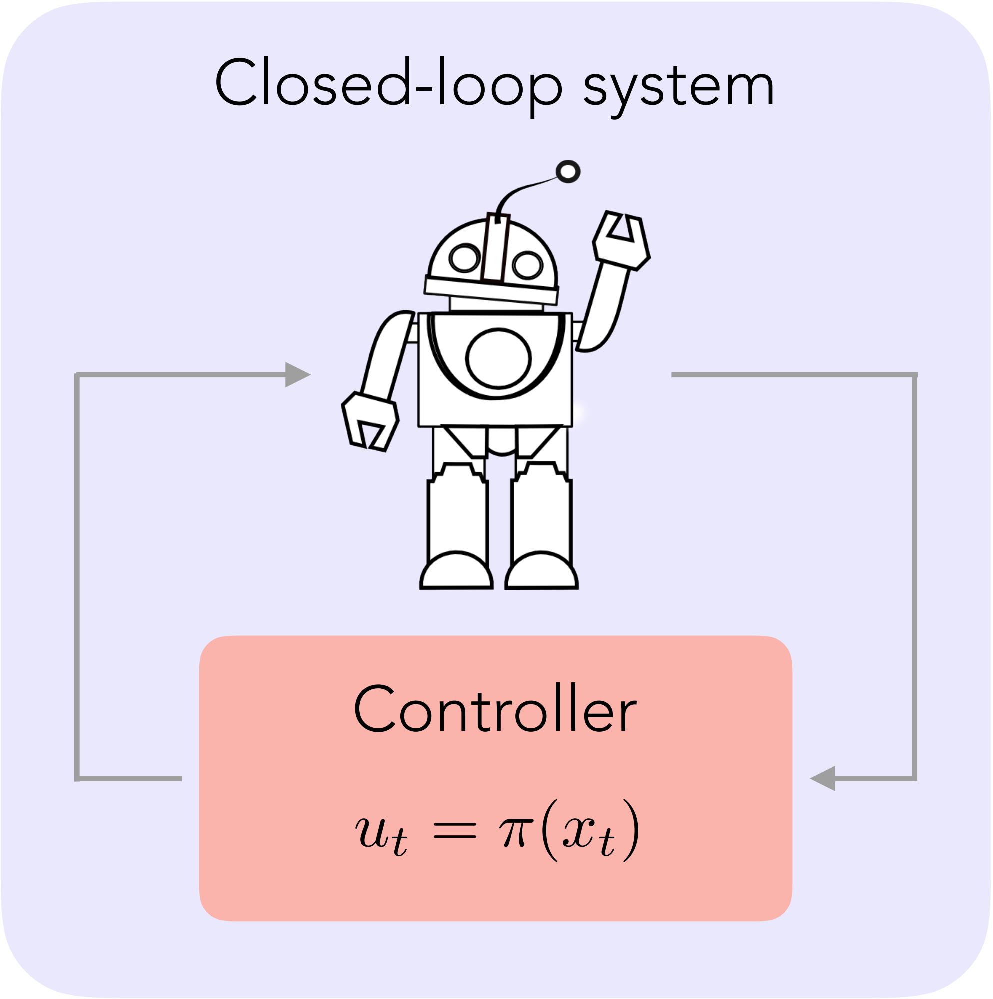

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

π : feedback controller

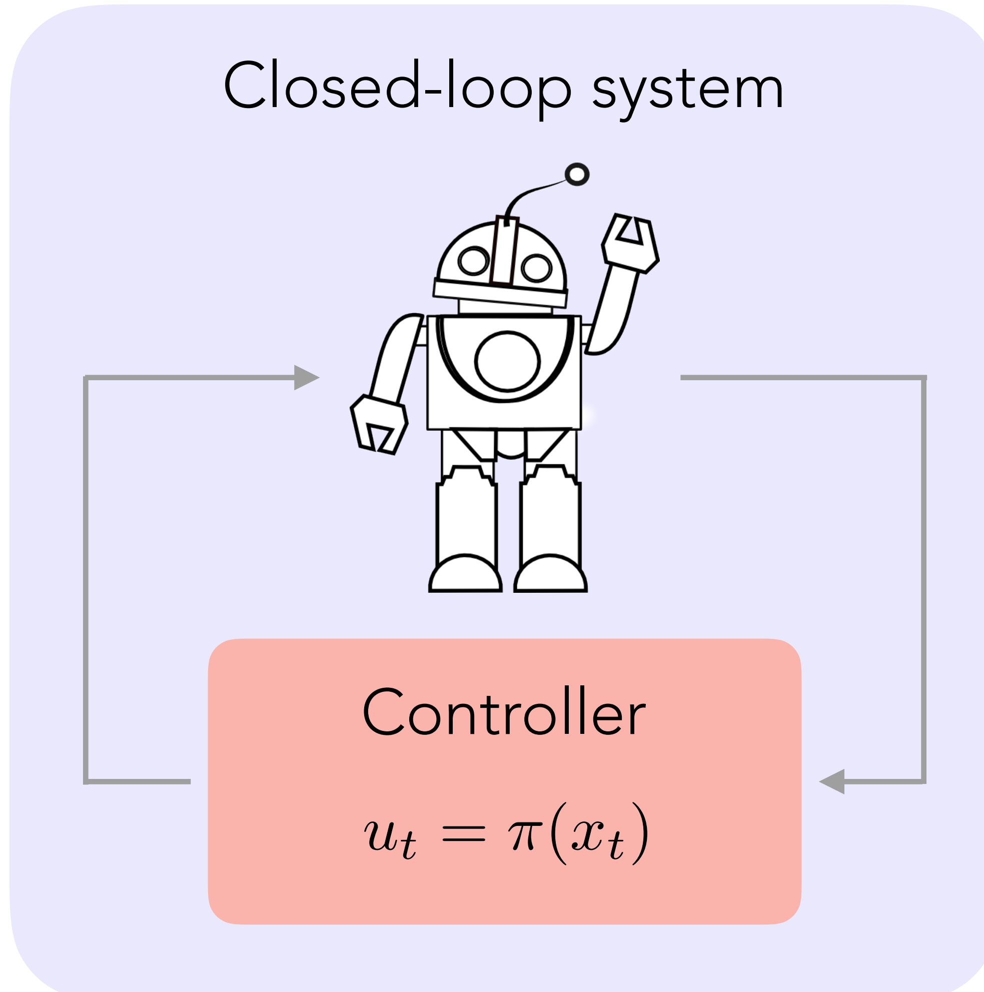

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

π : feedback controller

Linearized model $\tilde{x}_{t+1} = A\tilde{x}_t + B\tilde{x}_t + v_t$

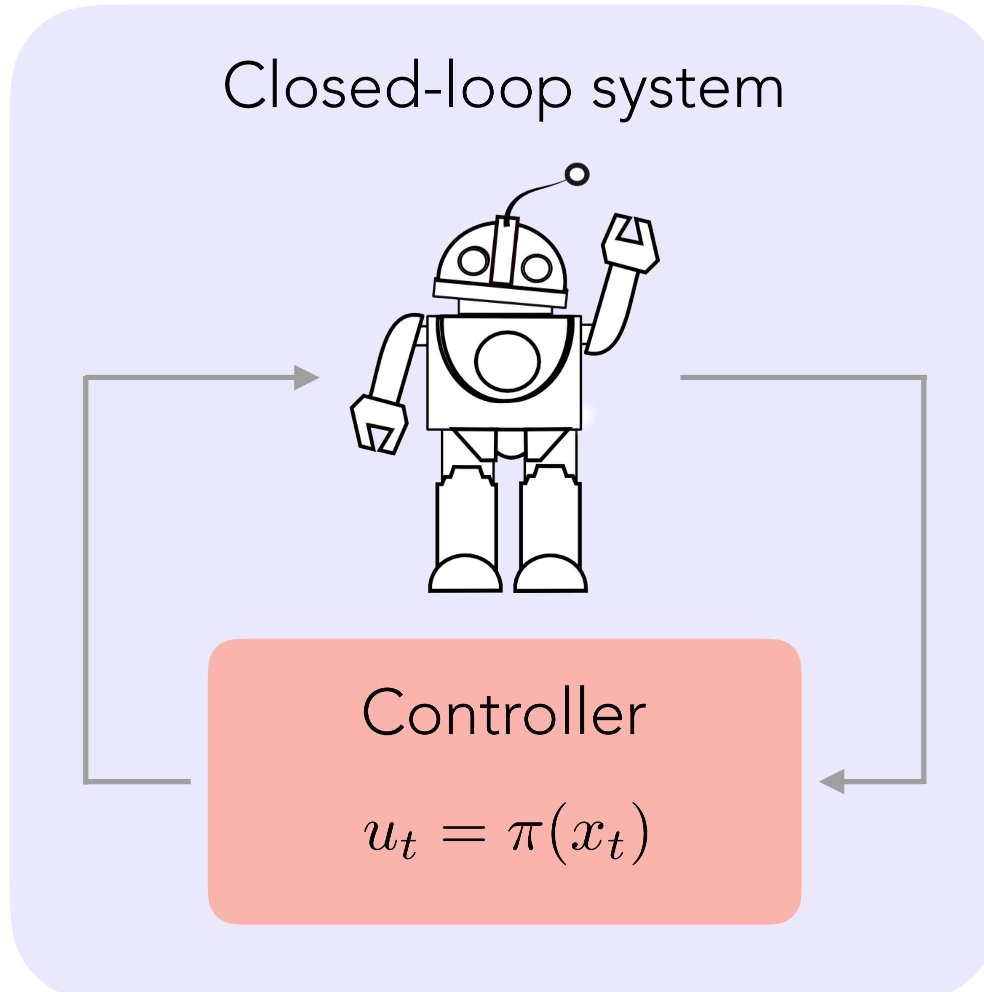

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

π : feedback controller

Linearized model $\quad \tilde{x}_{t+1} = A\tilde{x}_t + B\tilde{x}_t + v_t$

Quadratic cost $\quad J = \lim_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=0}^{T-1} \tilde{x}_t^{\mathrm{T}} Q \tilde{x}_t + \tilde{u}_t^{\mathrm{T}} R \tilde{u}_t \right]$

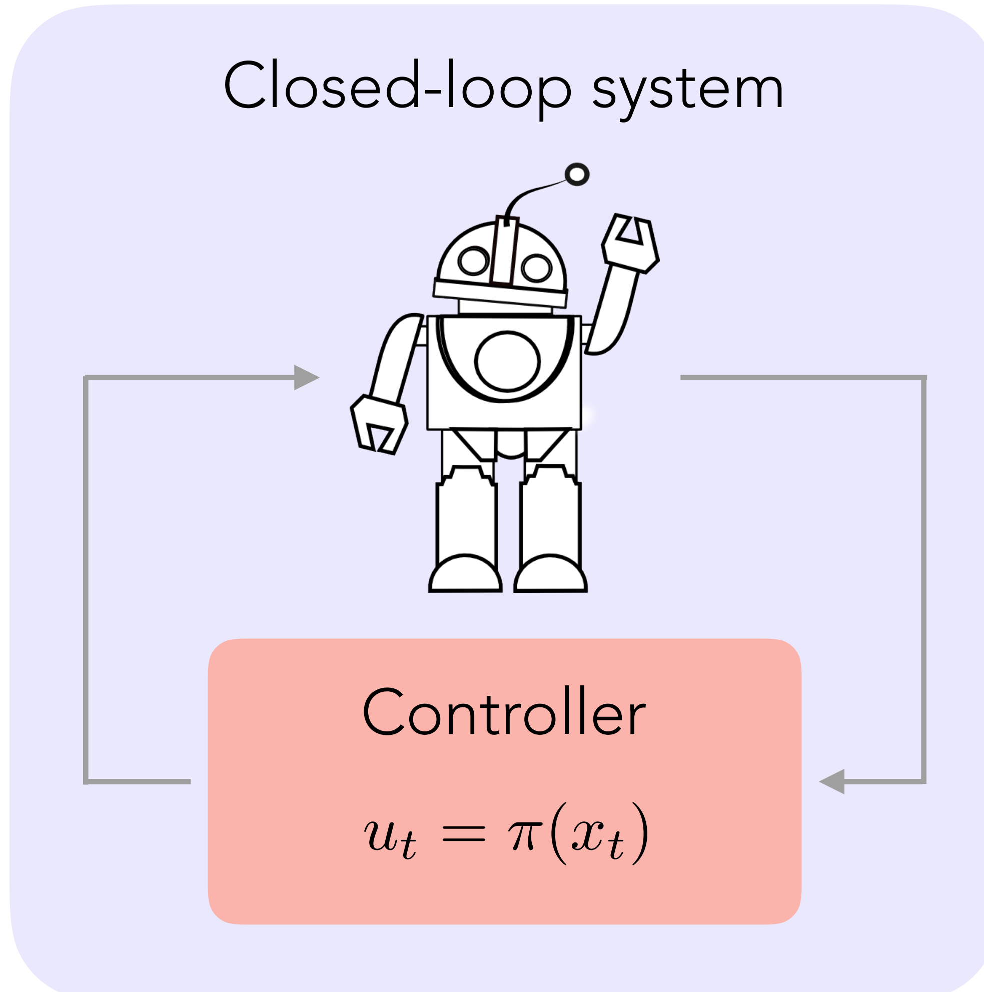

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

π : feedback controller

Linearized model $$\tilde{x}_{t+1} = A\tilde{x}_t + B\tilde{x}_t + v_t$$

Quadratic cost $$J = \lim_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=0}^{T-1} \tilde{x}_t^{\mathrm{T}} Q \tilde{x}_t + \tilde{u}_t^{\mathrm{T}} R \tilde{u}_t \right]$$

Which controller $u_t = \pi(x_t)$ minimizes J ?

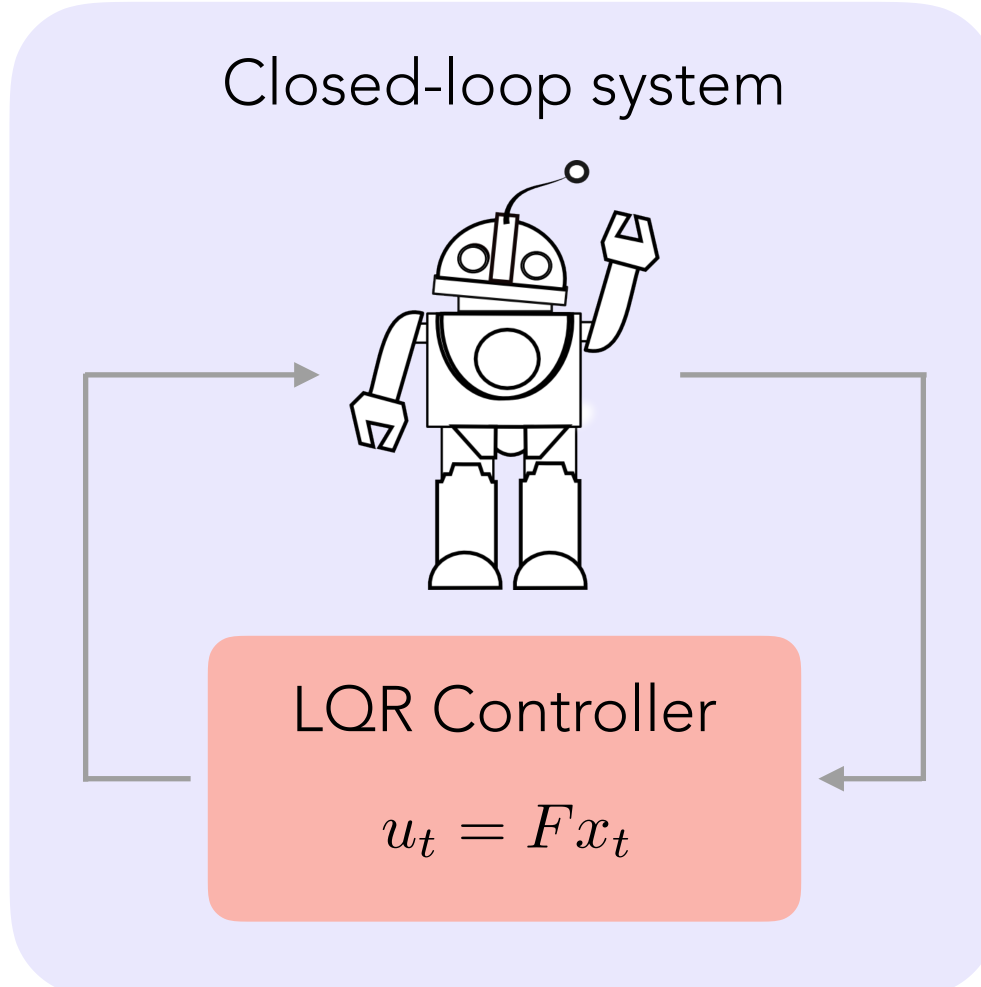

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

F : feedback gain

Linearized model $\quad \tilde{x}_{t+1} = A\tilde{x}_t + B\tilde{x}_t + v_t$

Quadratic cost $\quad J = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=0}^{T-1} \tilde{x}_t^{\mathrm{T}} Q \tilde{x}_t + \tilde{u}_t^{\mathrm{T}} R \tilde{u}_t\right]$

If 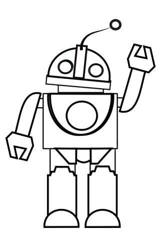 was linear (A, B) ⇒ LQR is optimal $\begin{cases} u_t = F x_t \\ F = \mathrm{lqr}(A, B, Q, R) \end{cases}$

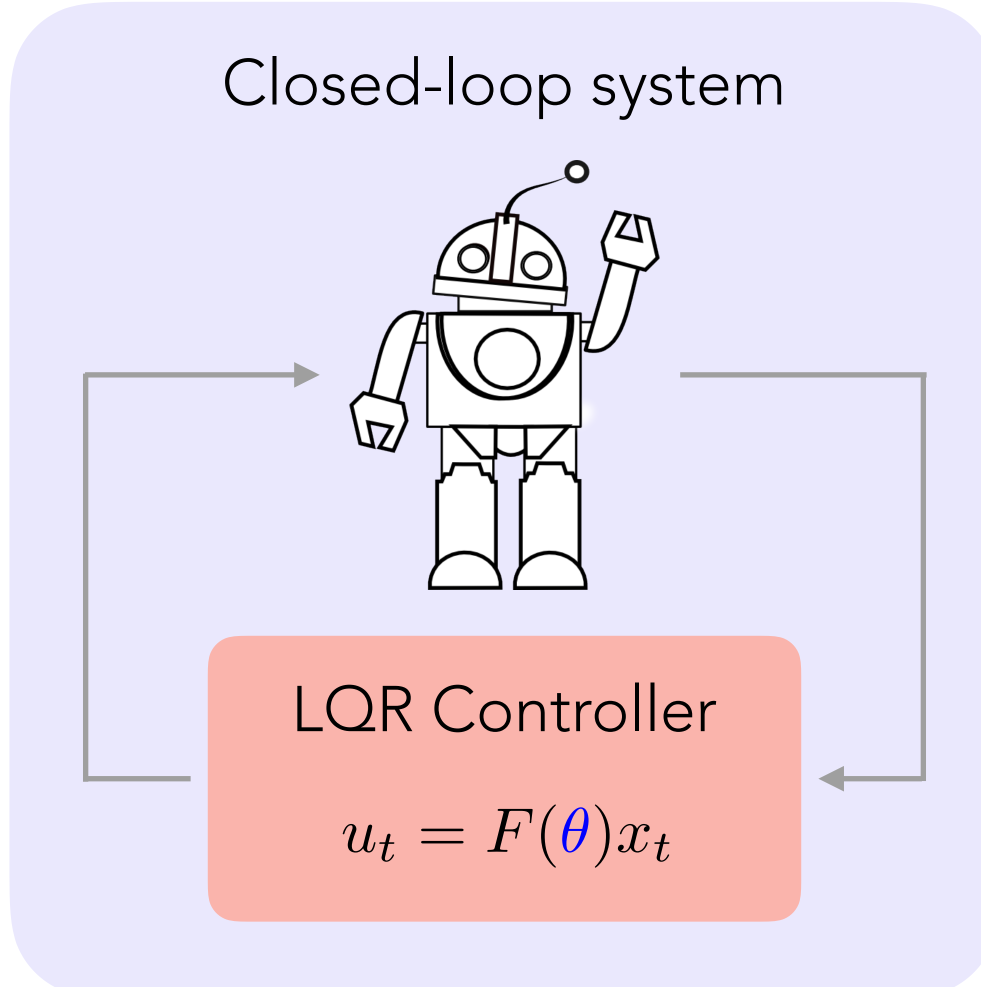

$$x_{t+1} = h(x_t, u_t, v_t)$$

x_t : states

u_t : control input

v_t : process noise

$F(\theta)$: feedback gain

θ : parameters

Linearized model $\quad \tilde{x}_{t+1} = A\tilde{x}_t + B\tilde{x}_t + v_t$

Quadratic cost $\quad J = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=0}^{T-1} \tilde{x}_t^{\mathrm{T}} Q \tilde{x}_t + \tilde{u}_t^{\mathrm{T}} R \tilde{u}_t\right]$

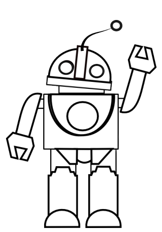
is non-linear → LQR is suboptimal → Parametrize $F(\theta)$

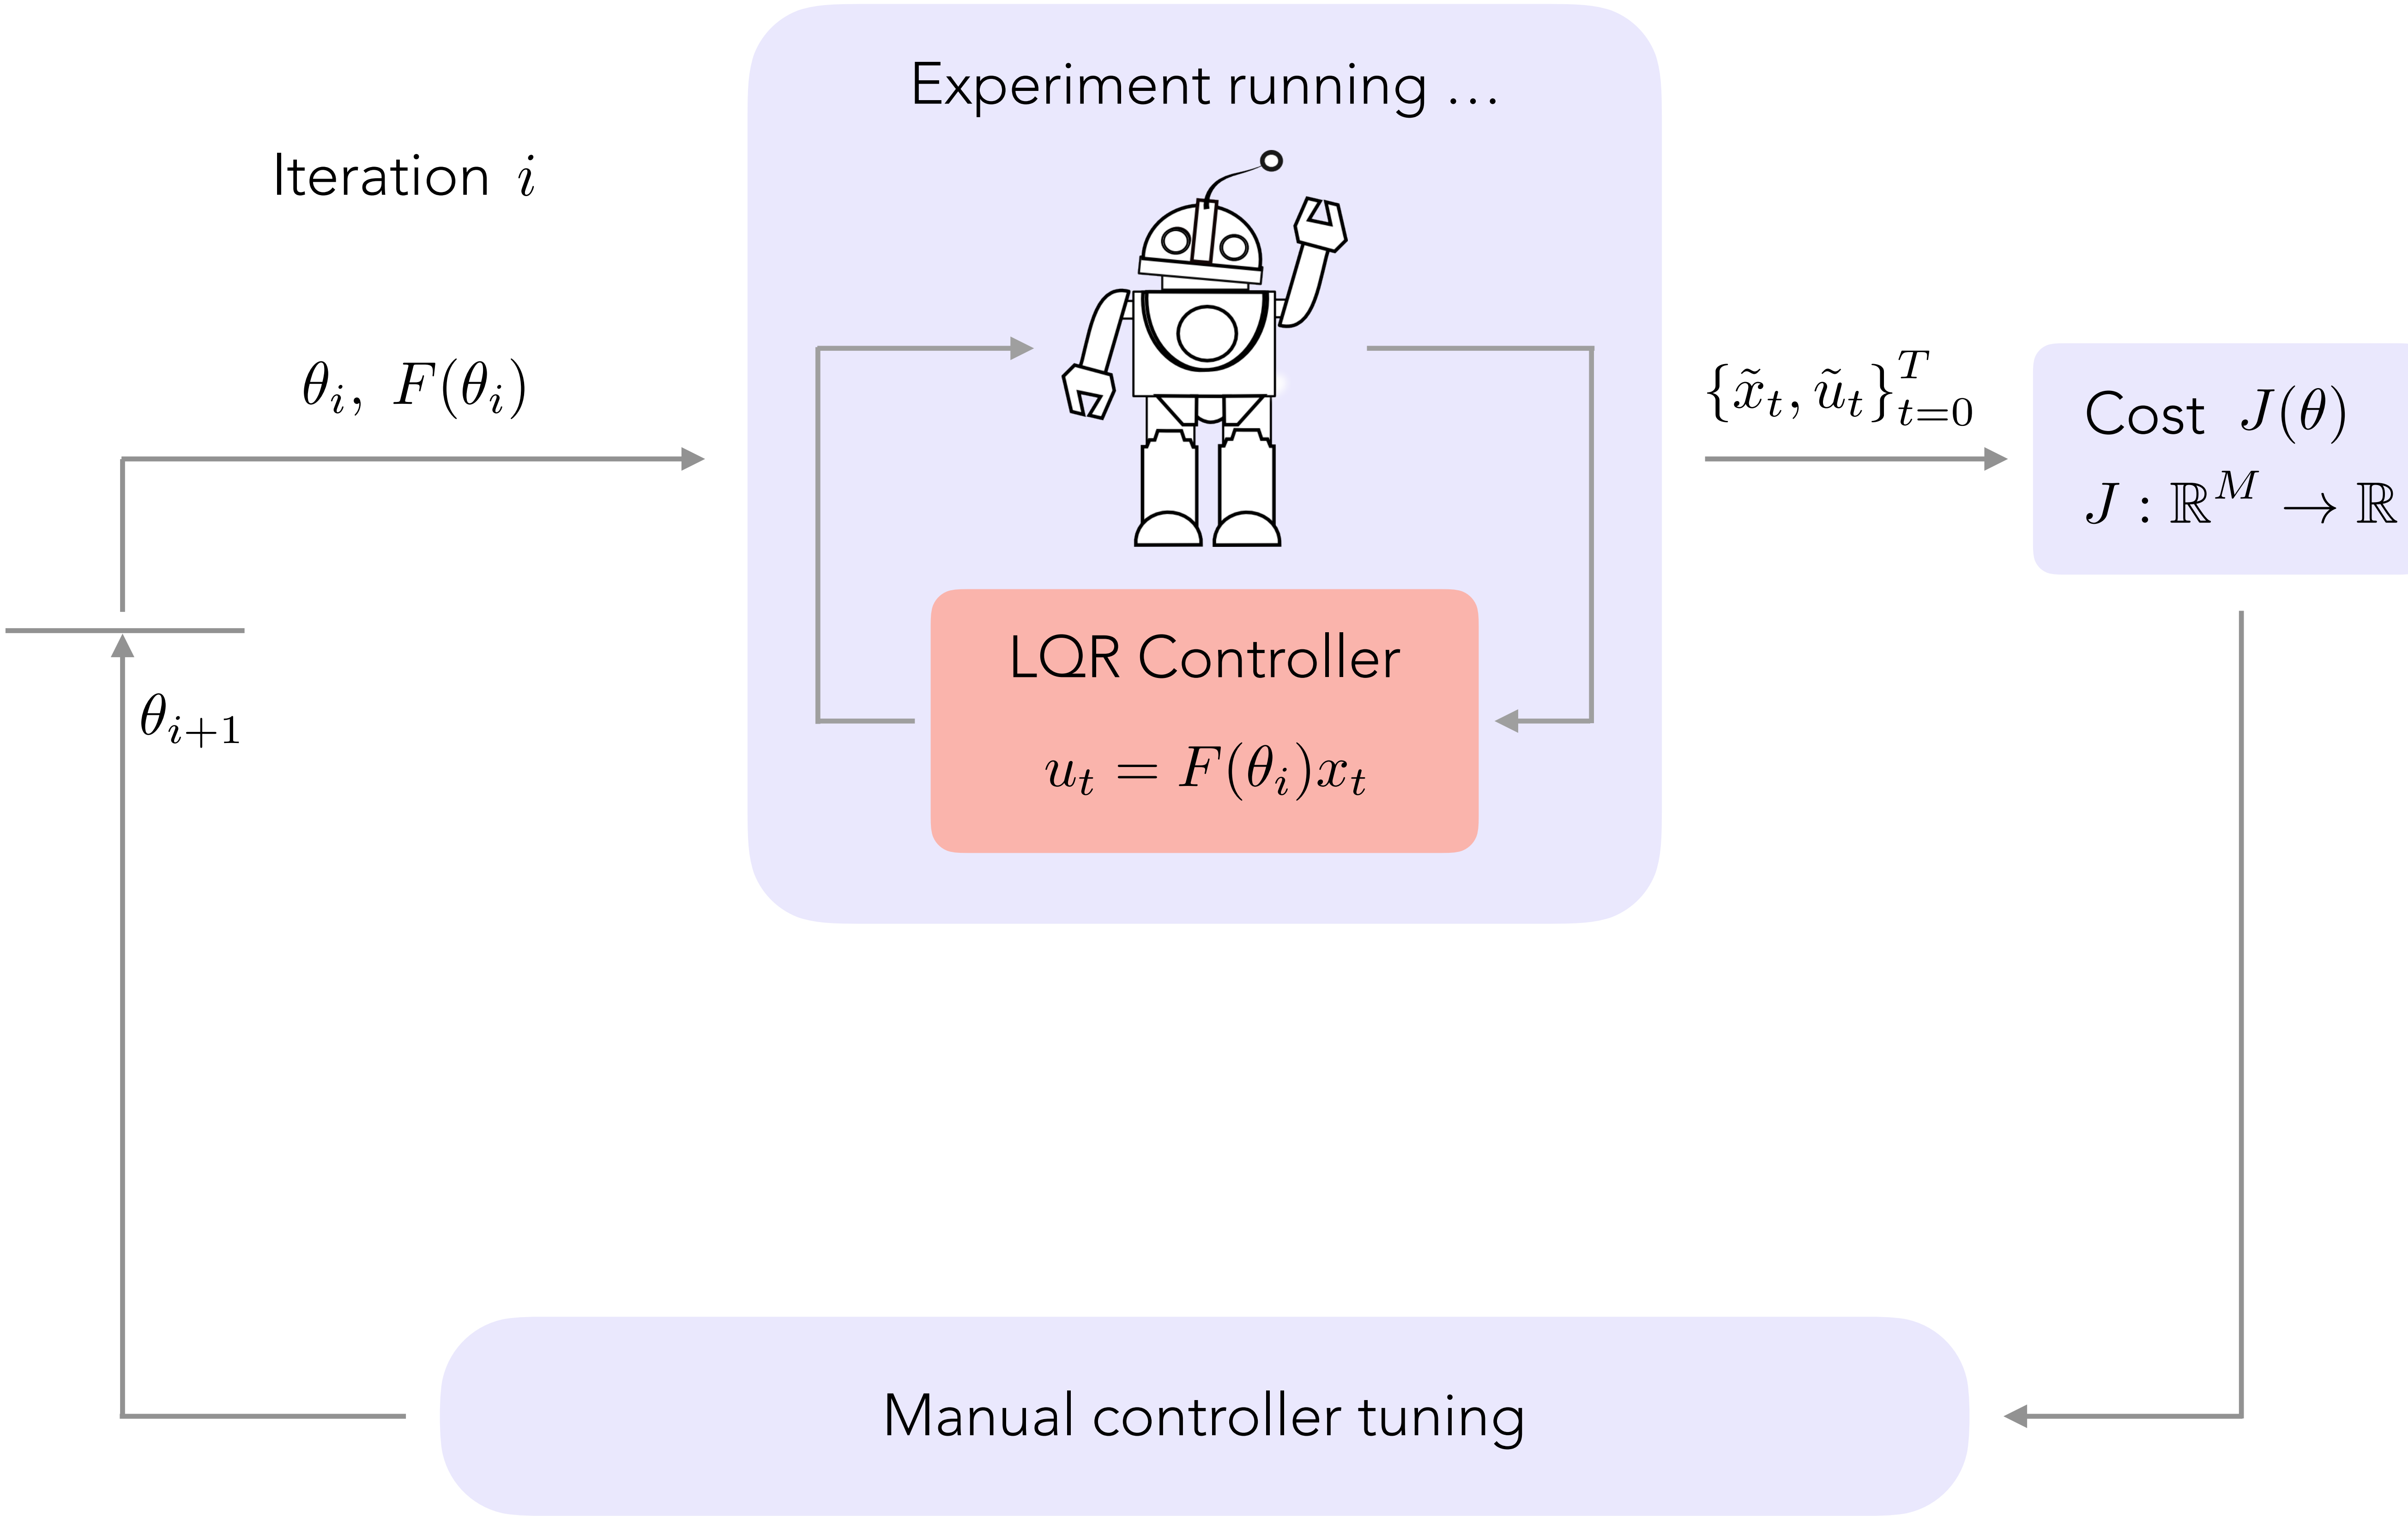
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \rightarrow \mathbb{R}$
LQR Controller
$u_t = F(\theta_i)x_t$
θ_{i+1}
Manual controller tuning

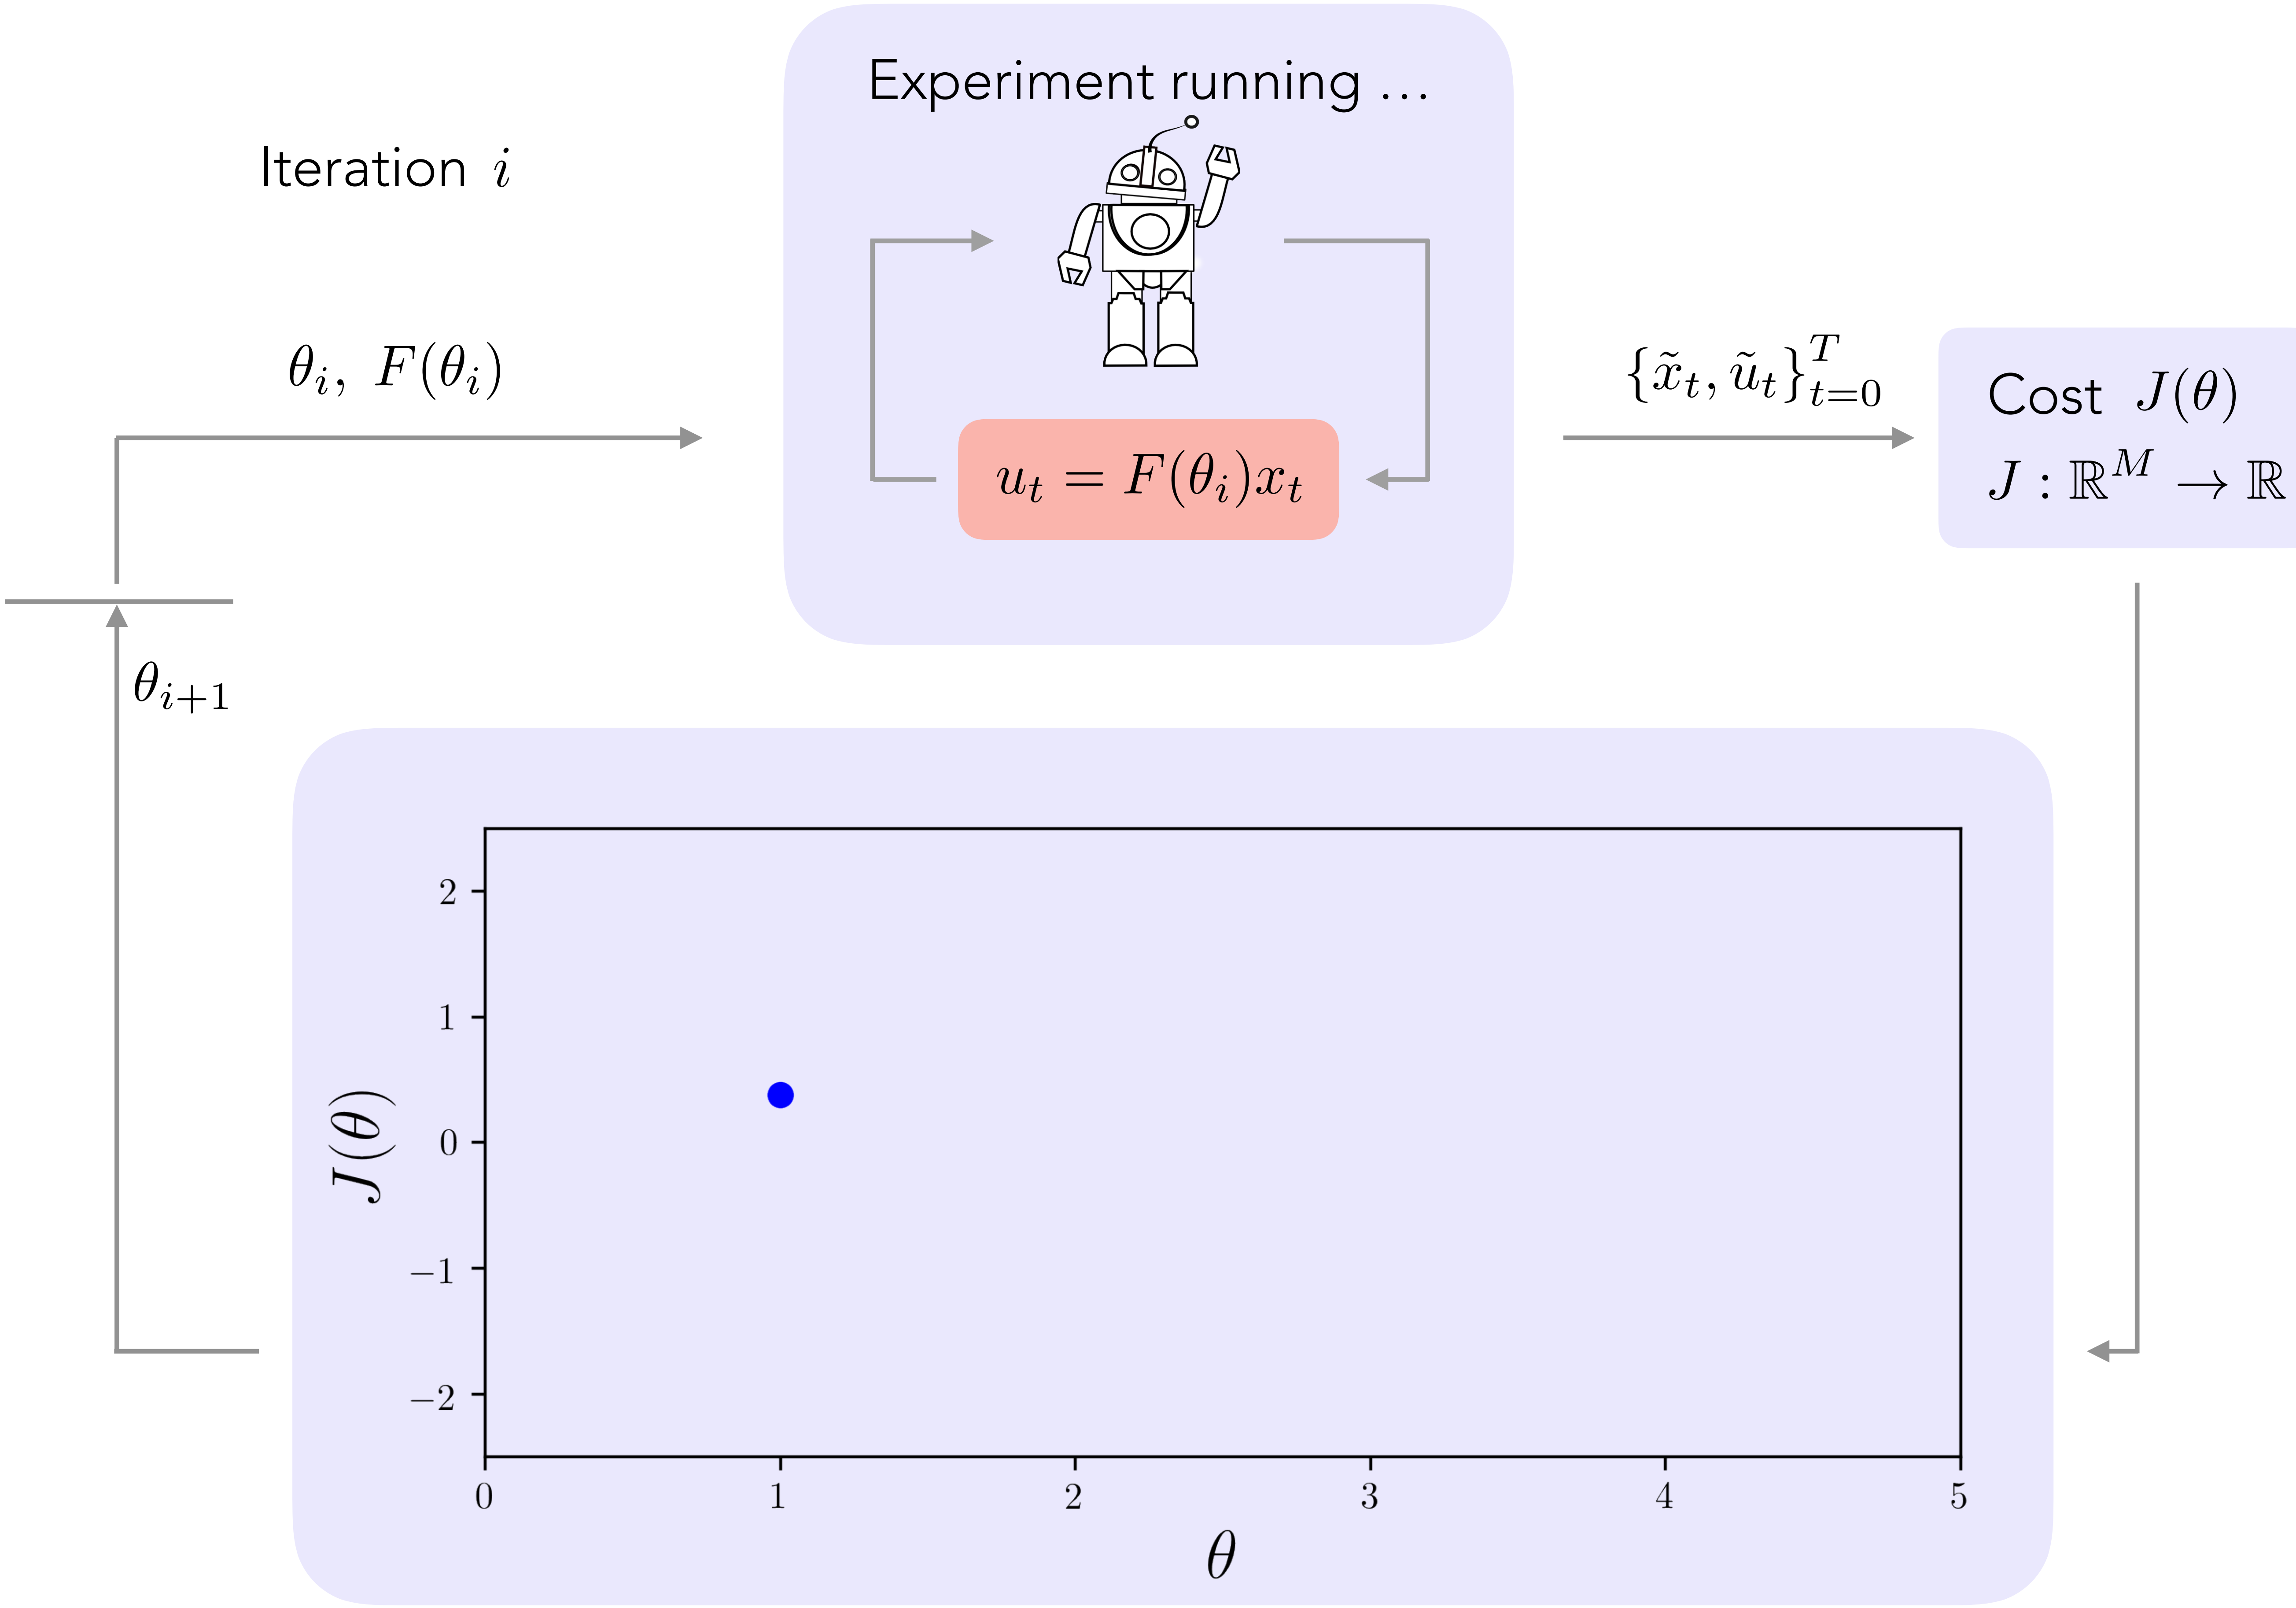
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \rightarrow \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

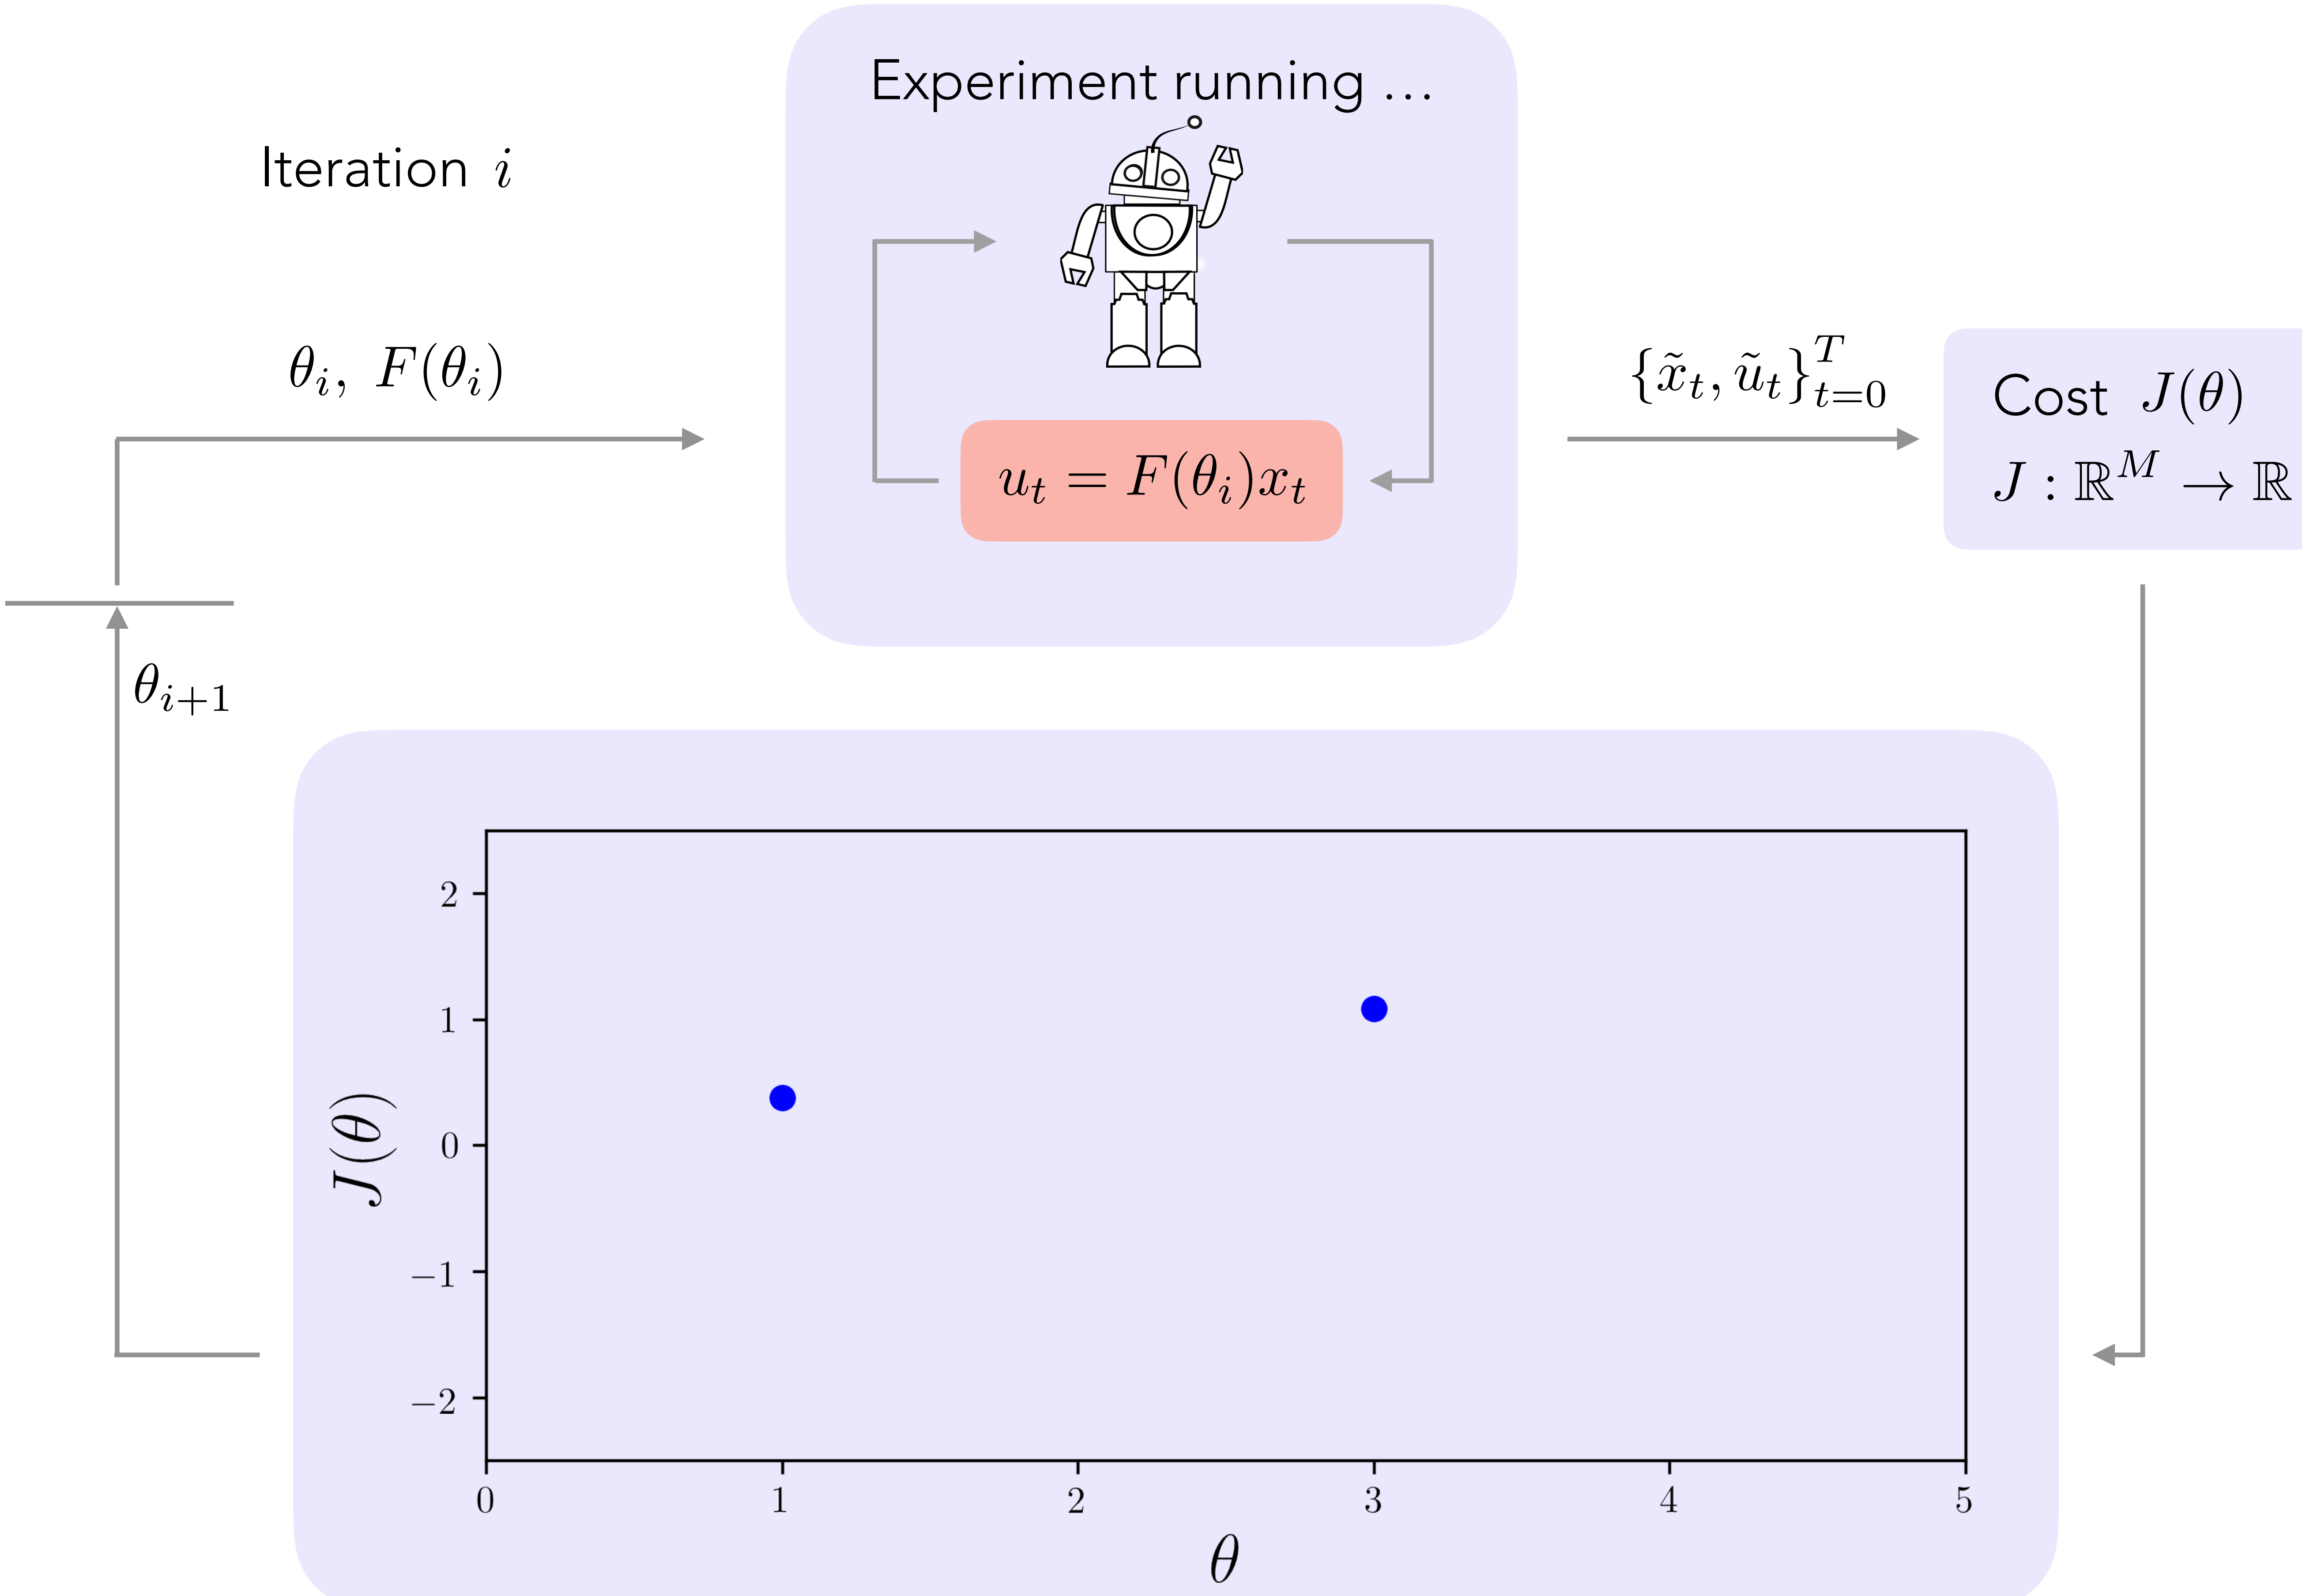
Iteration i
Experiment running ...
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

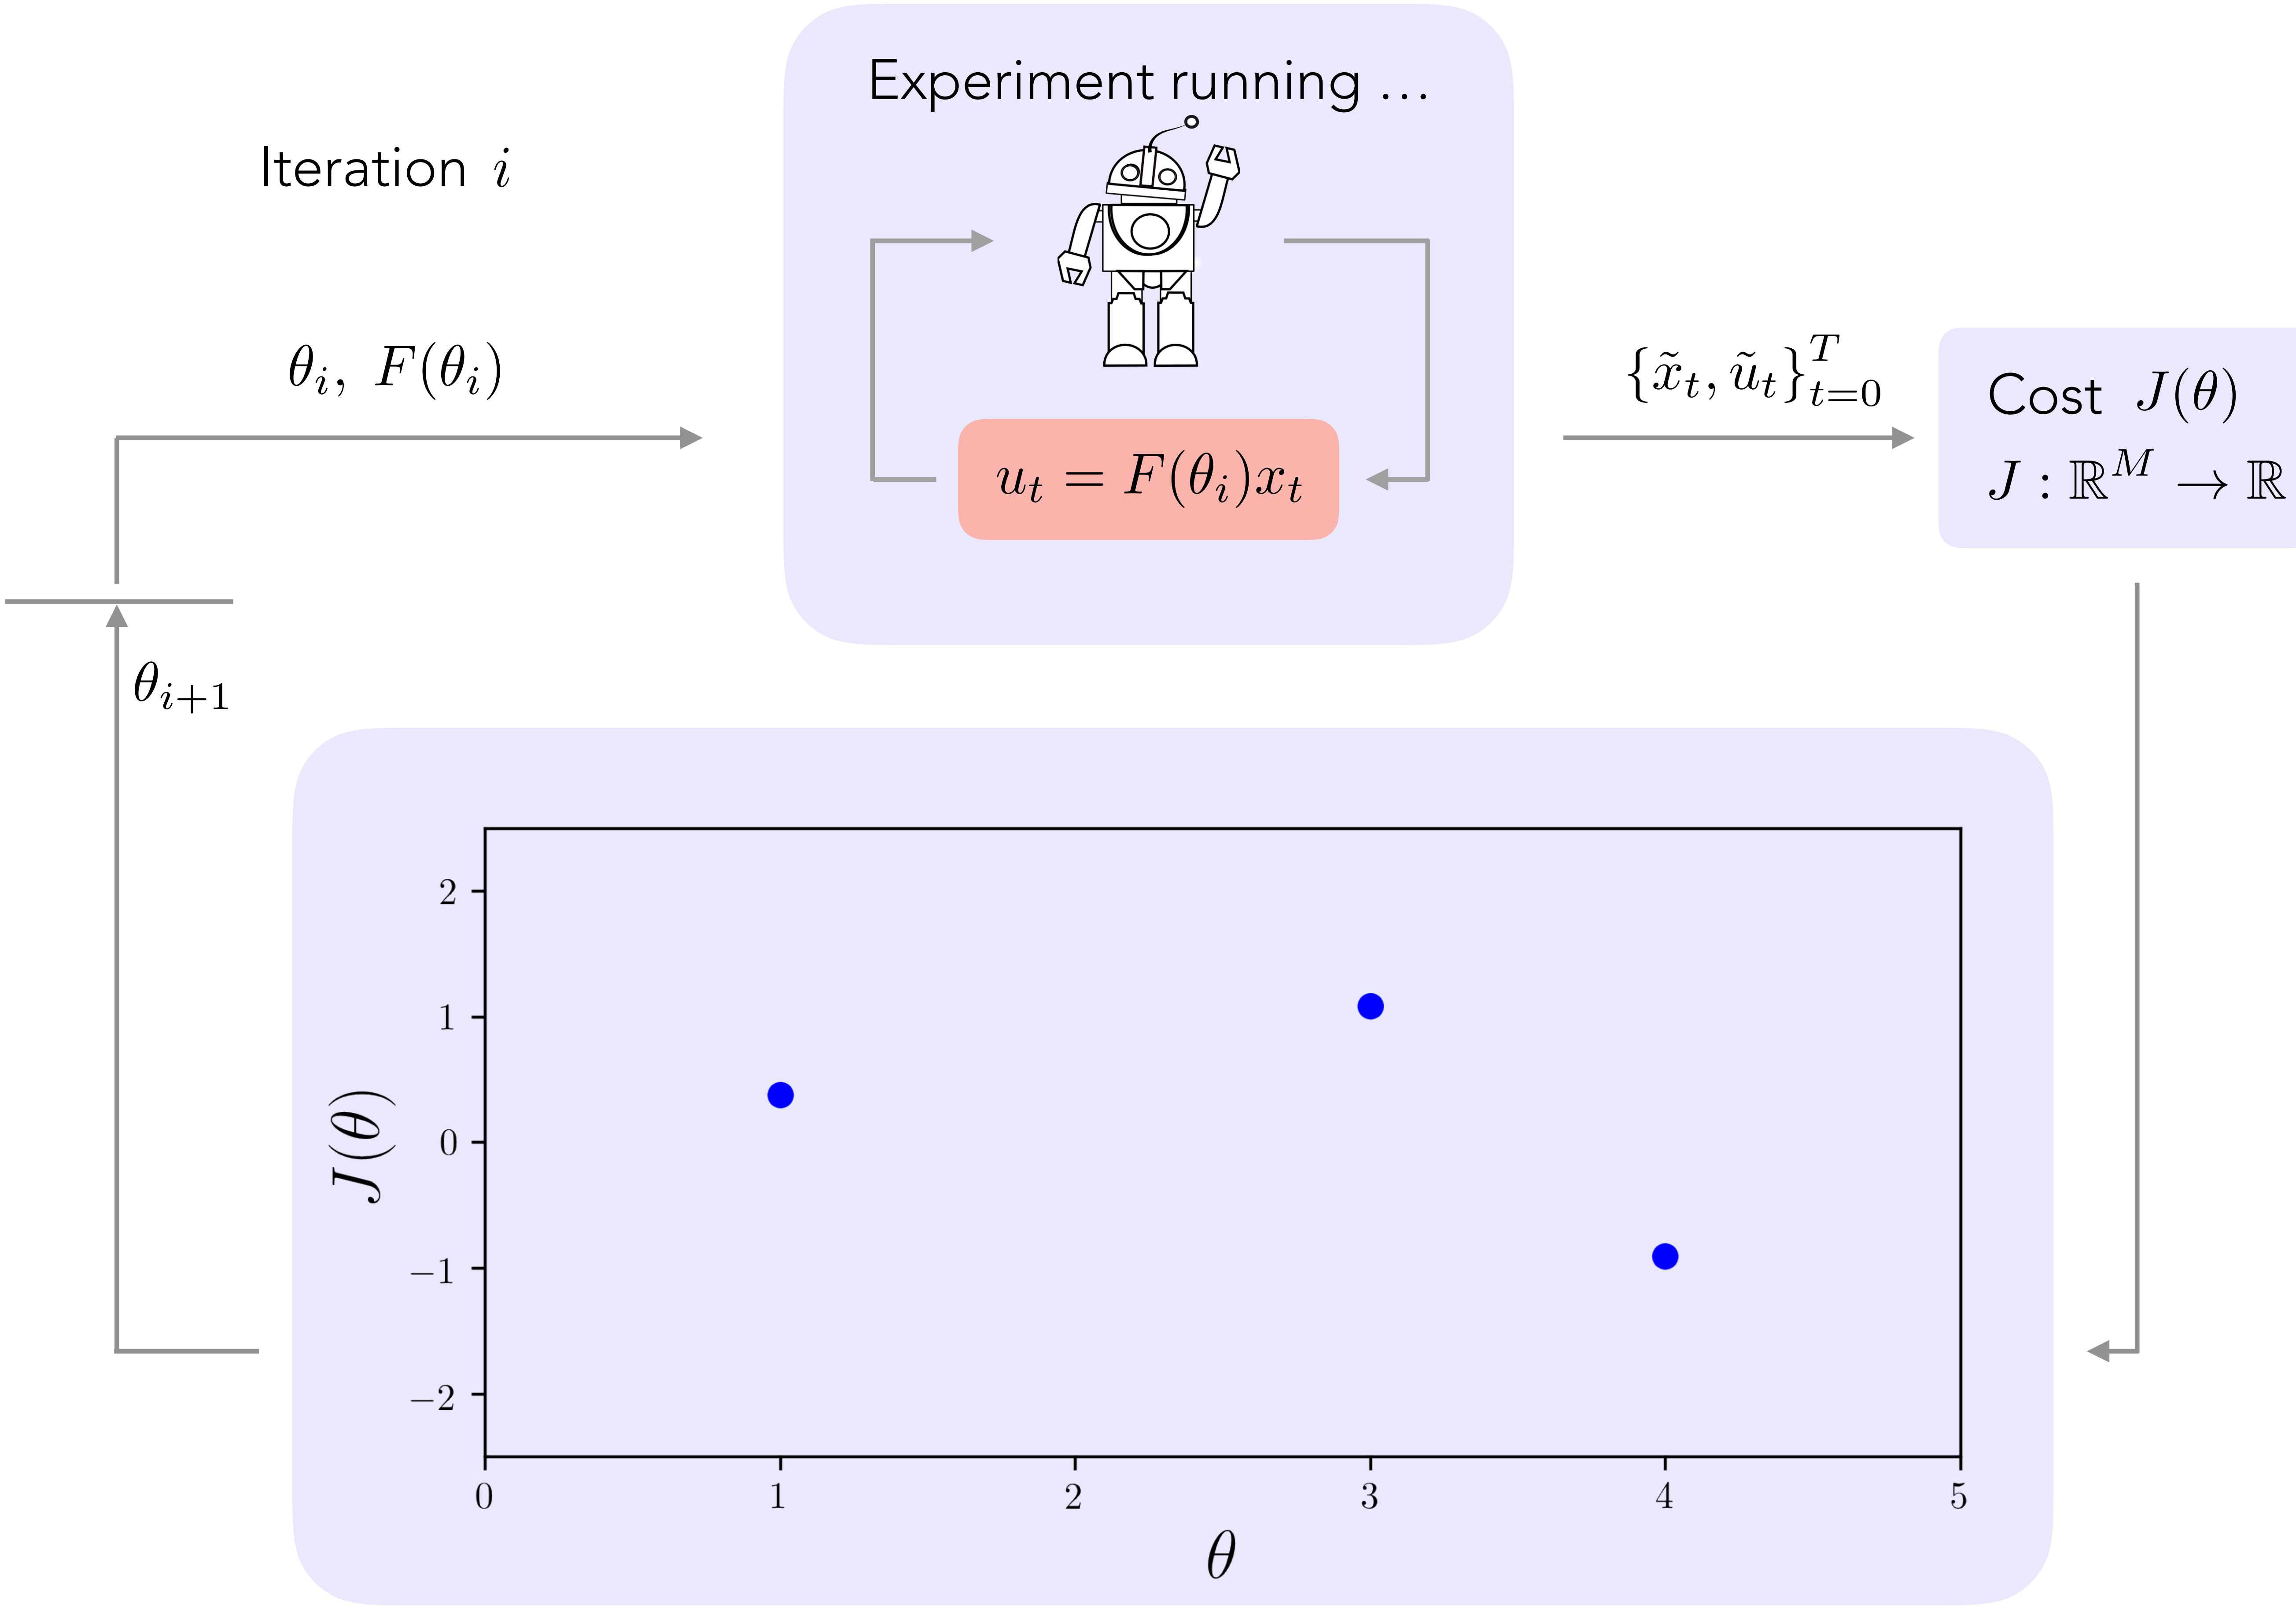
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

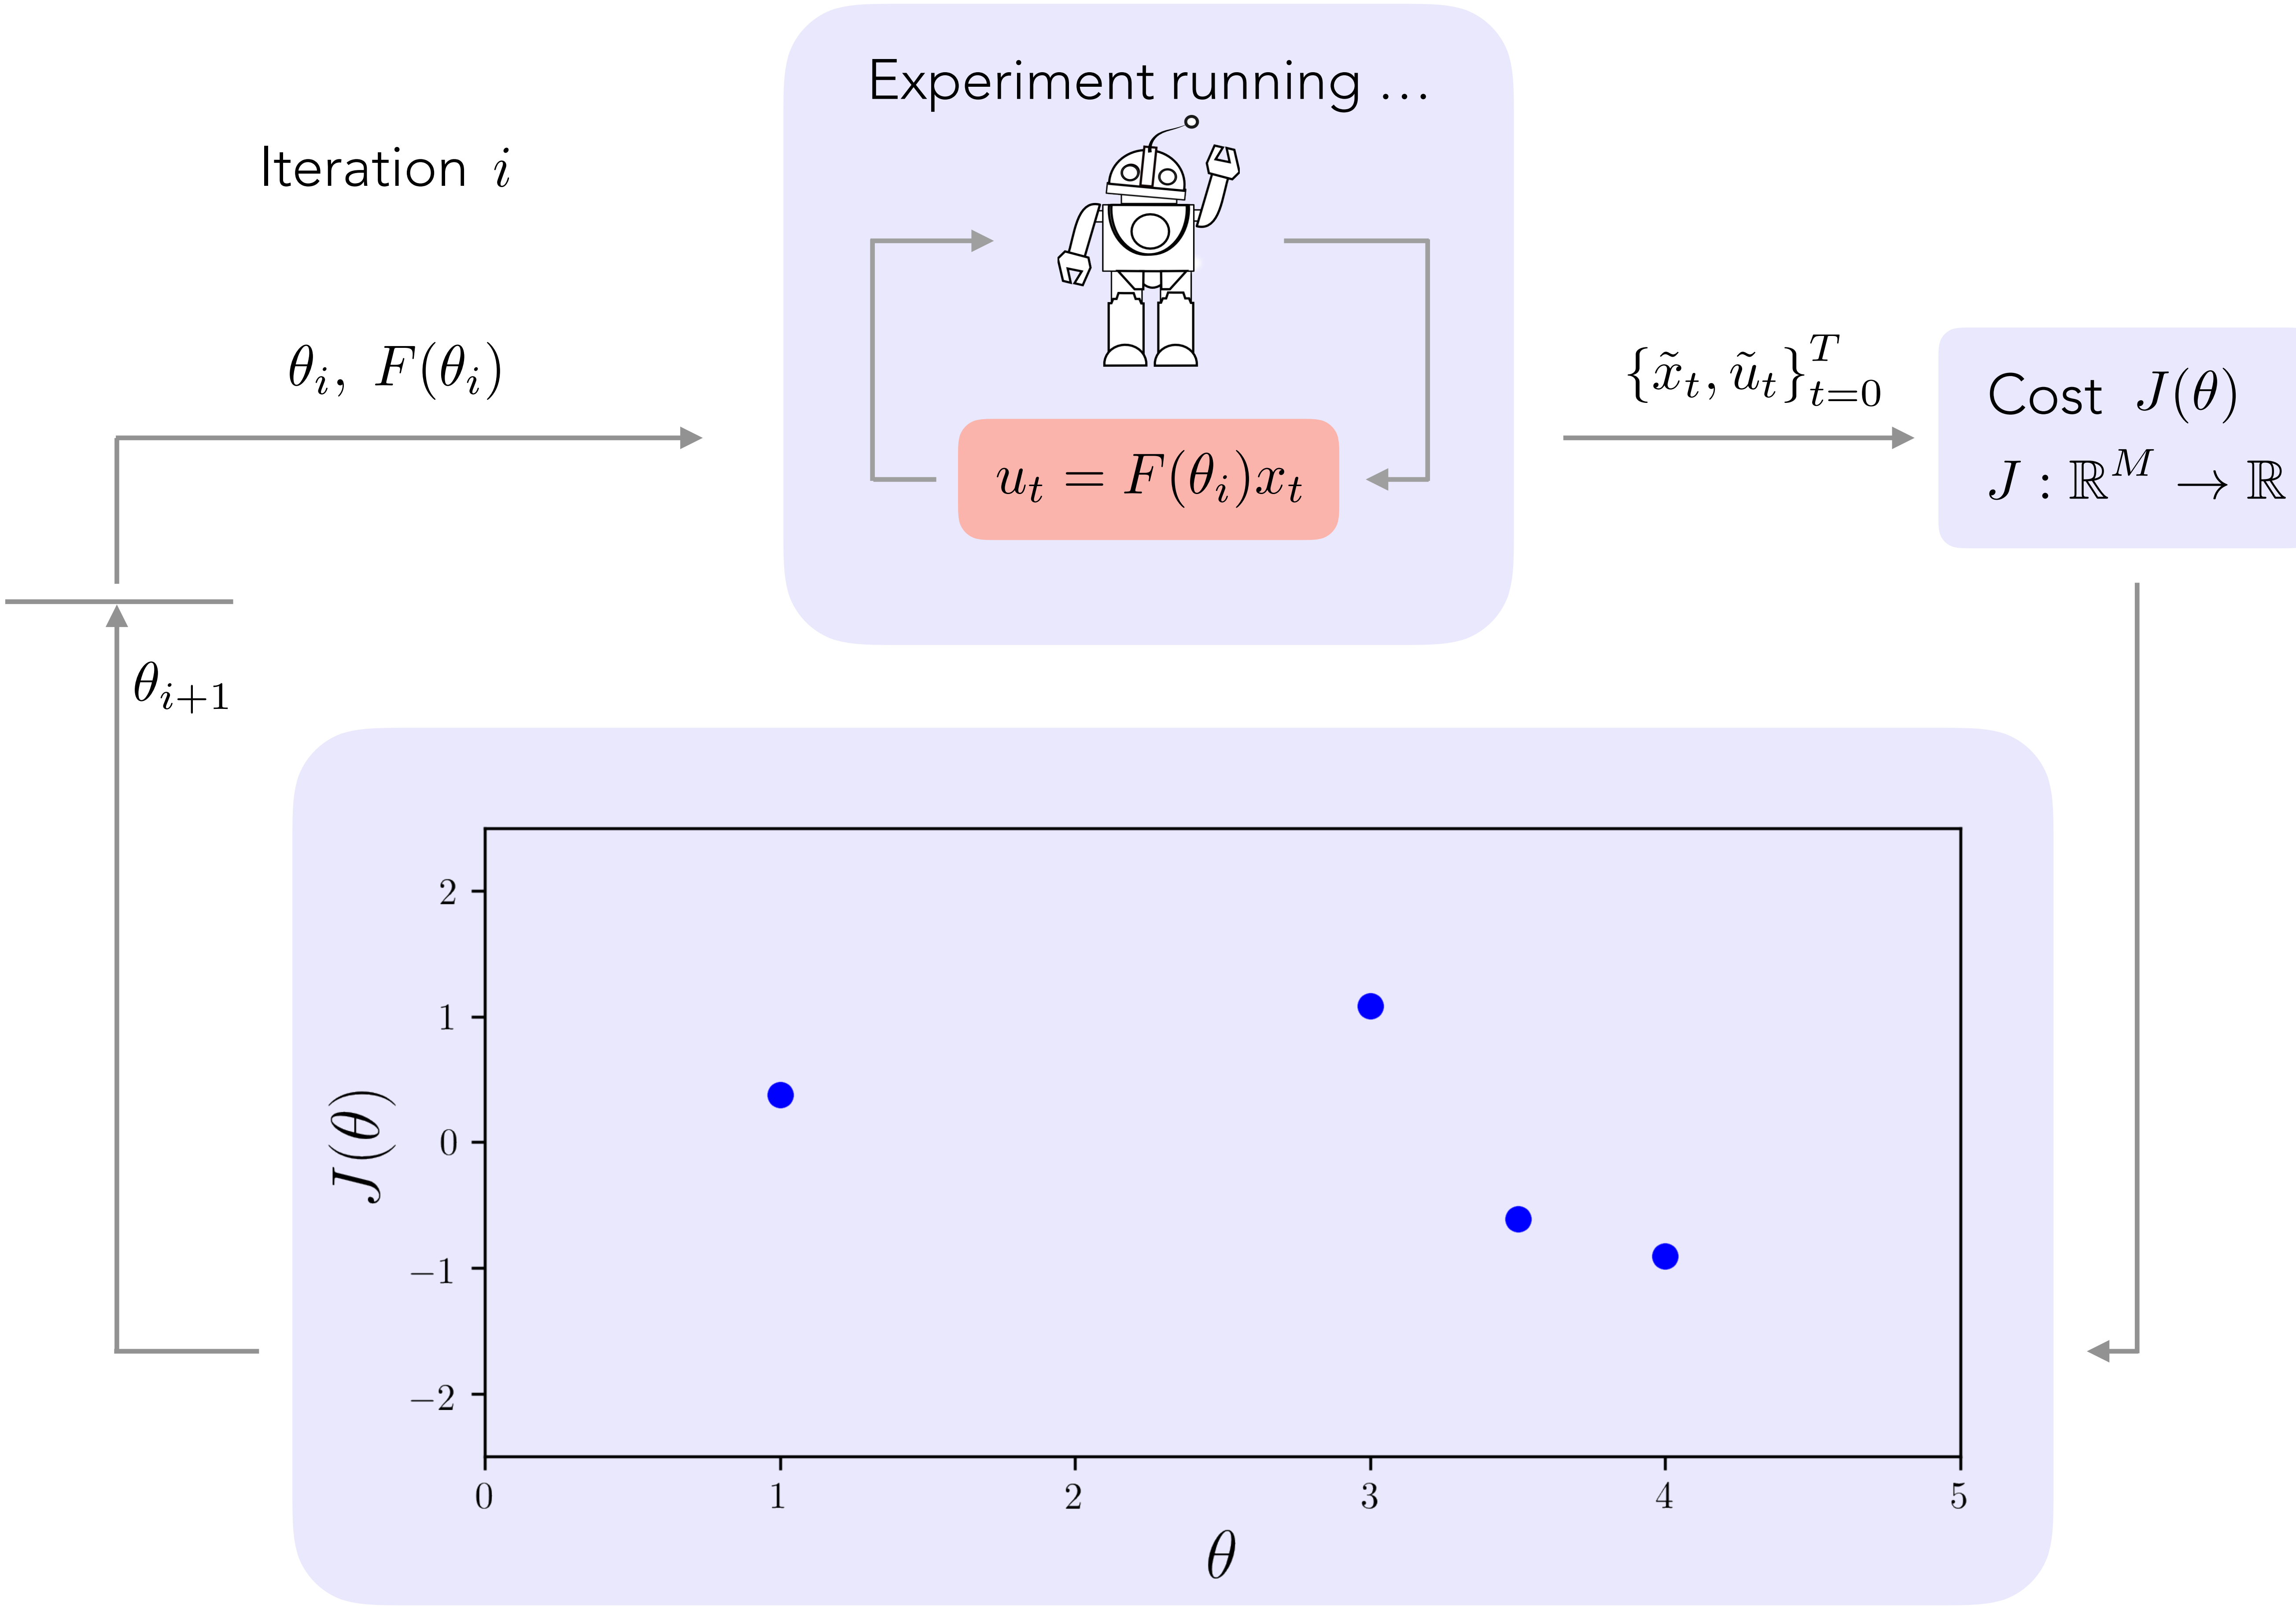
Experiment running …
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

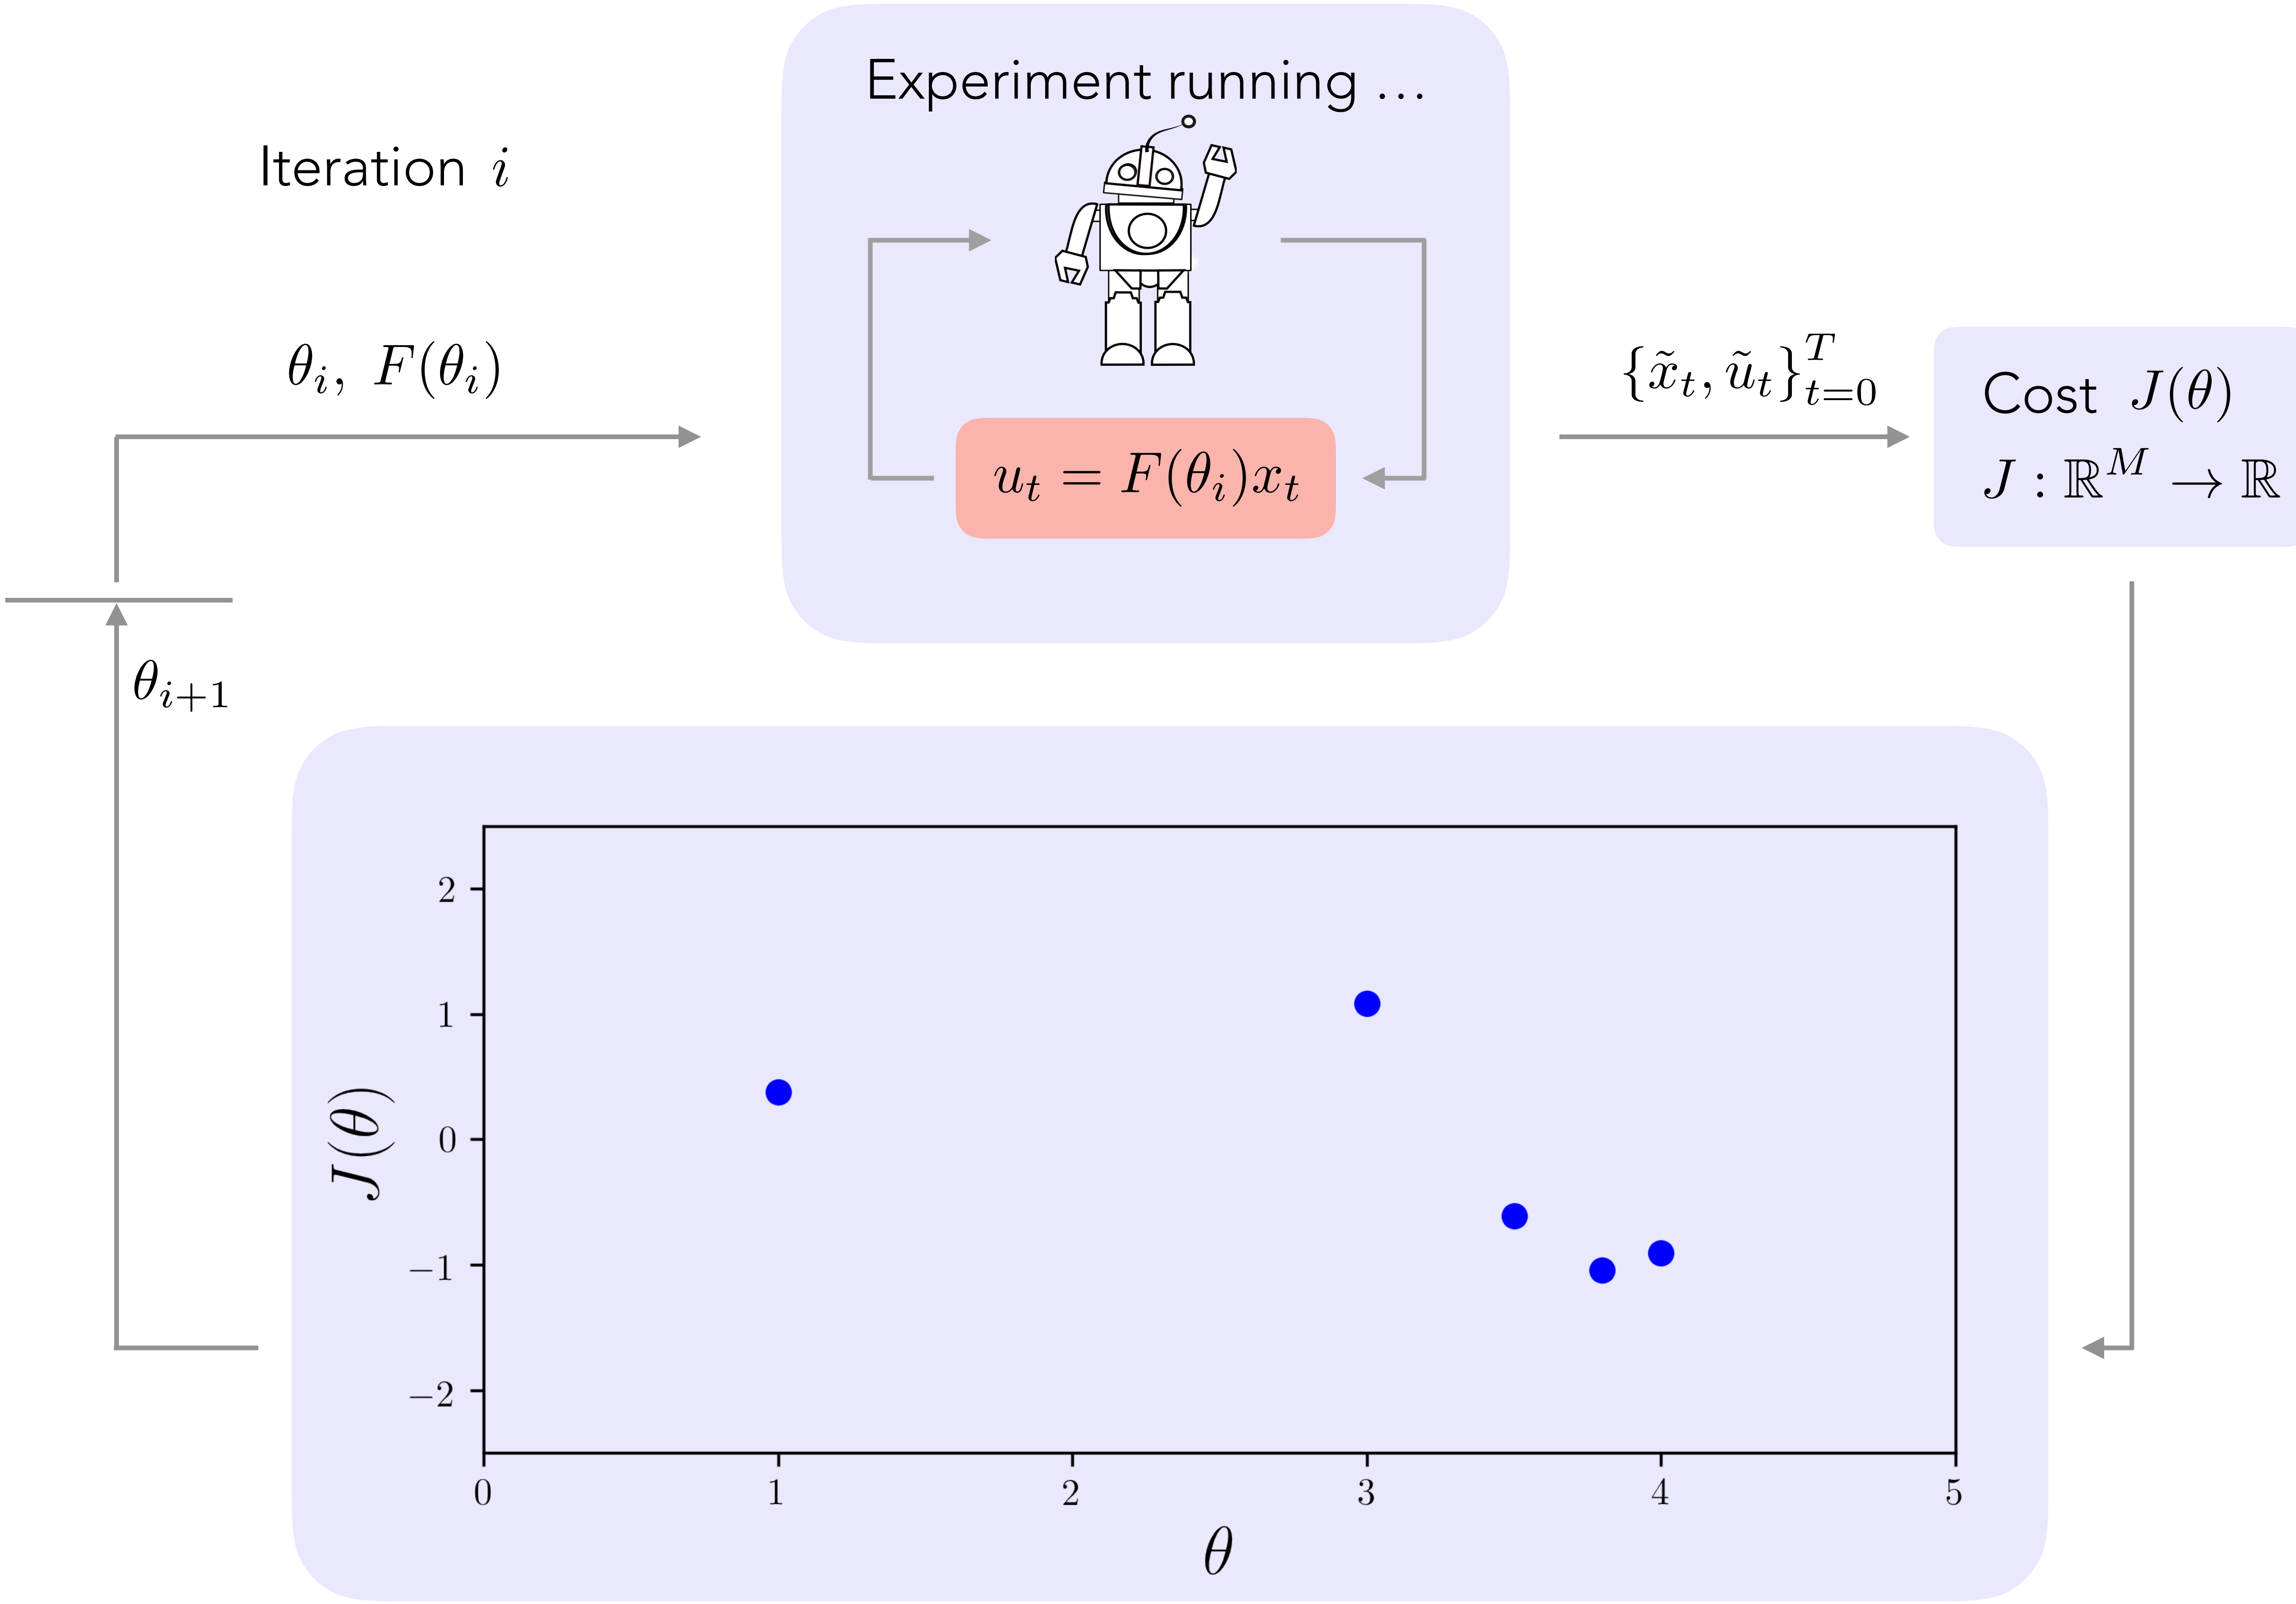
Experiment running …
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

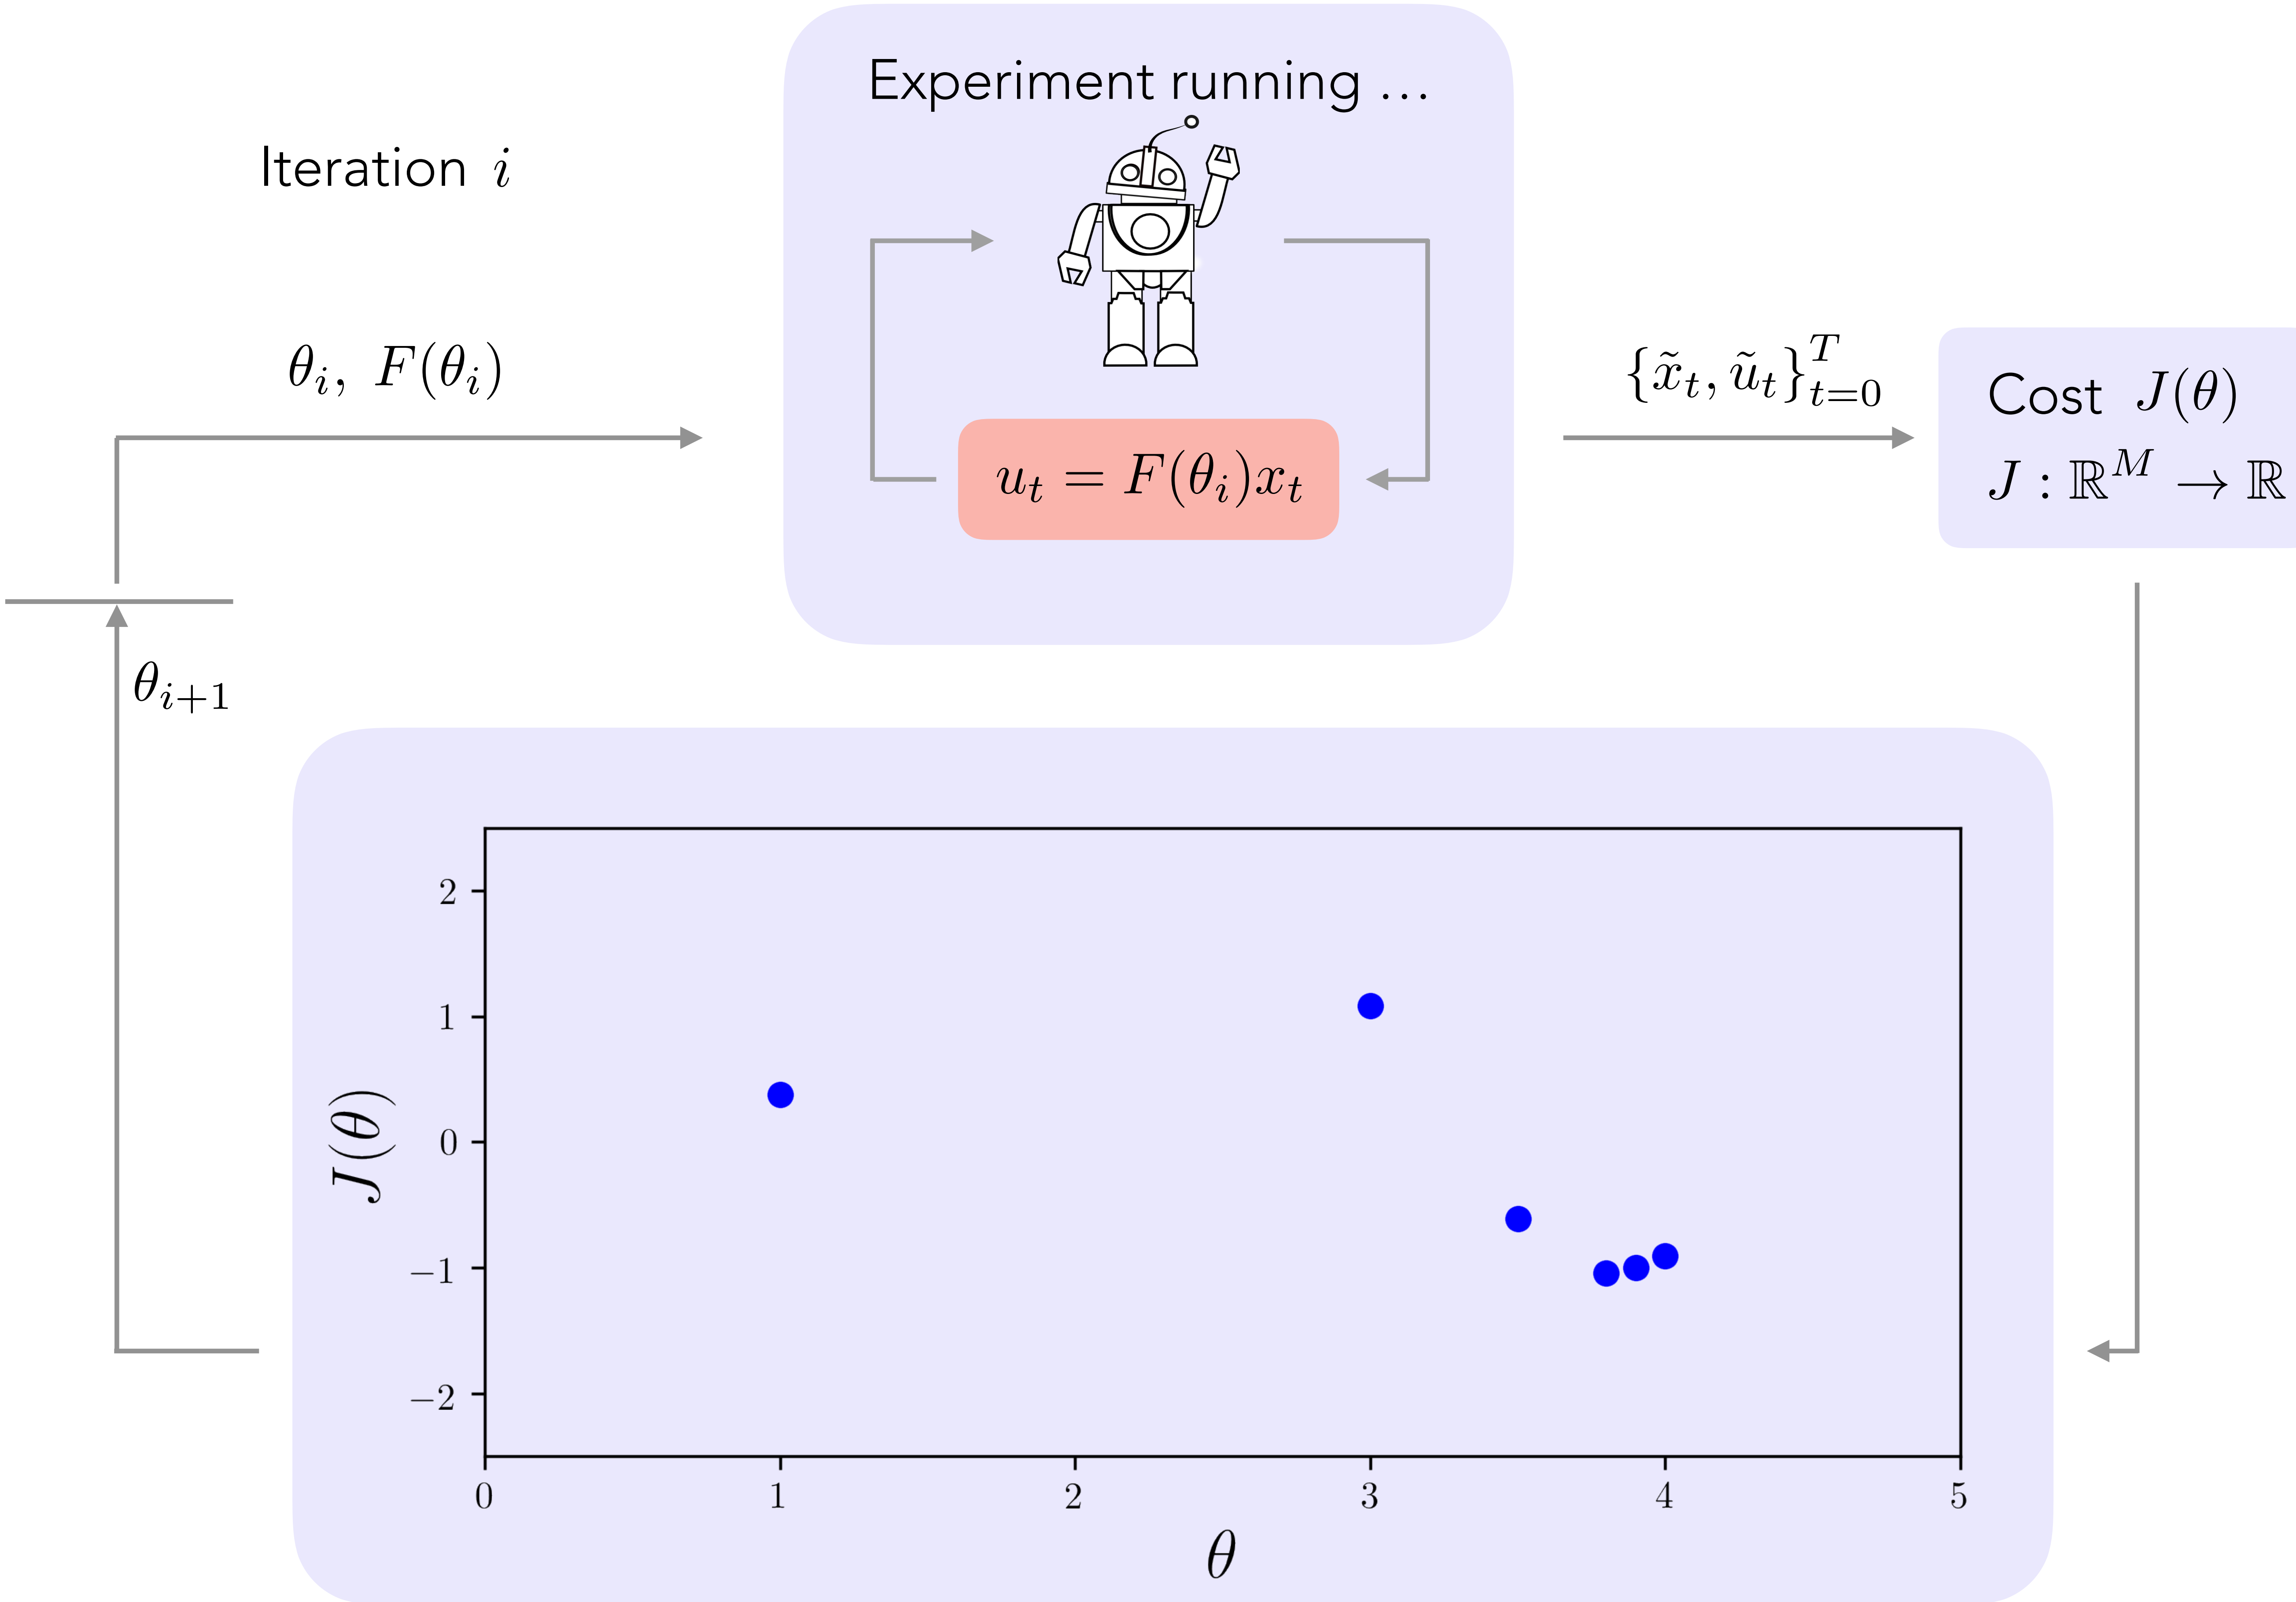
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

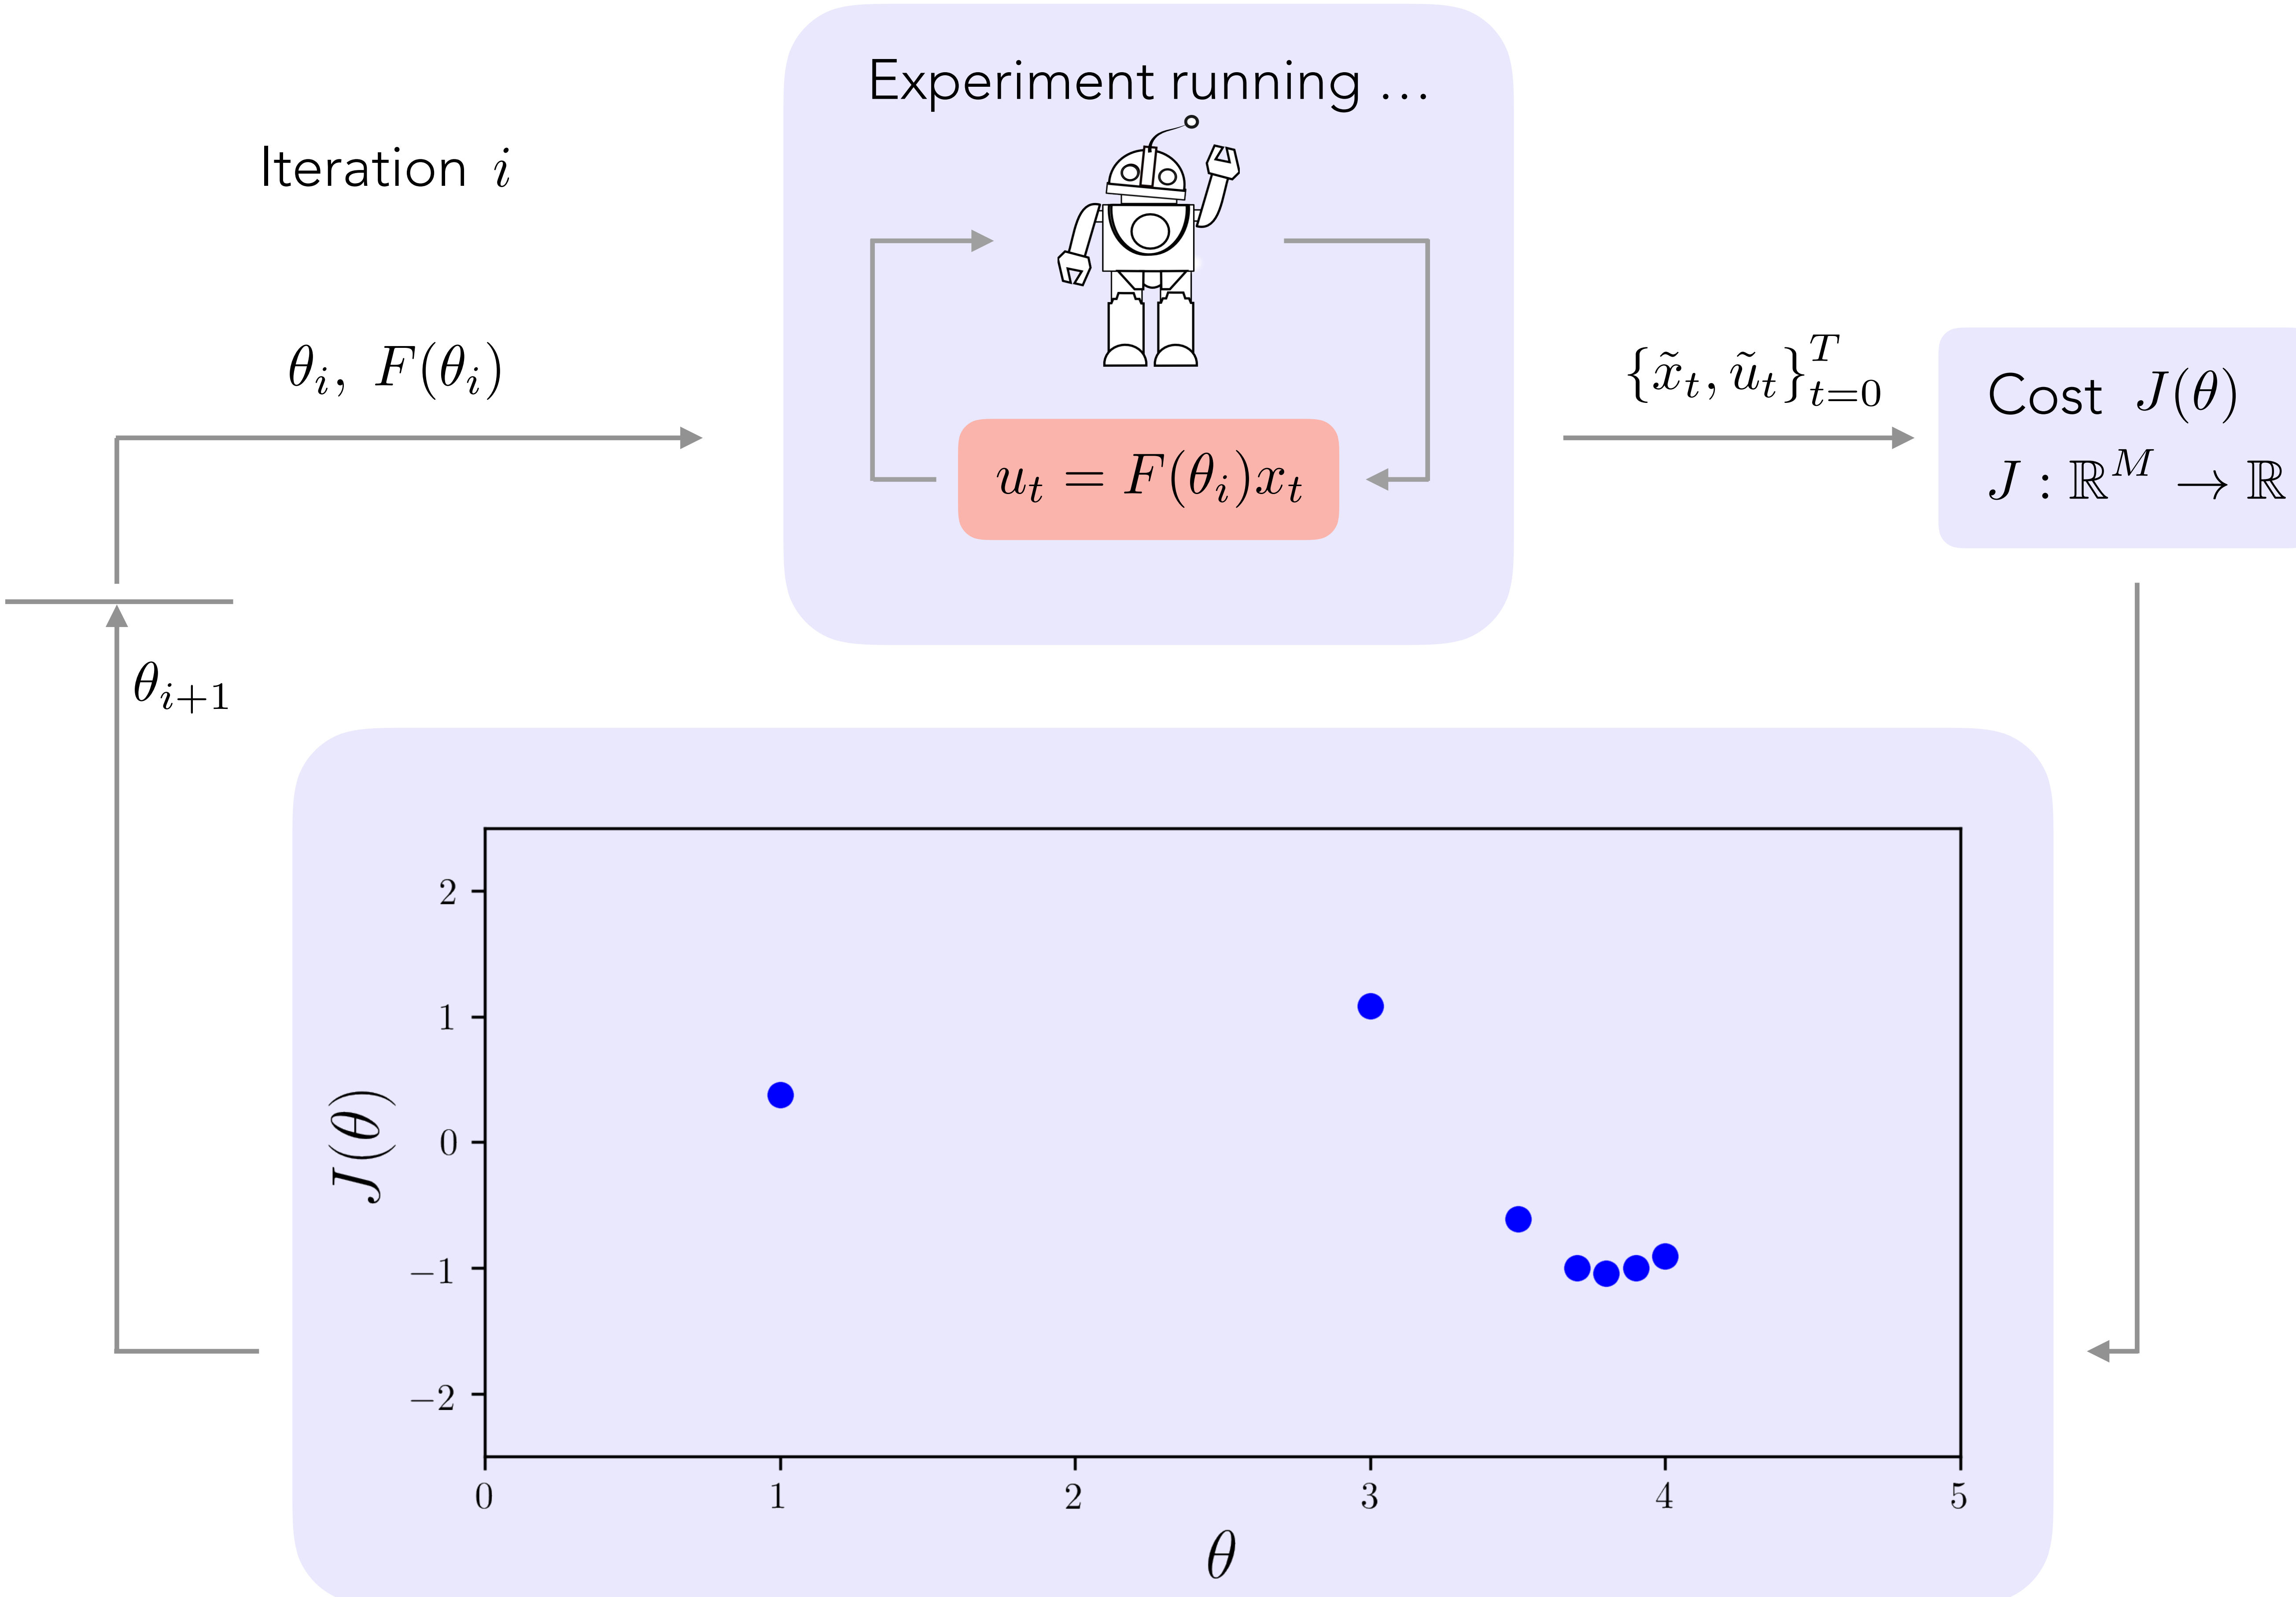
Iteration i
Experiment running ...
$\theta_i, \; F(\theta_i)$
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

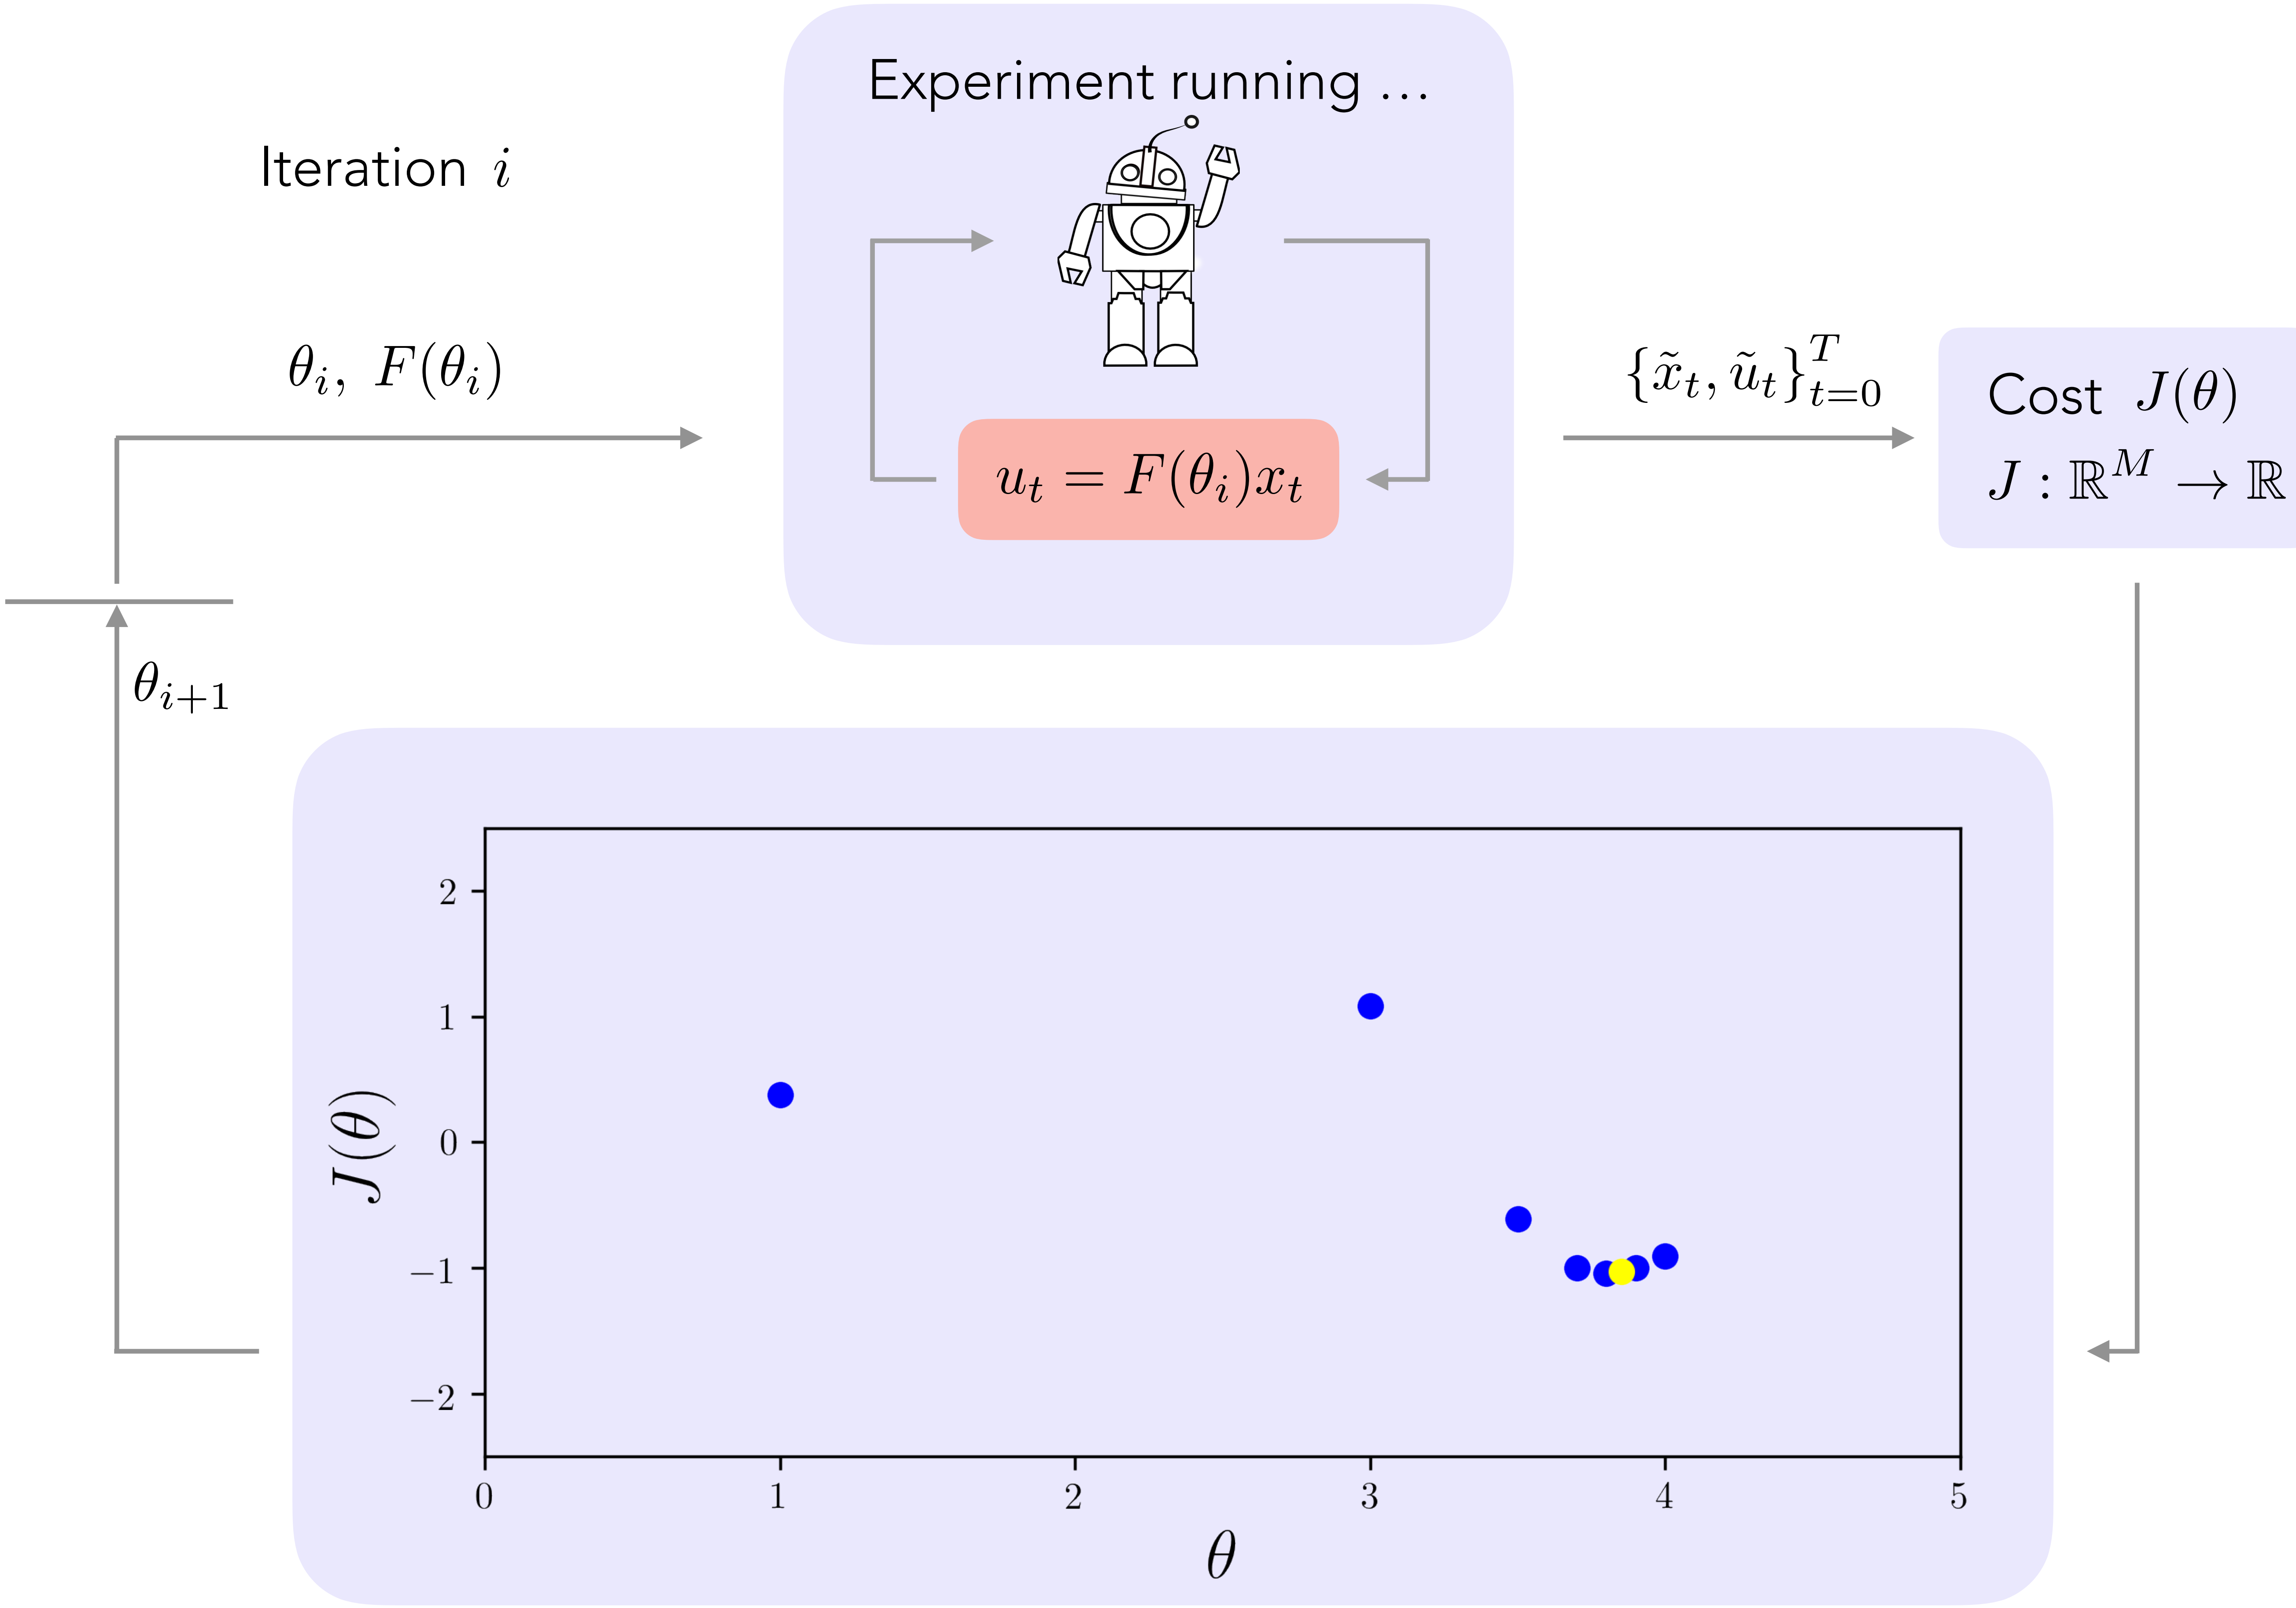
Iteration i
Experiment running …
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

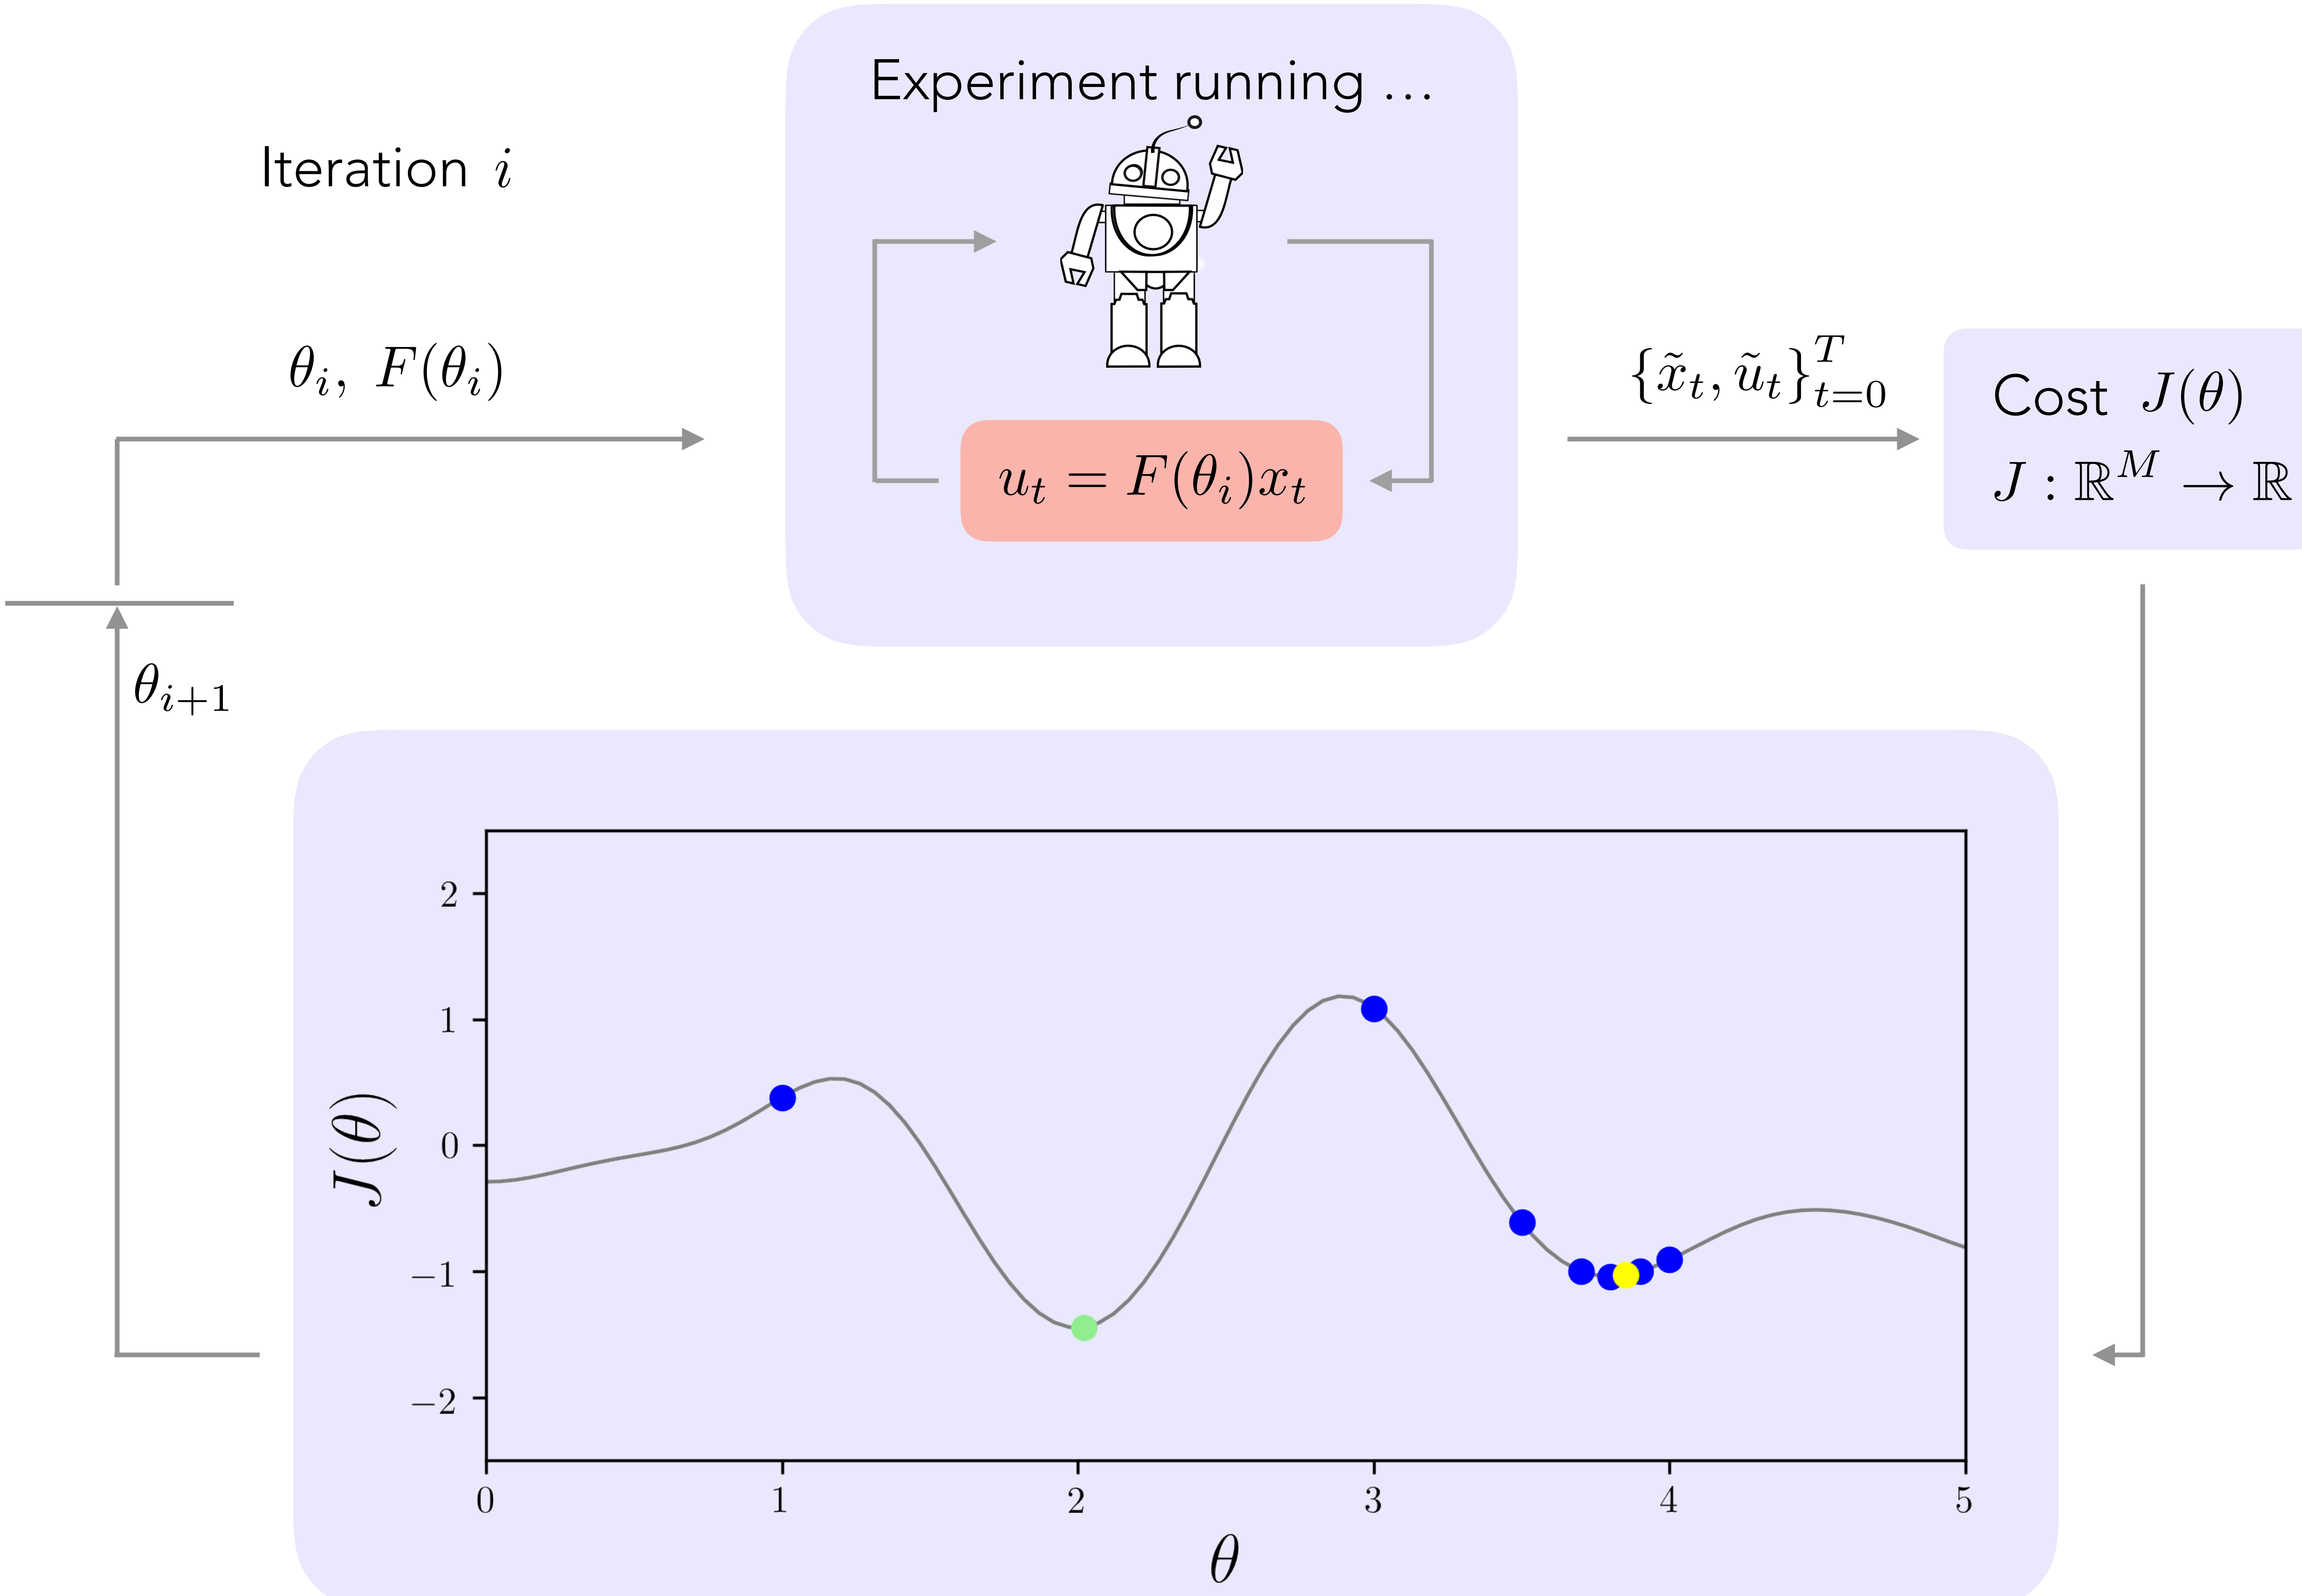
Experiment running …
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
$J(\theta)$
θ
2
1
0
−1
−2
0
1
2
3
4
5

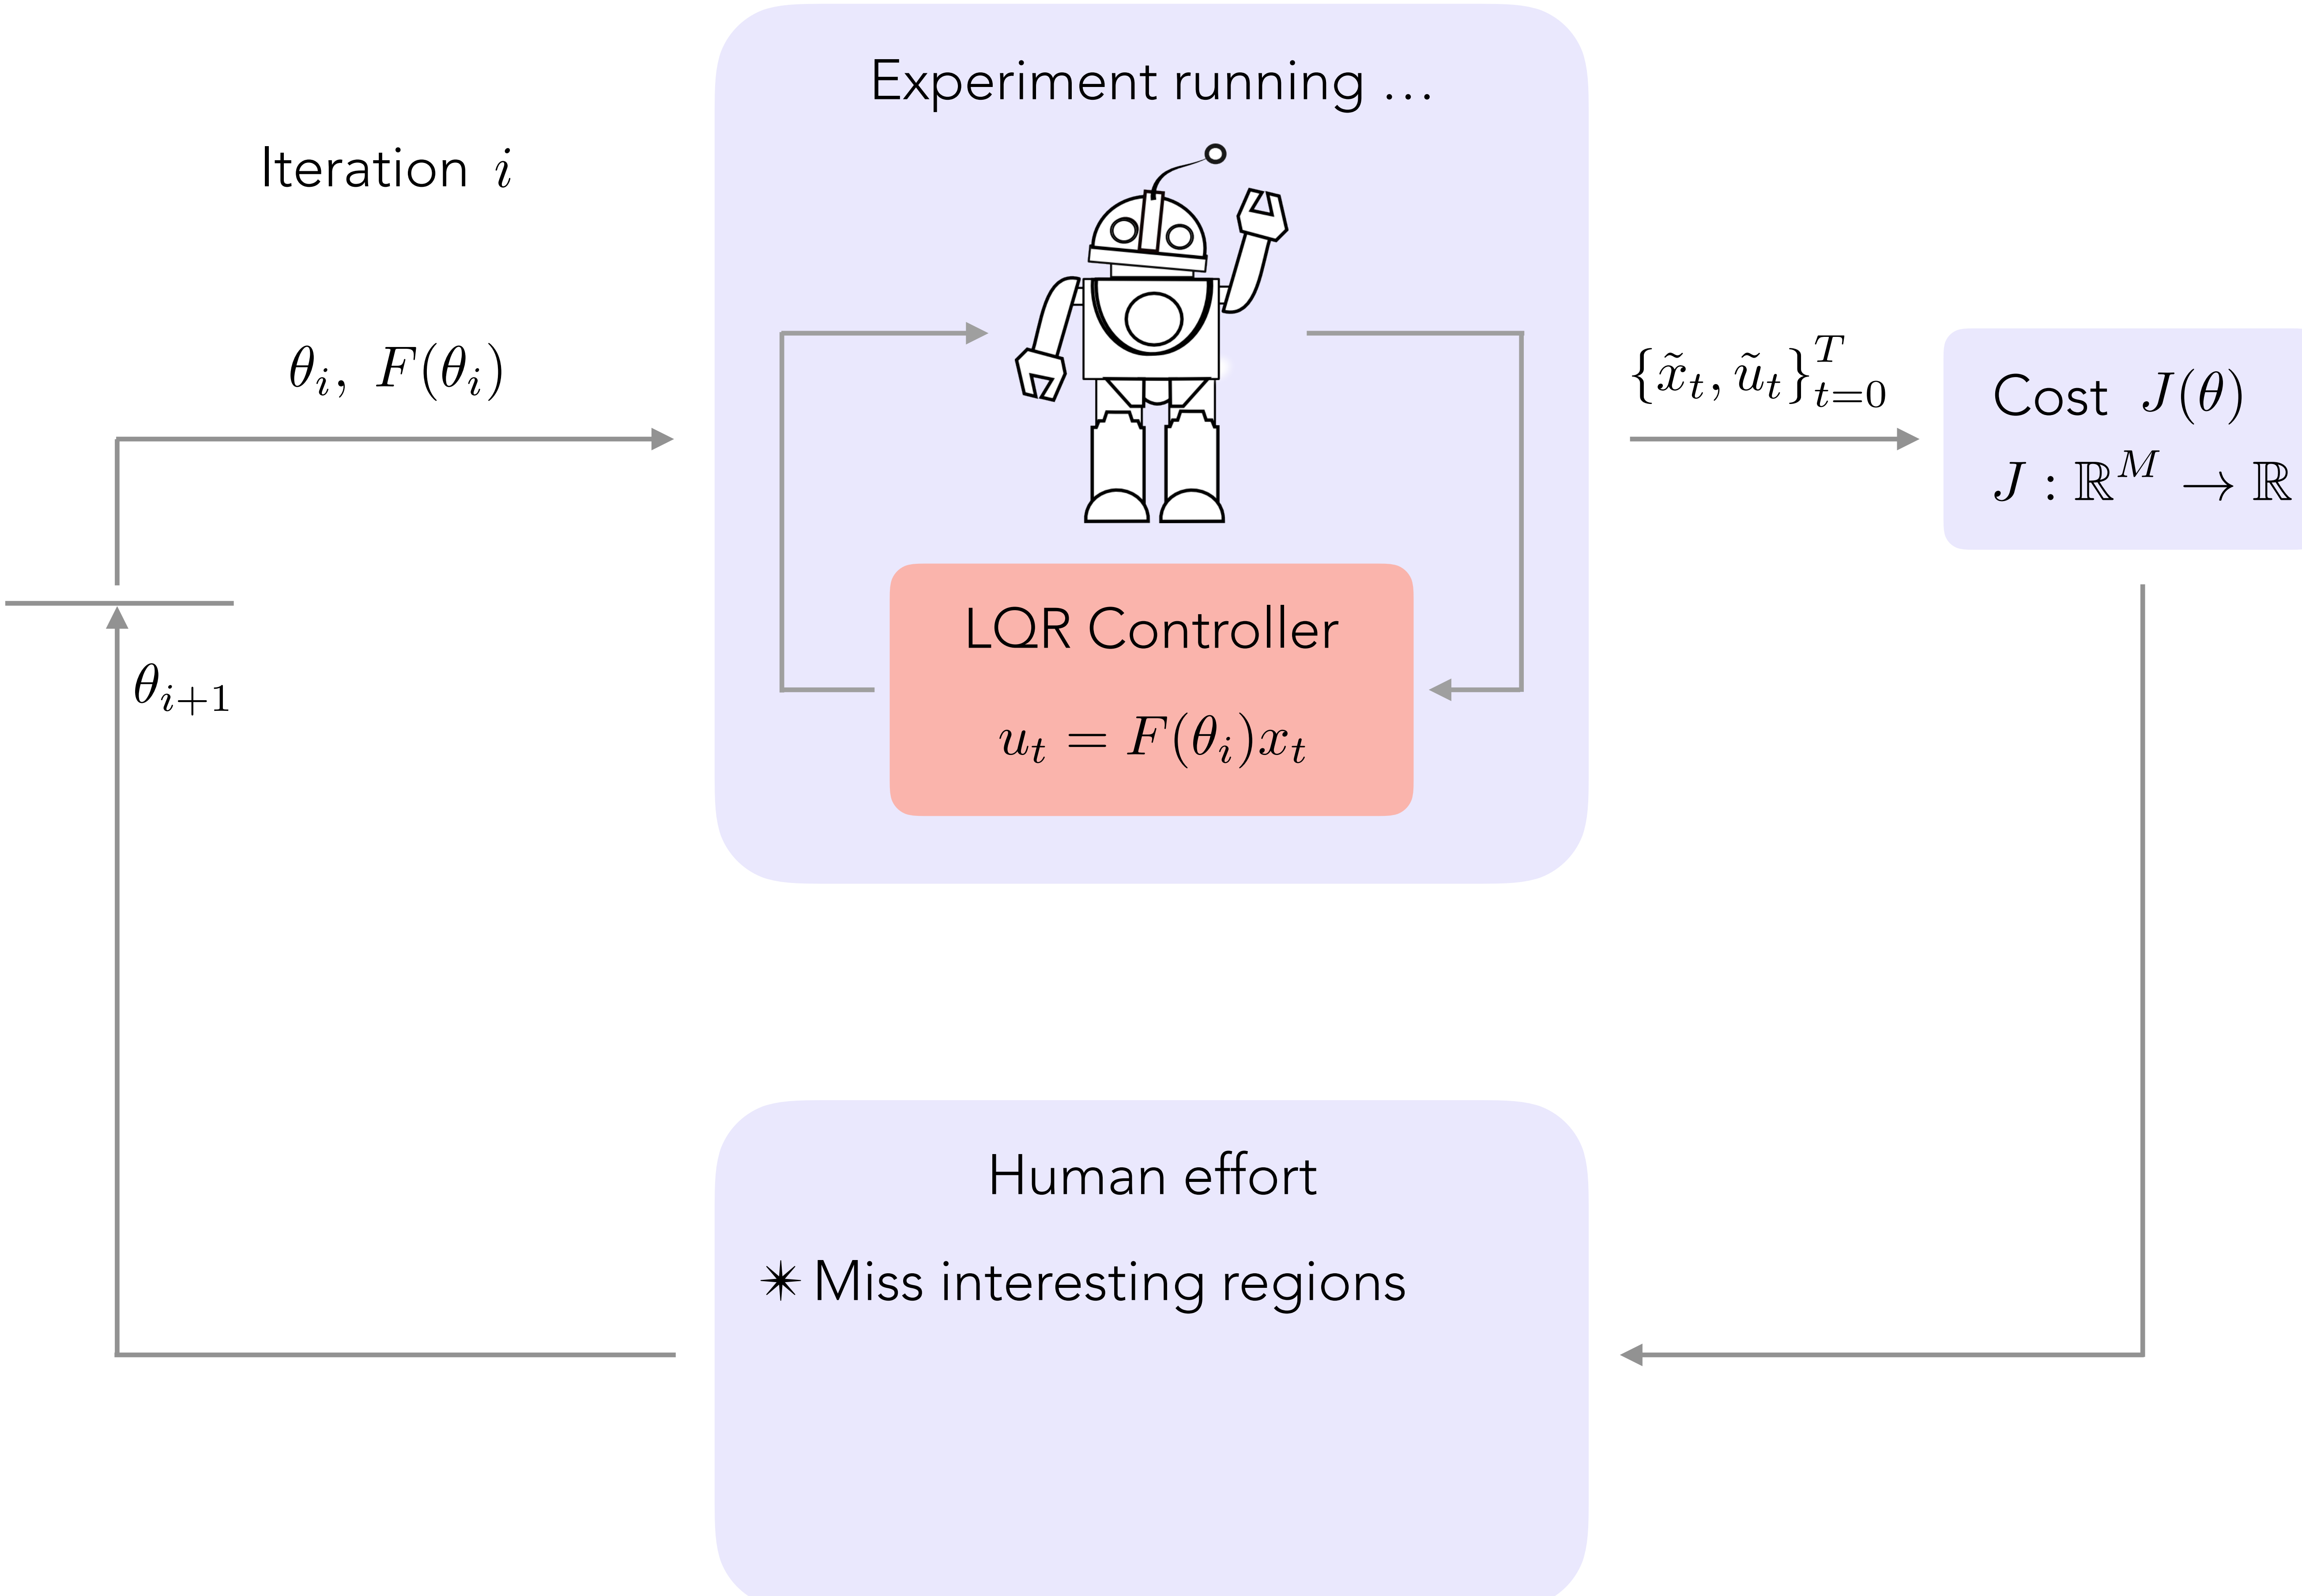
Iteration i
Experiment running ...
$\theta_i,\ F(\theta_i)$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J: \mathbb{R}^M \to \mathbb{R}$
LQR Controller
$u_t = F(\theta_i)x_t$
θ_{i+1}
Human effort
✳ Miss interesting regions

Iteration i

$\theta_i,\ F(\theta_i)$

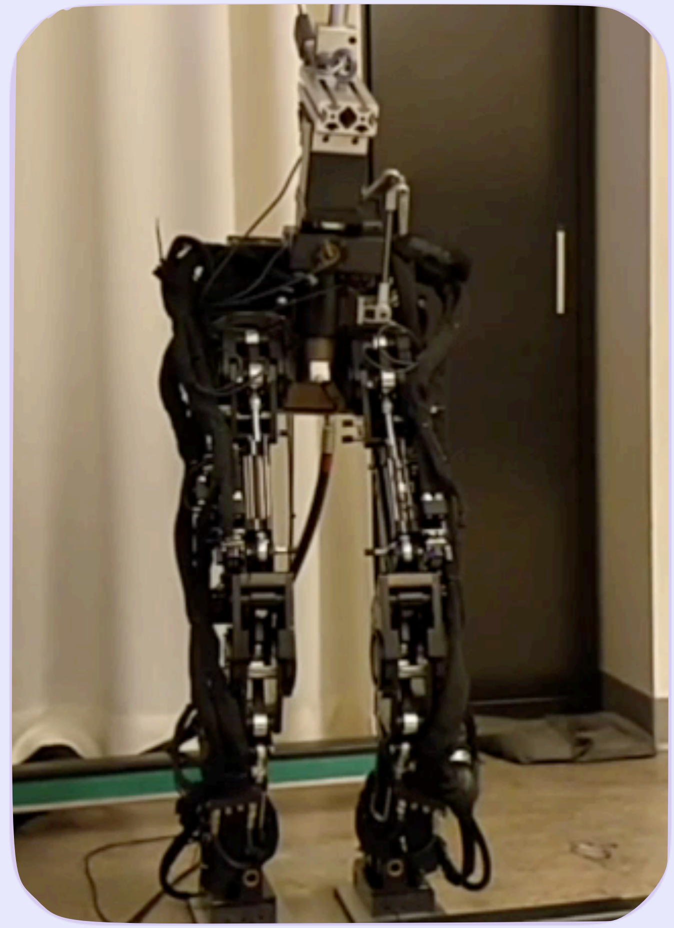

$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

θ_{i+1}

Human effort

* Miss interesting regions
* $\theta \in \mathbb{R}^M$, with M large

Iteration i

$\theta_i,\ F(\theta_i)$

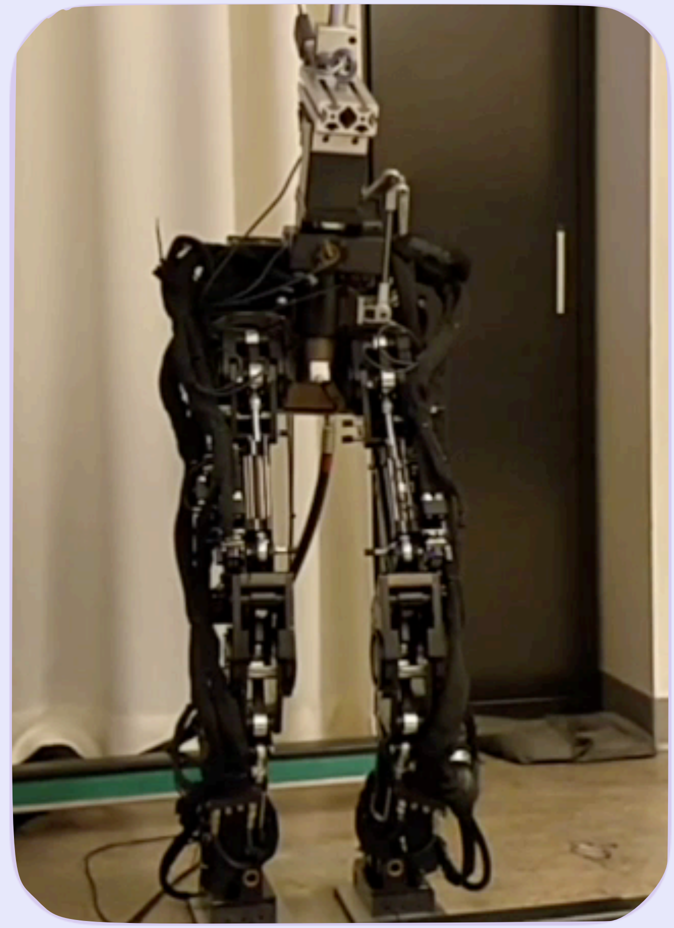

$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

θ_{i+1}

Human effort

✴ Miss interesting regions

✴ $\theta \in \mathbb{R}^M$, with M large

✴ Wait for experiments

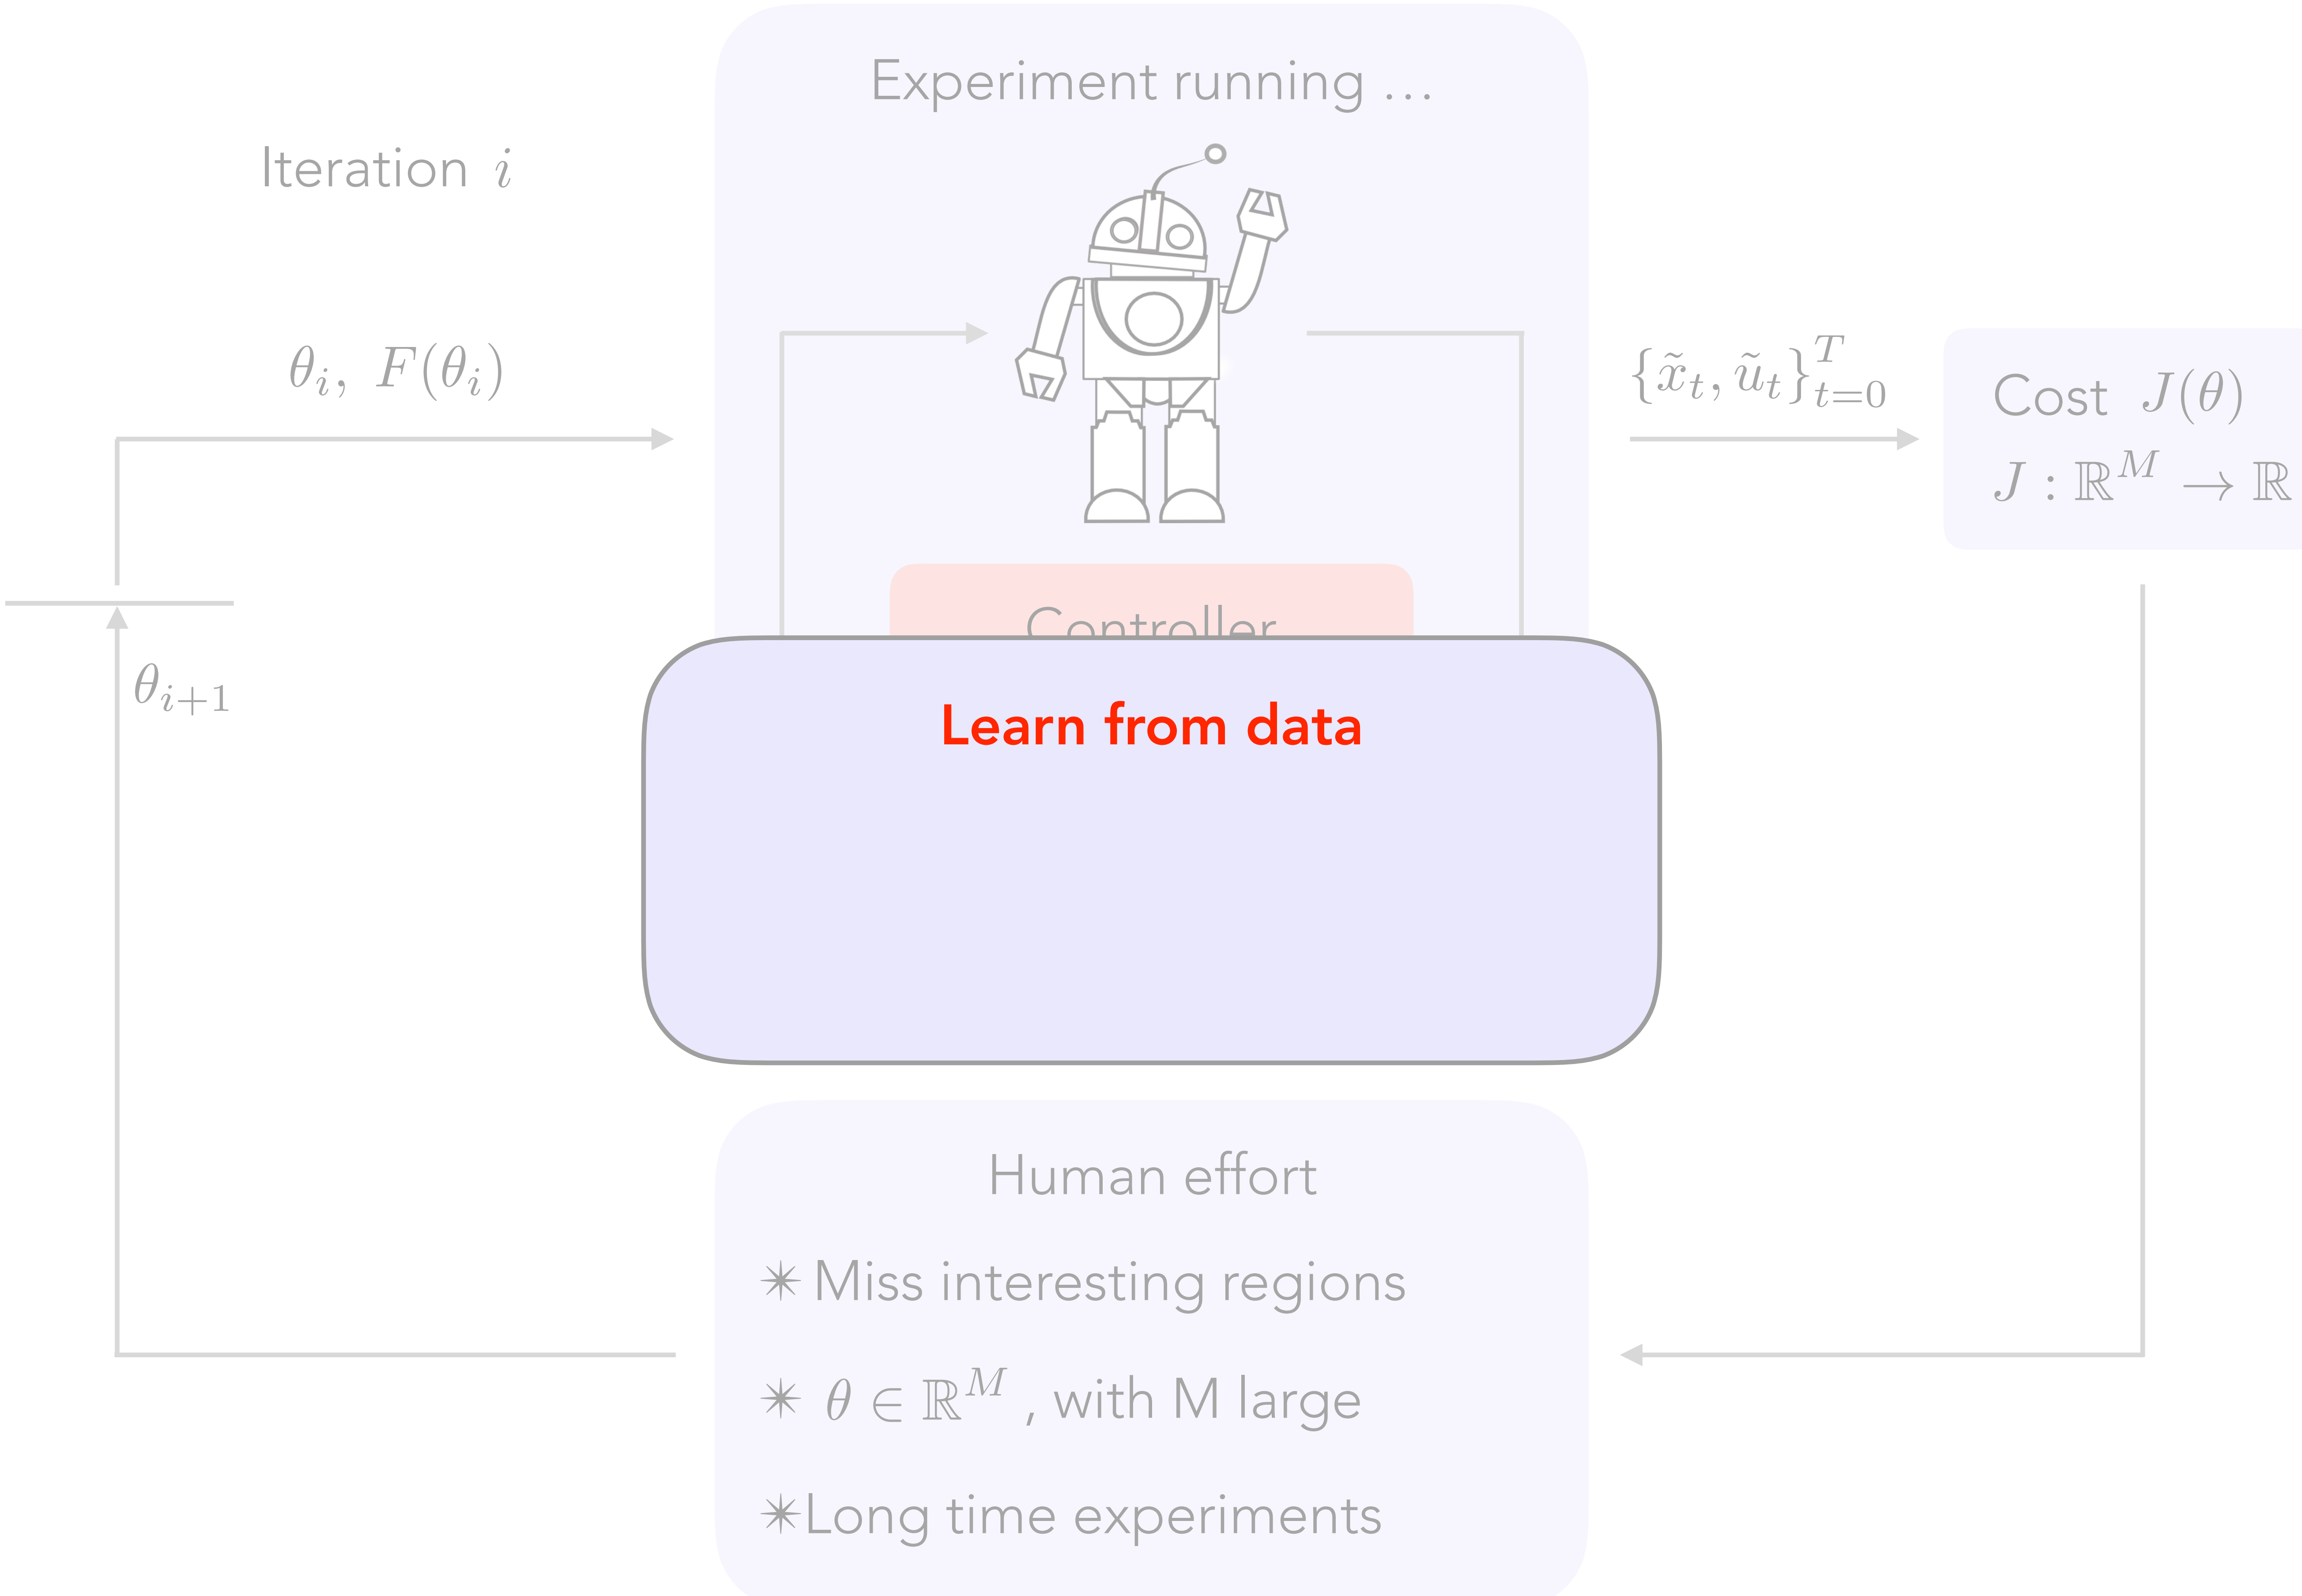
Experiment running …
Iteration i
$\theta_i,\ F(\theta_i)$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \rightarrow \mathbb{R}$
Controller
θ_{i+1}
Learn from data
Human effort
✳ Miss interesting regions
✳ $\theta \in \mathbb{R}^M$, with M large
✳ Long time experiments

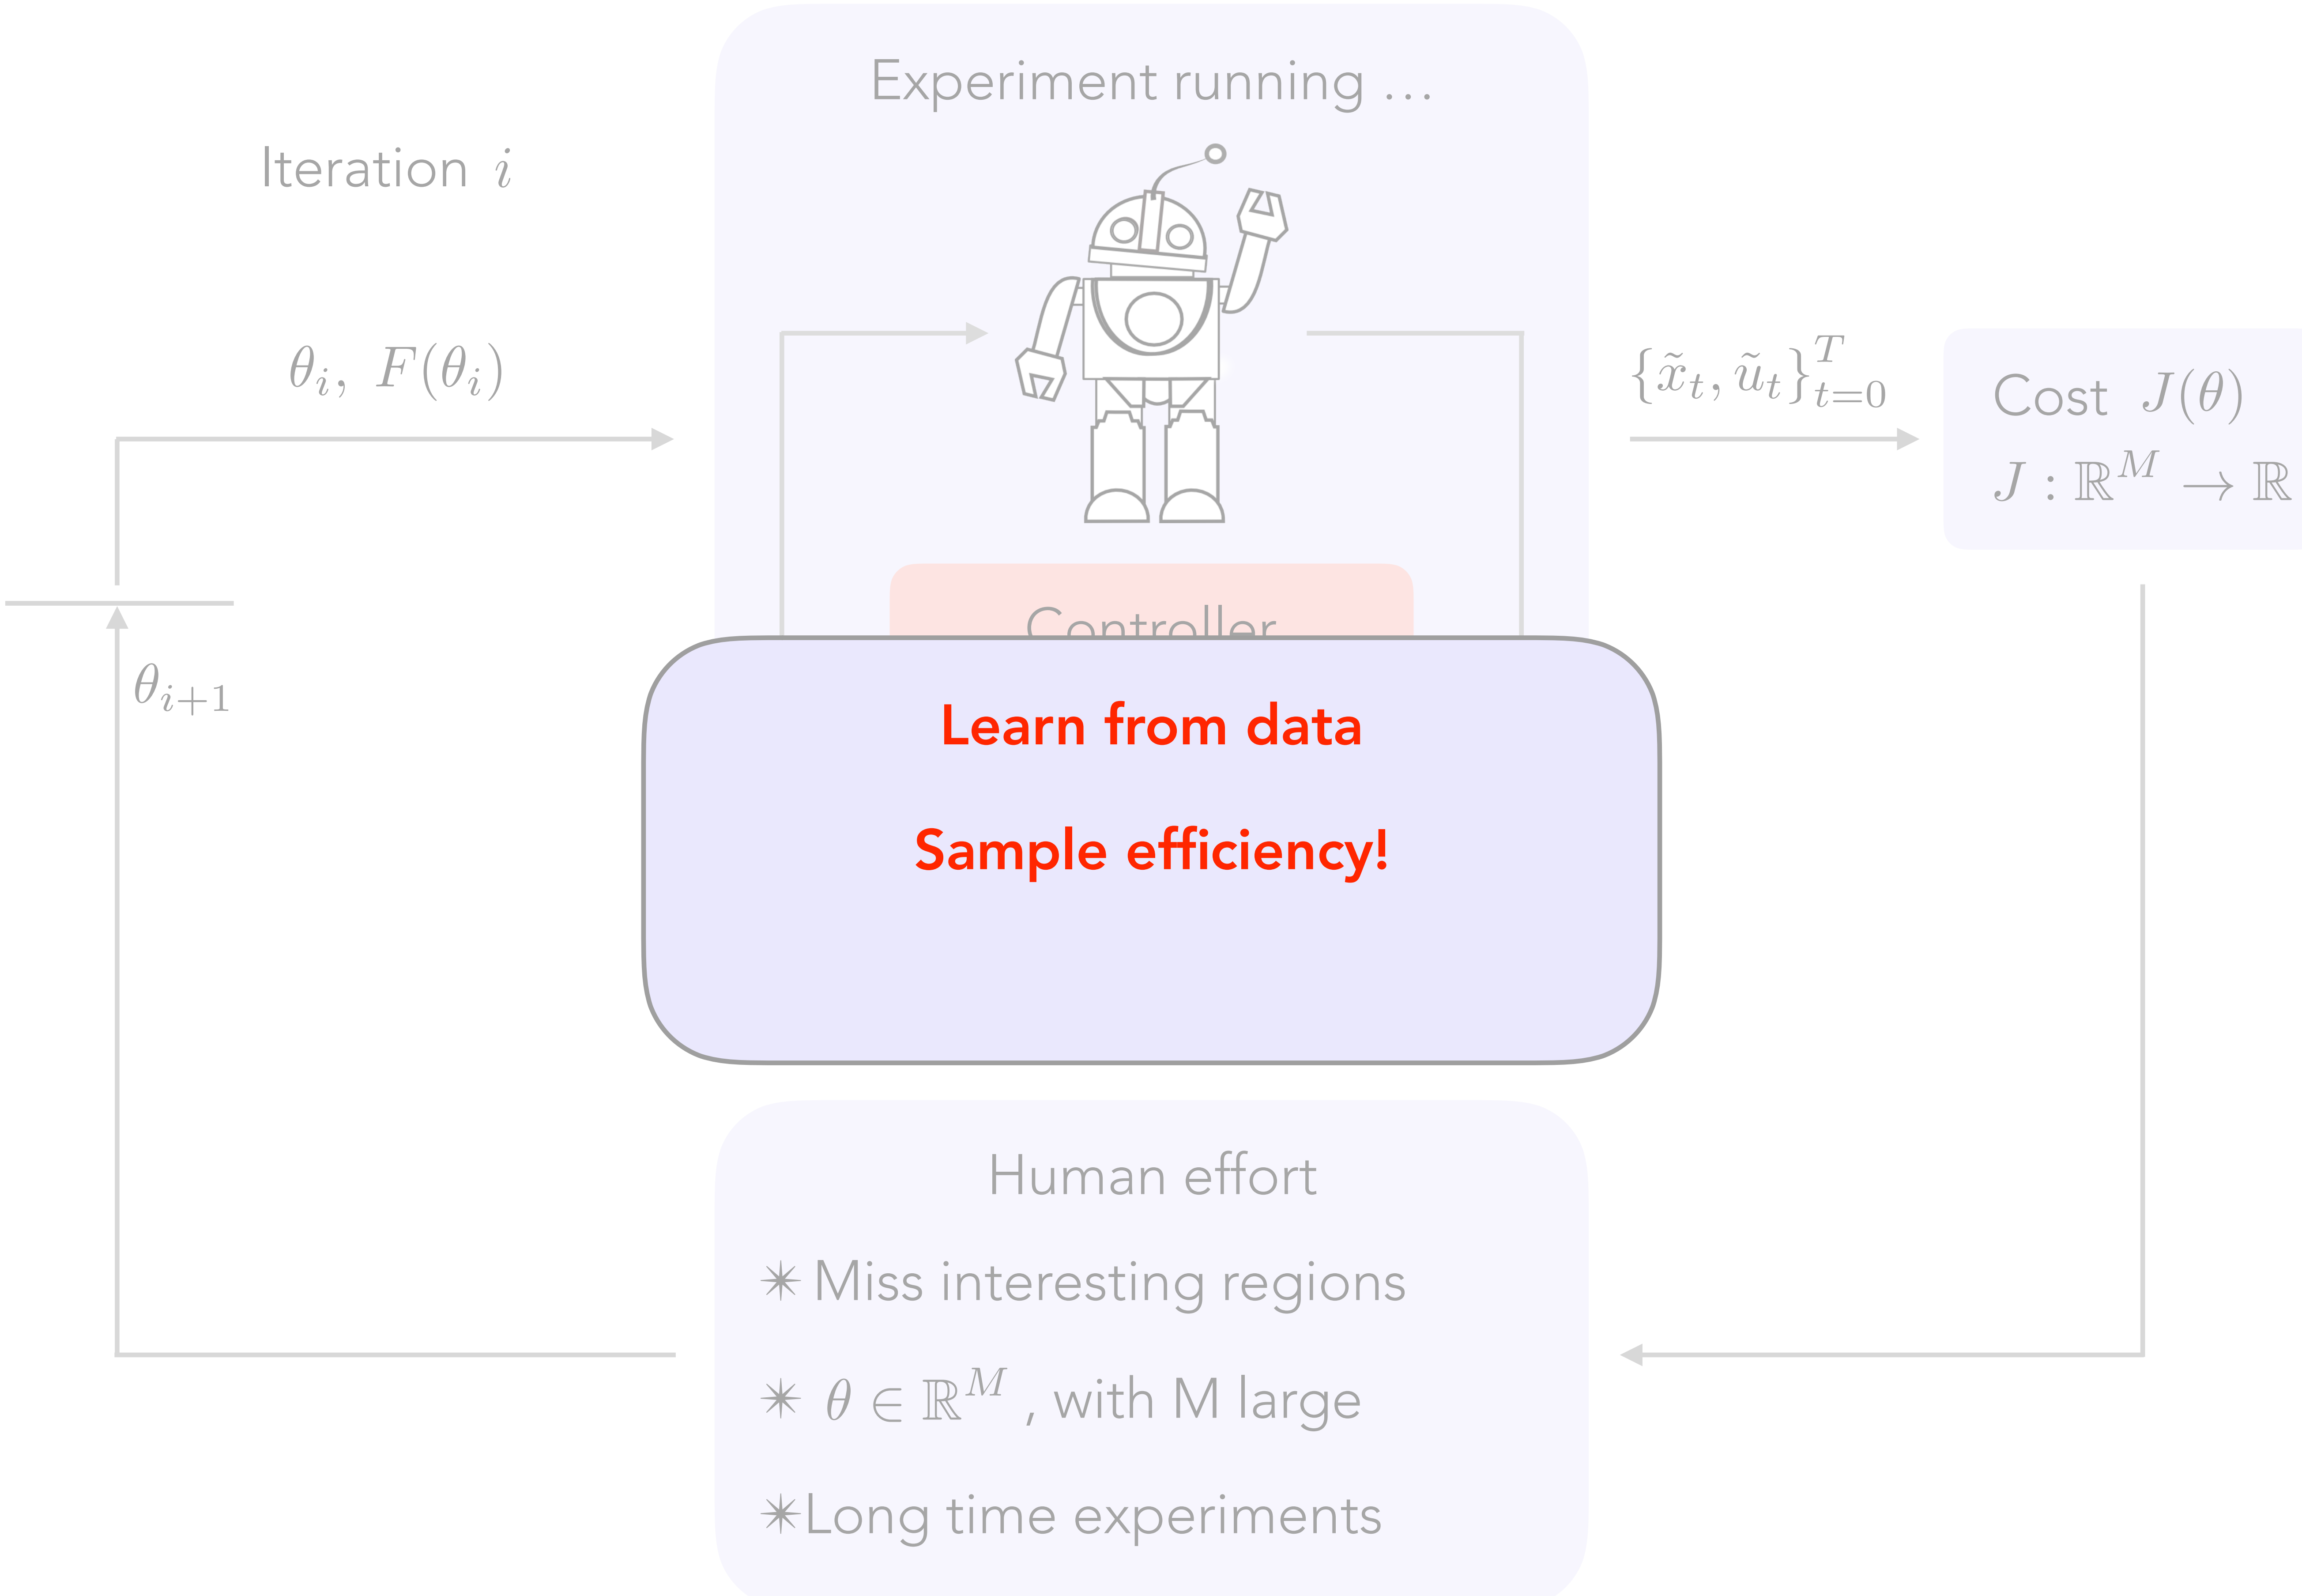
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \rightarrow \mathbb{R}$
Controller
θ_{i+1}
Learn from data
Sample efficiency!
Human effort
✳ Miss interesting regions
✳ $\theta \in \mathbb{R}^M$, with M large
✳Long time experiments

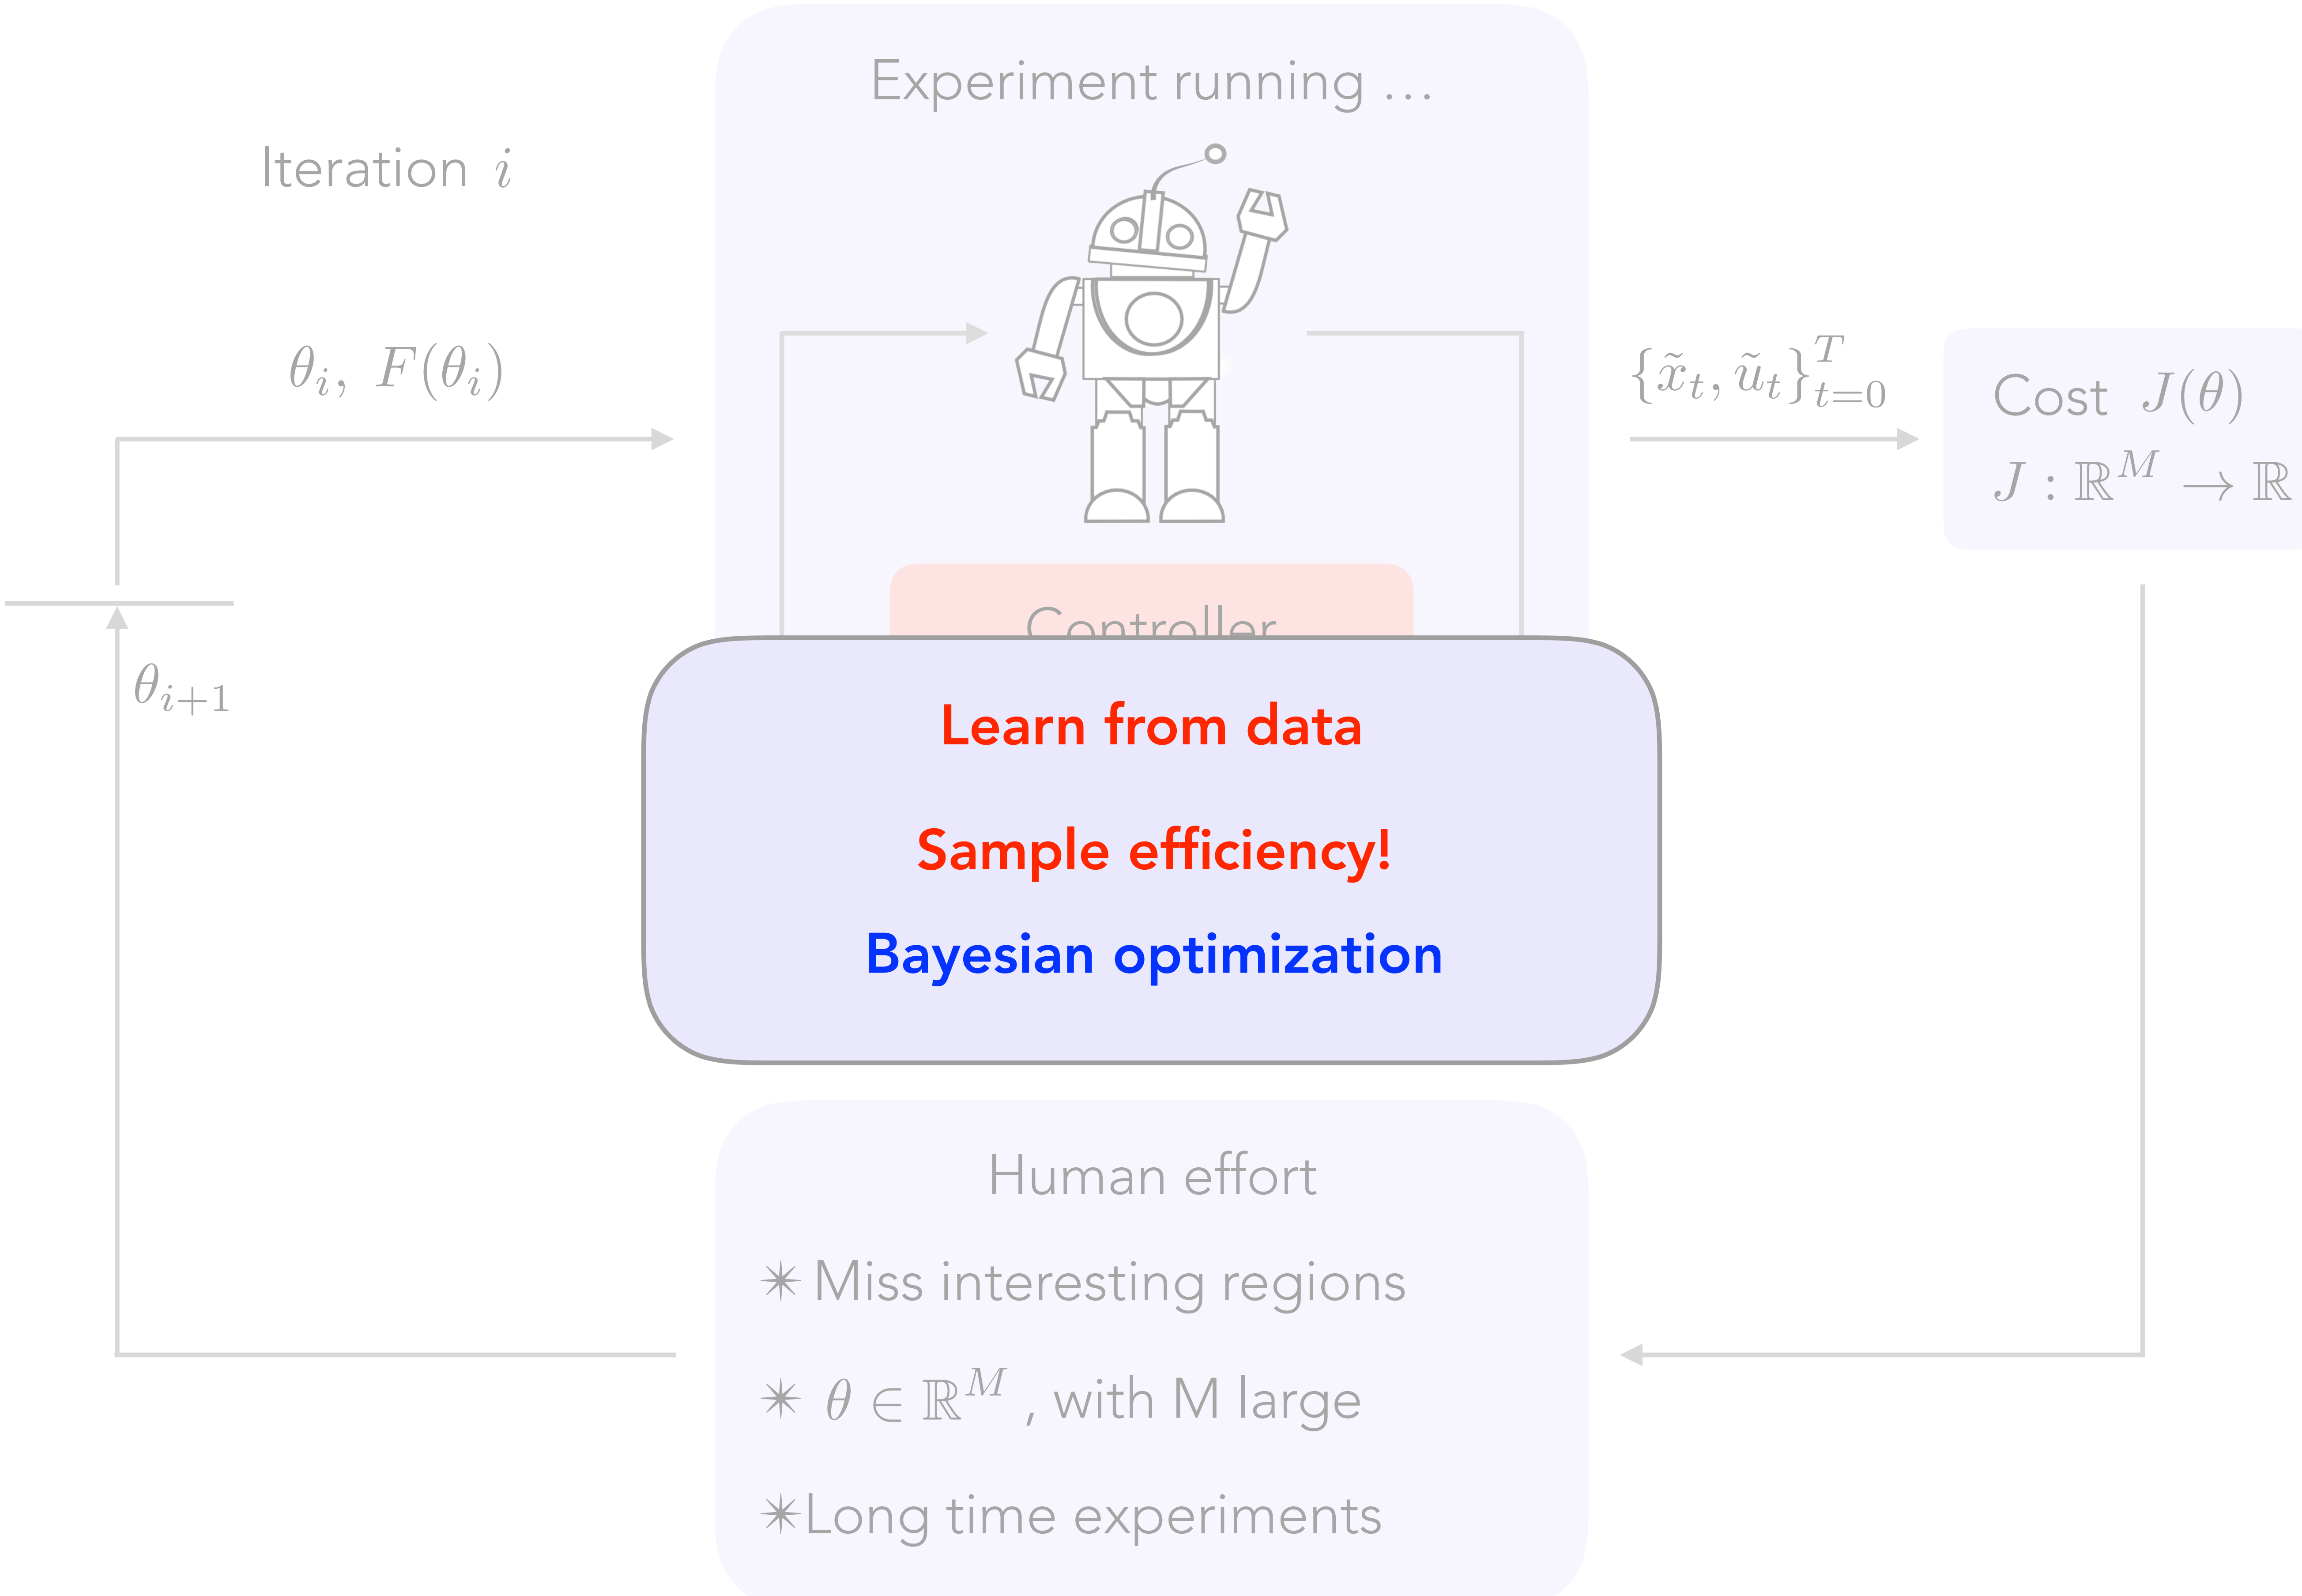
Iteration i
Experiment running …
$\theta_i, \; F(\theta_i)$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
Controller
θ_{i+1}
Learn from data
Sample efficiency!
Bayesian optimization
Human effort
✳ Miss interesting regions
✳ $\theta \in \mathbb{R}^M$, with M large
✳Long time experiments

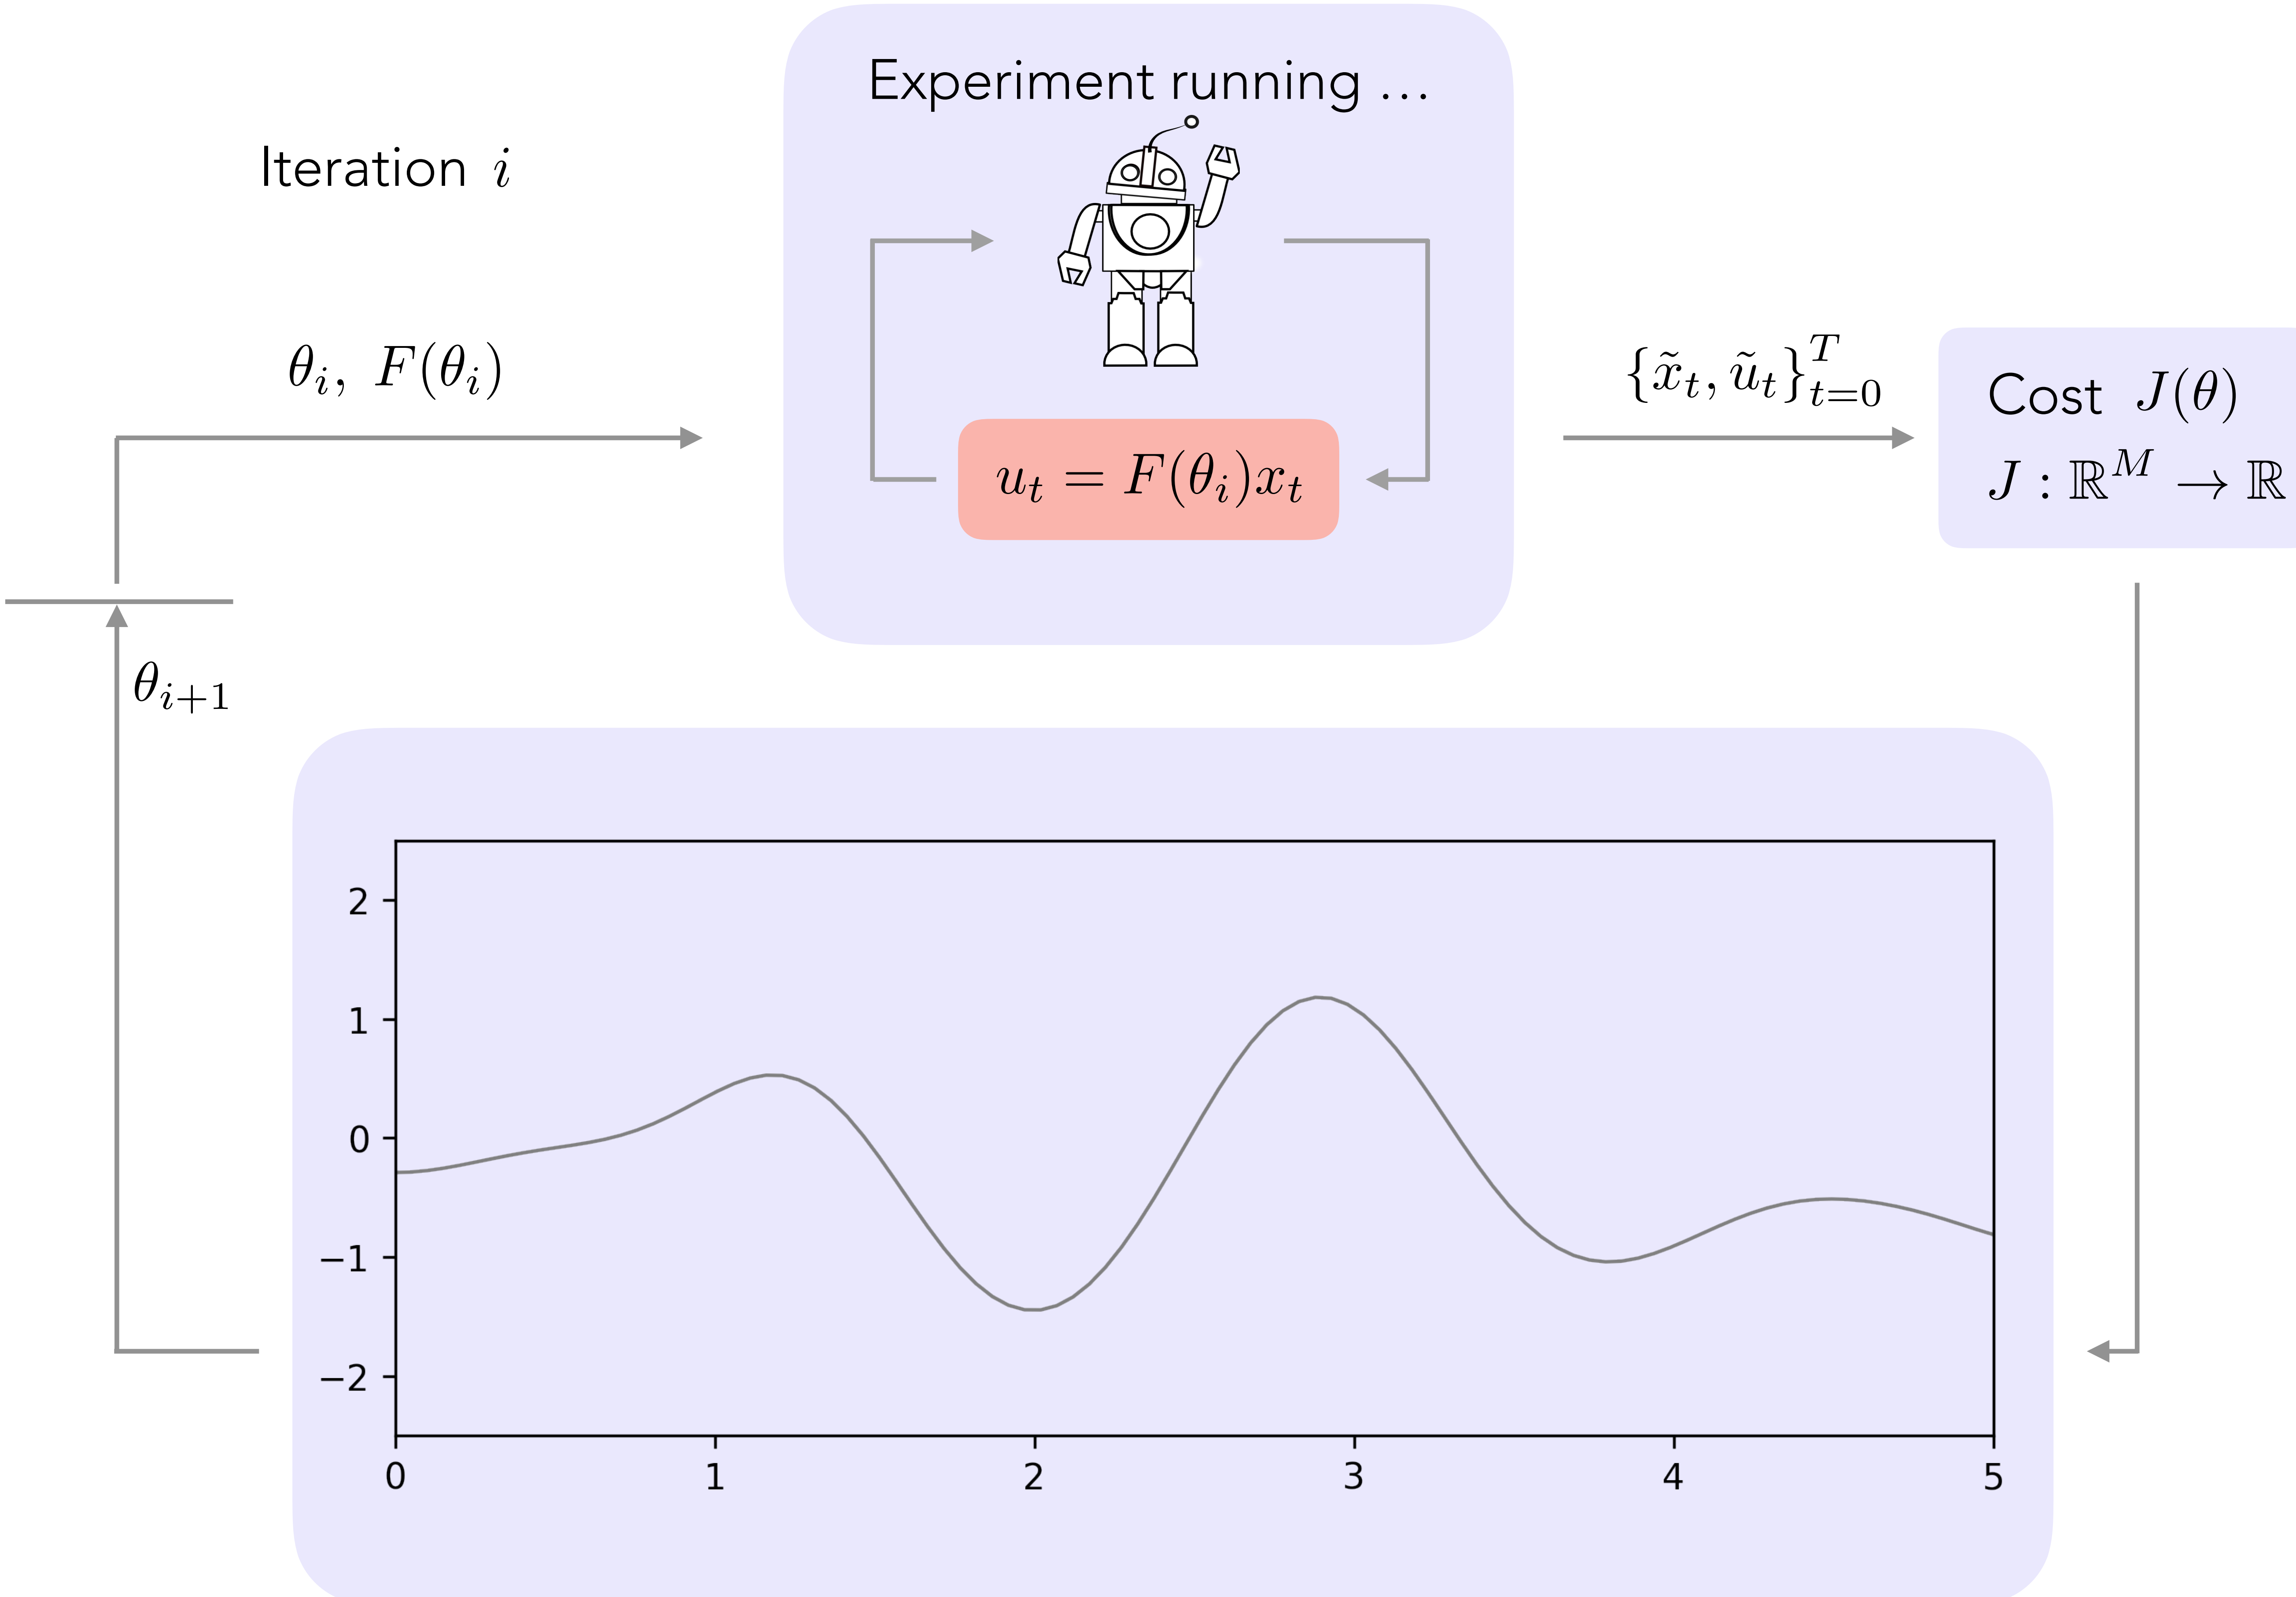
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
2
1
0
−1
−2
0
1
2
3
4
5

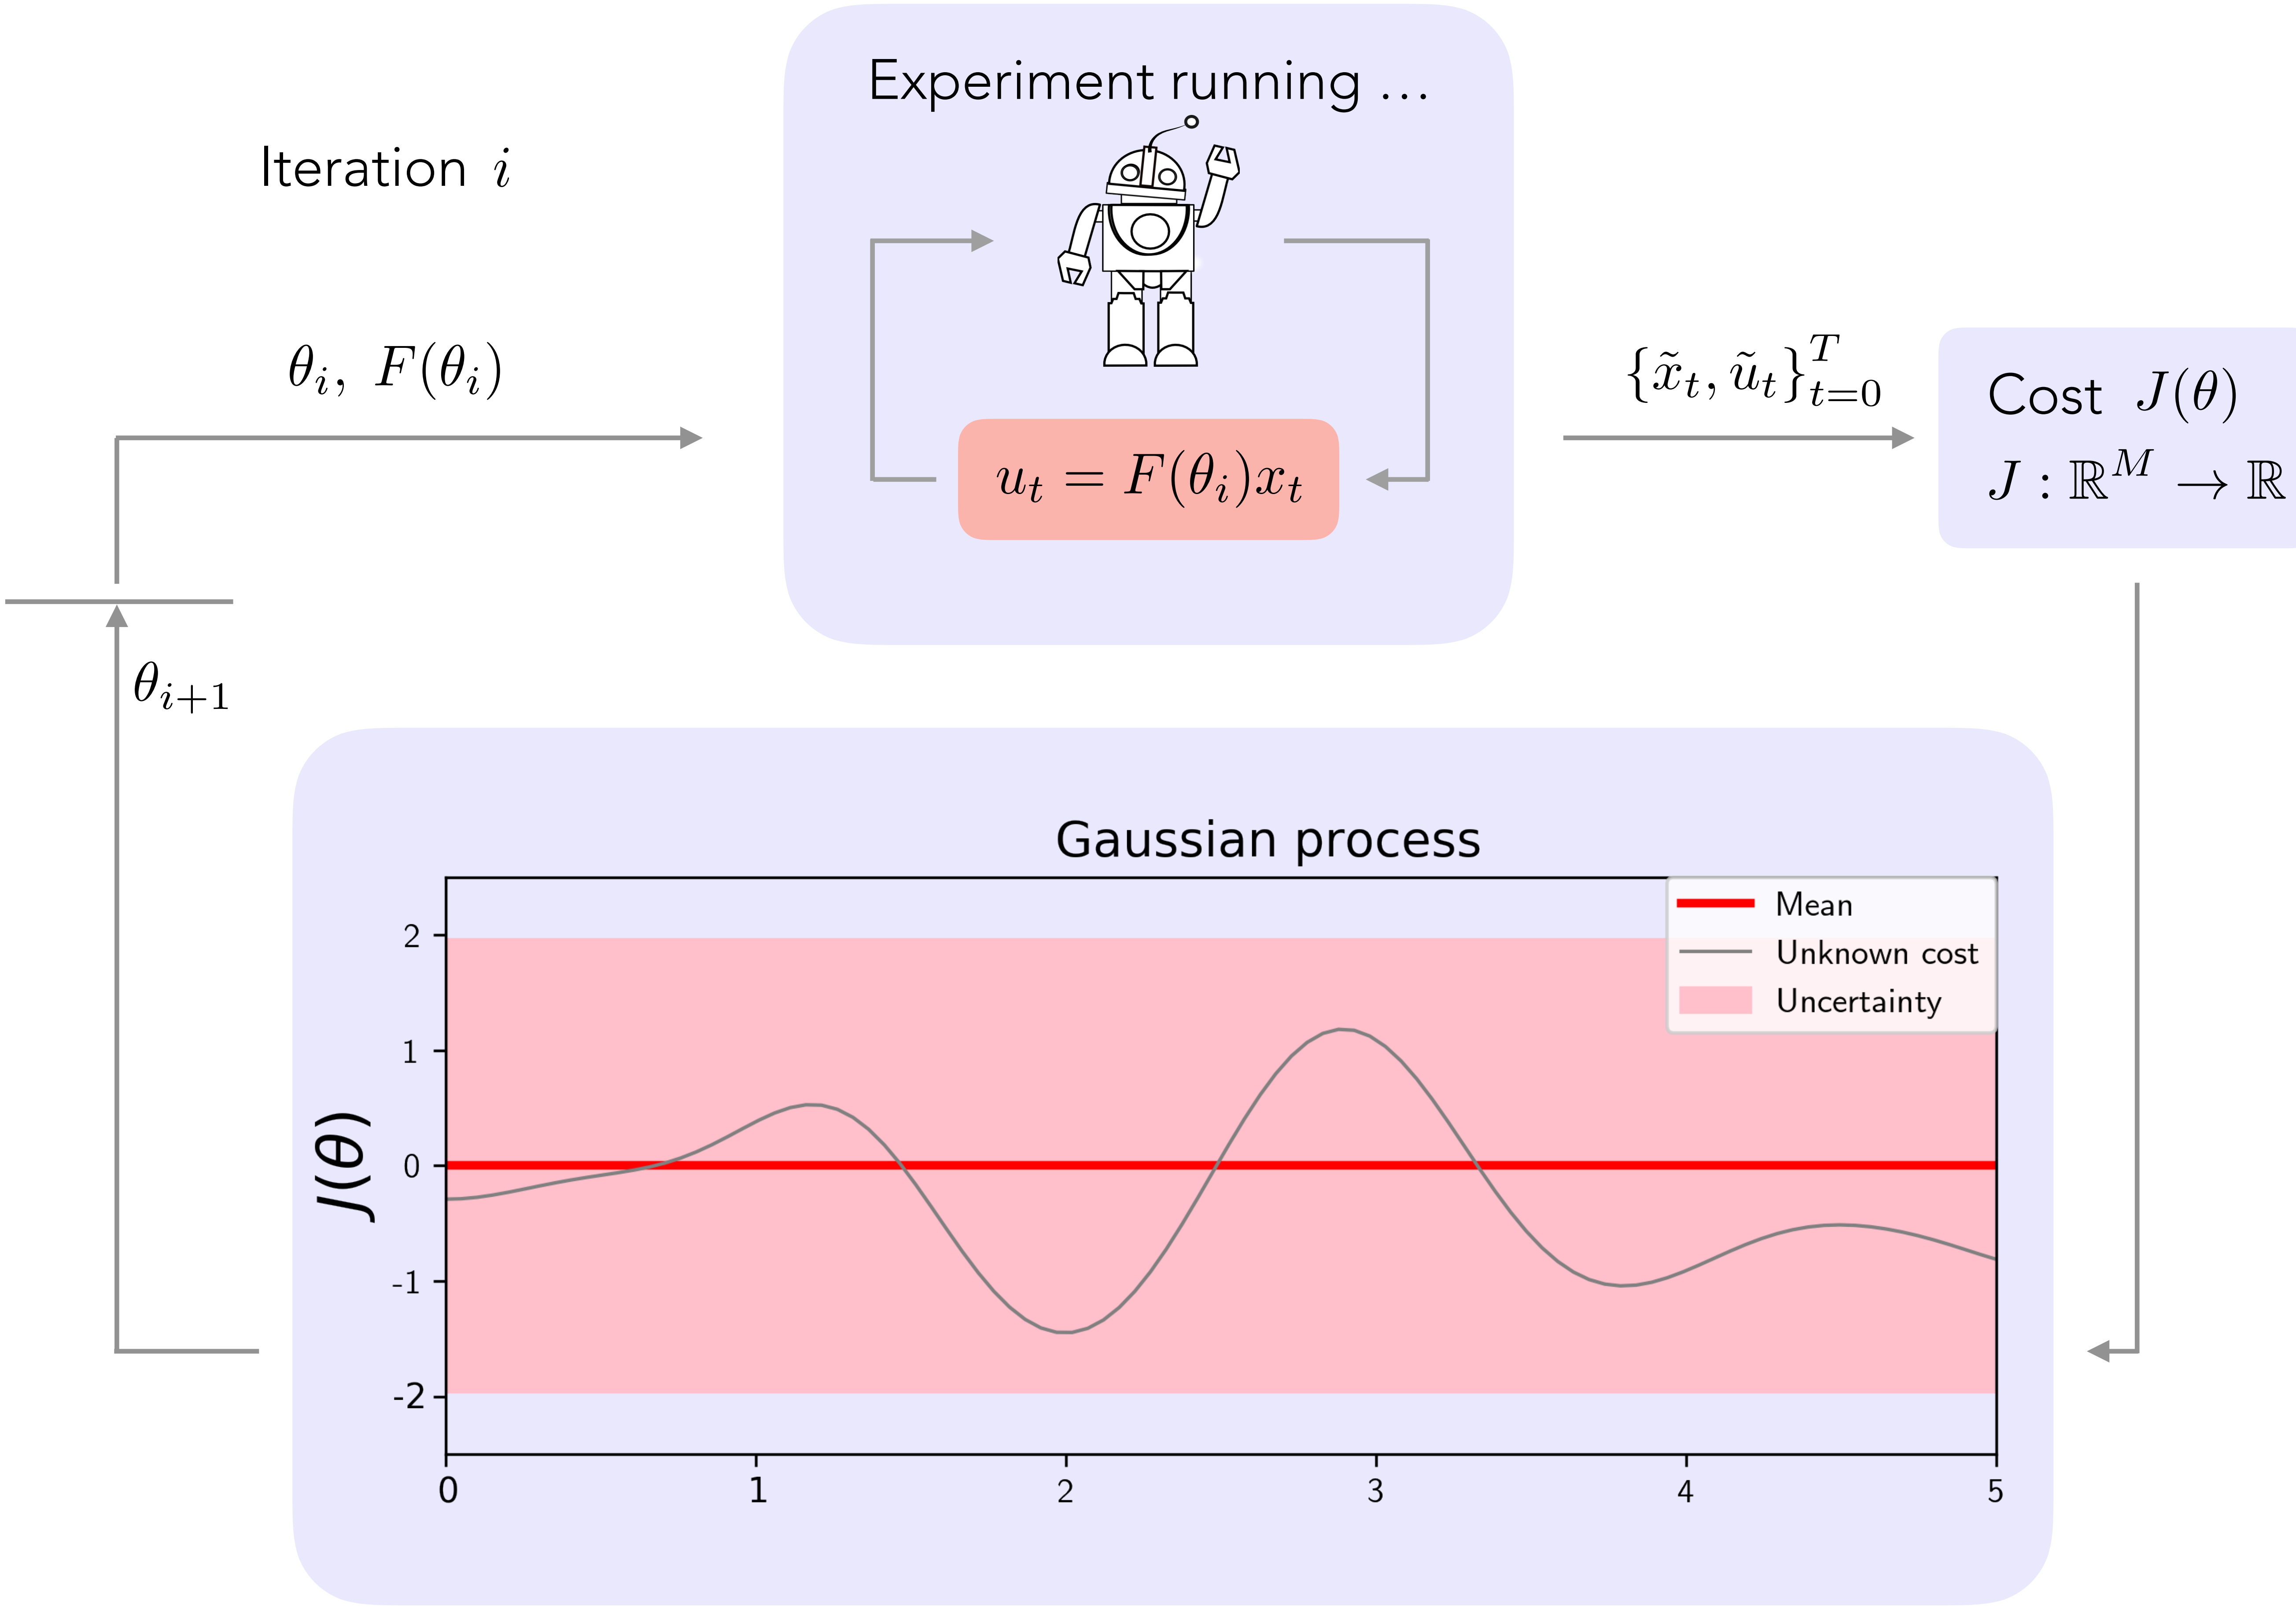
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Unknown cost
Uncertainty
J(θ)
2
1
0
-1
-2
0
1
2
3
4
5

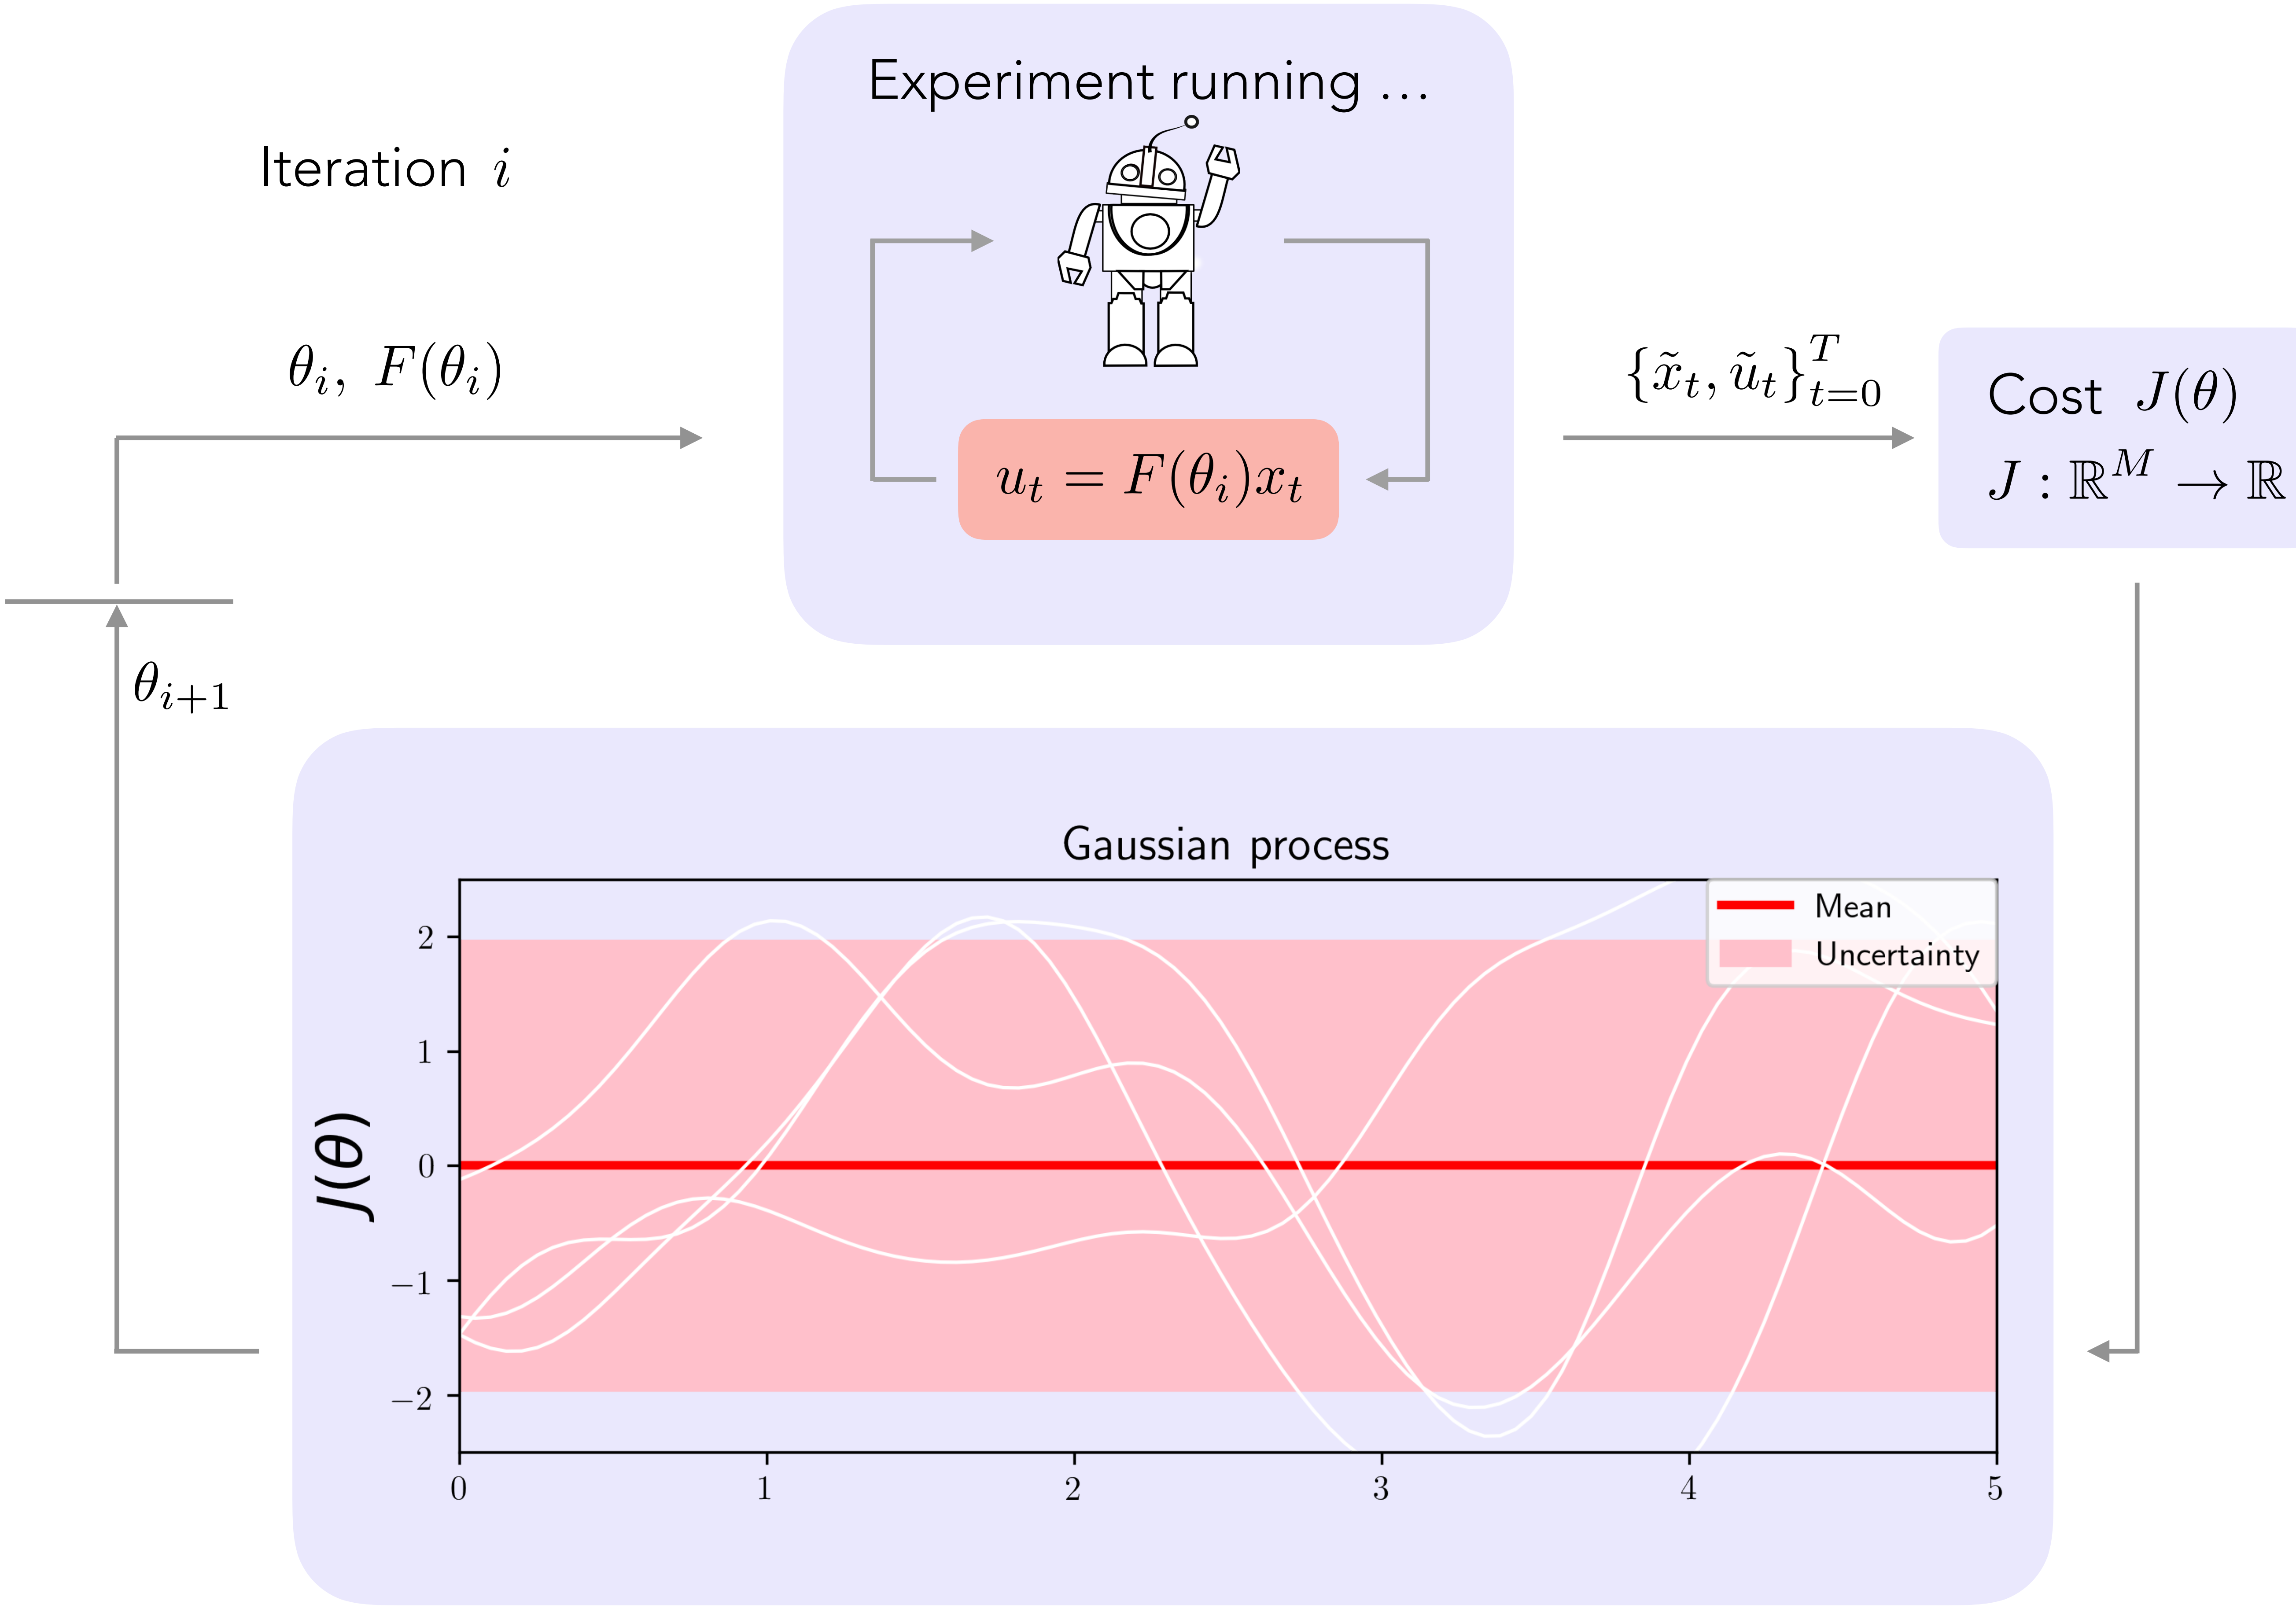

Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5

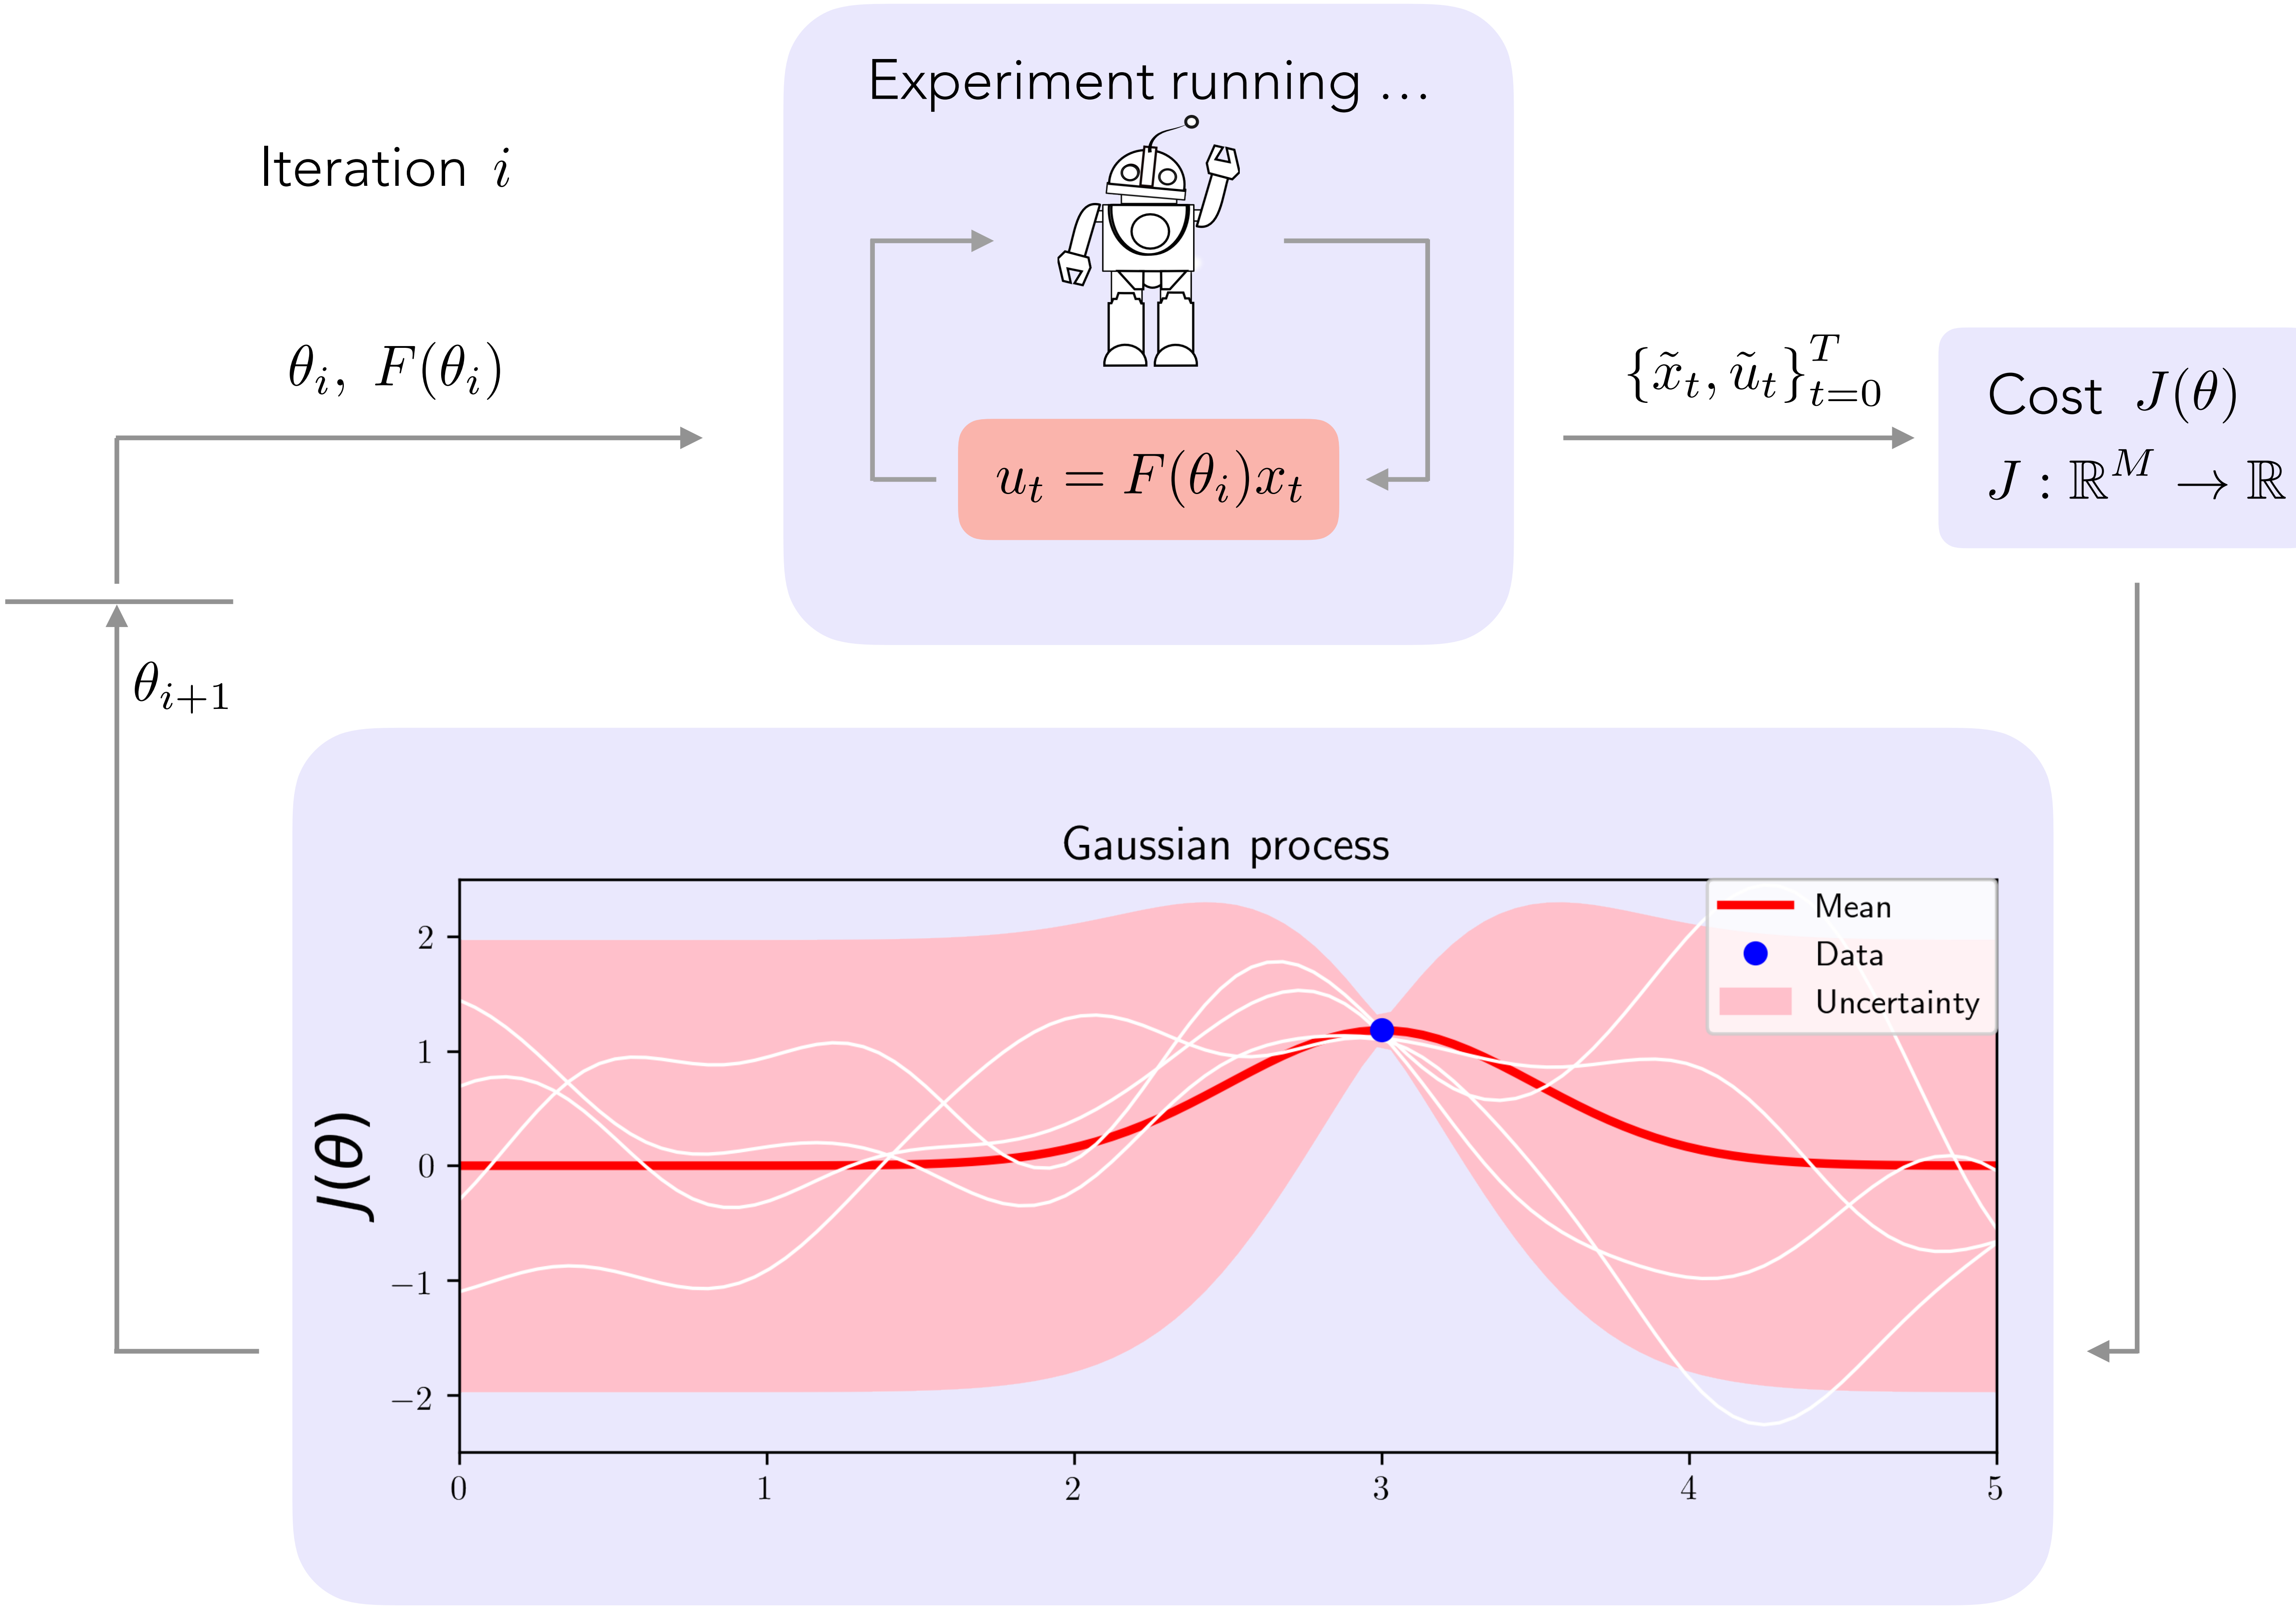
Iteration i
Experiment running ...
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
-1
-2
0
1
2
3
4
5

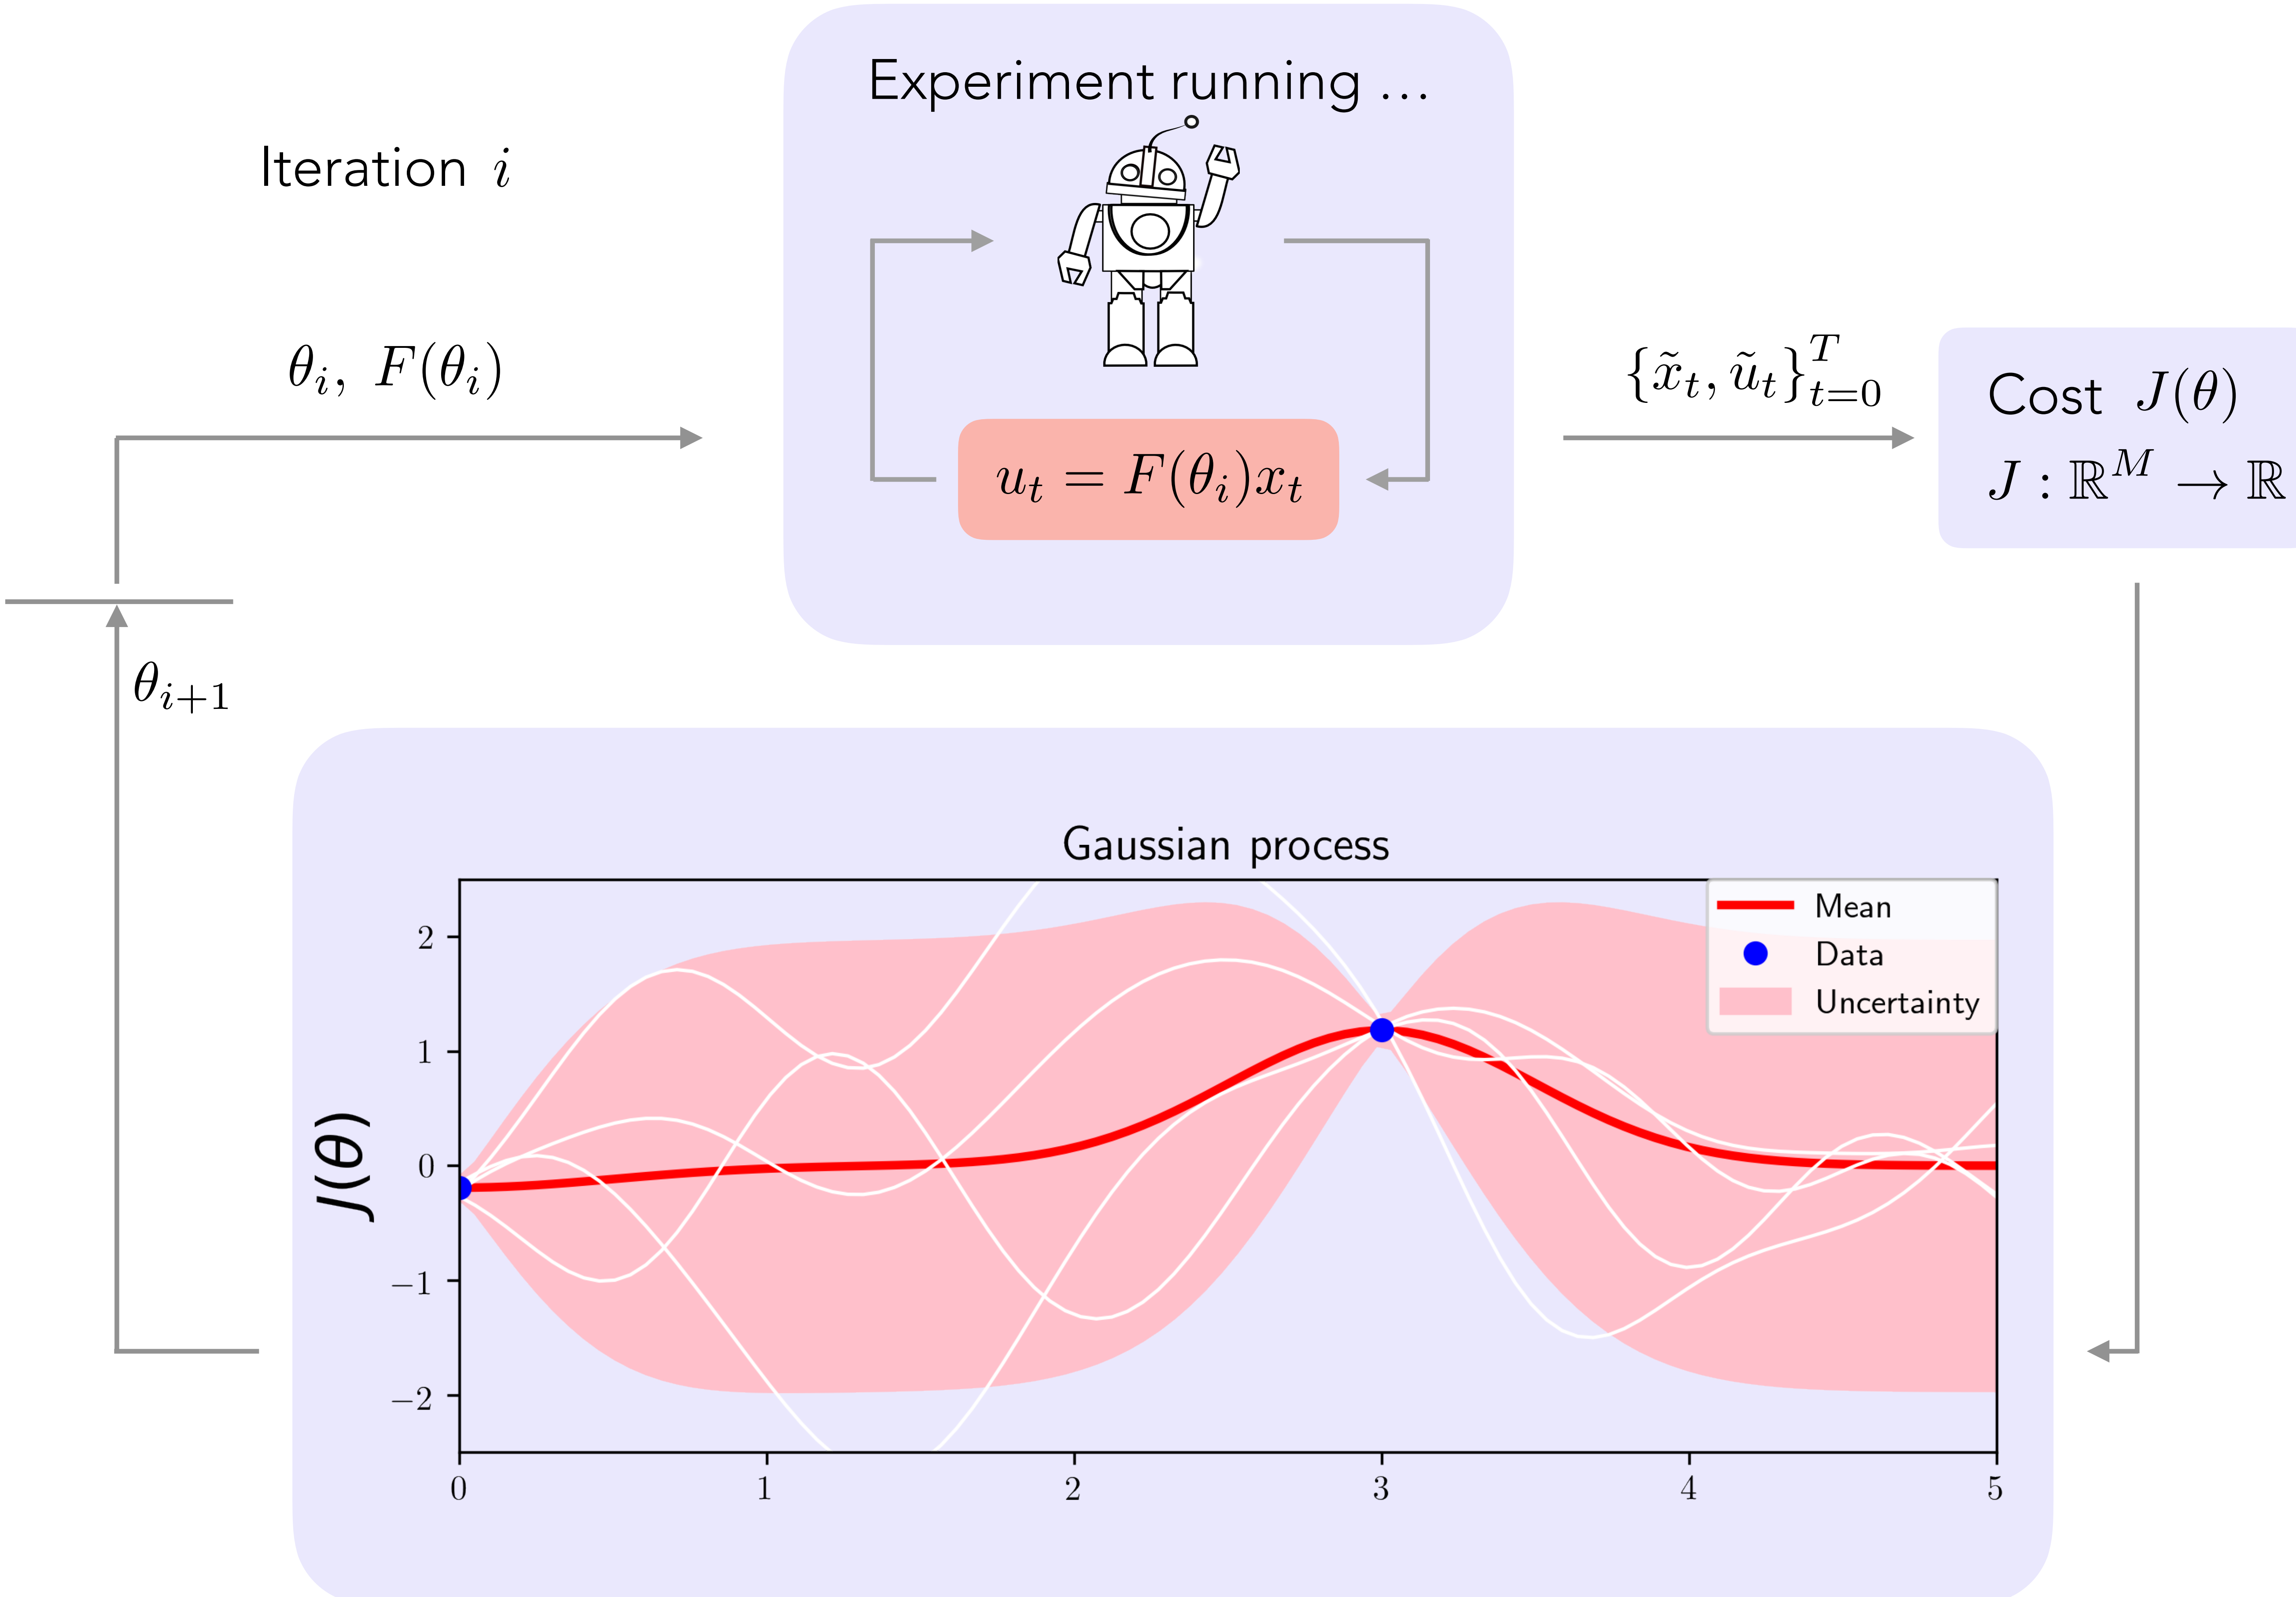
Iteration i
Experiment running ...
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5

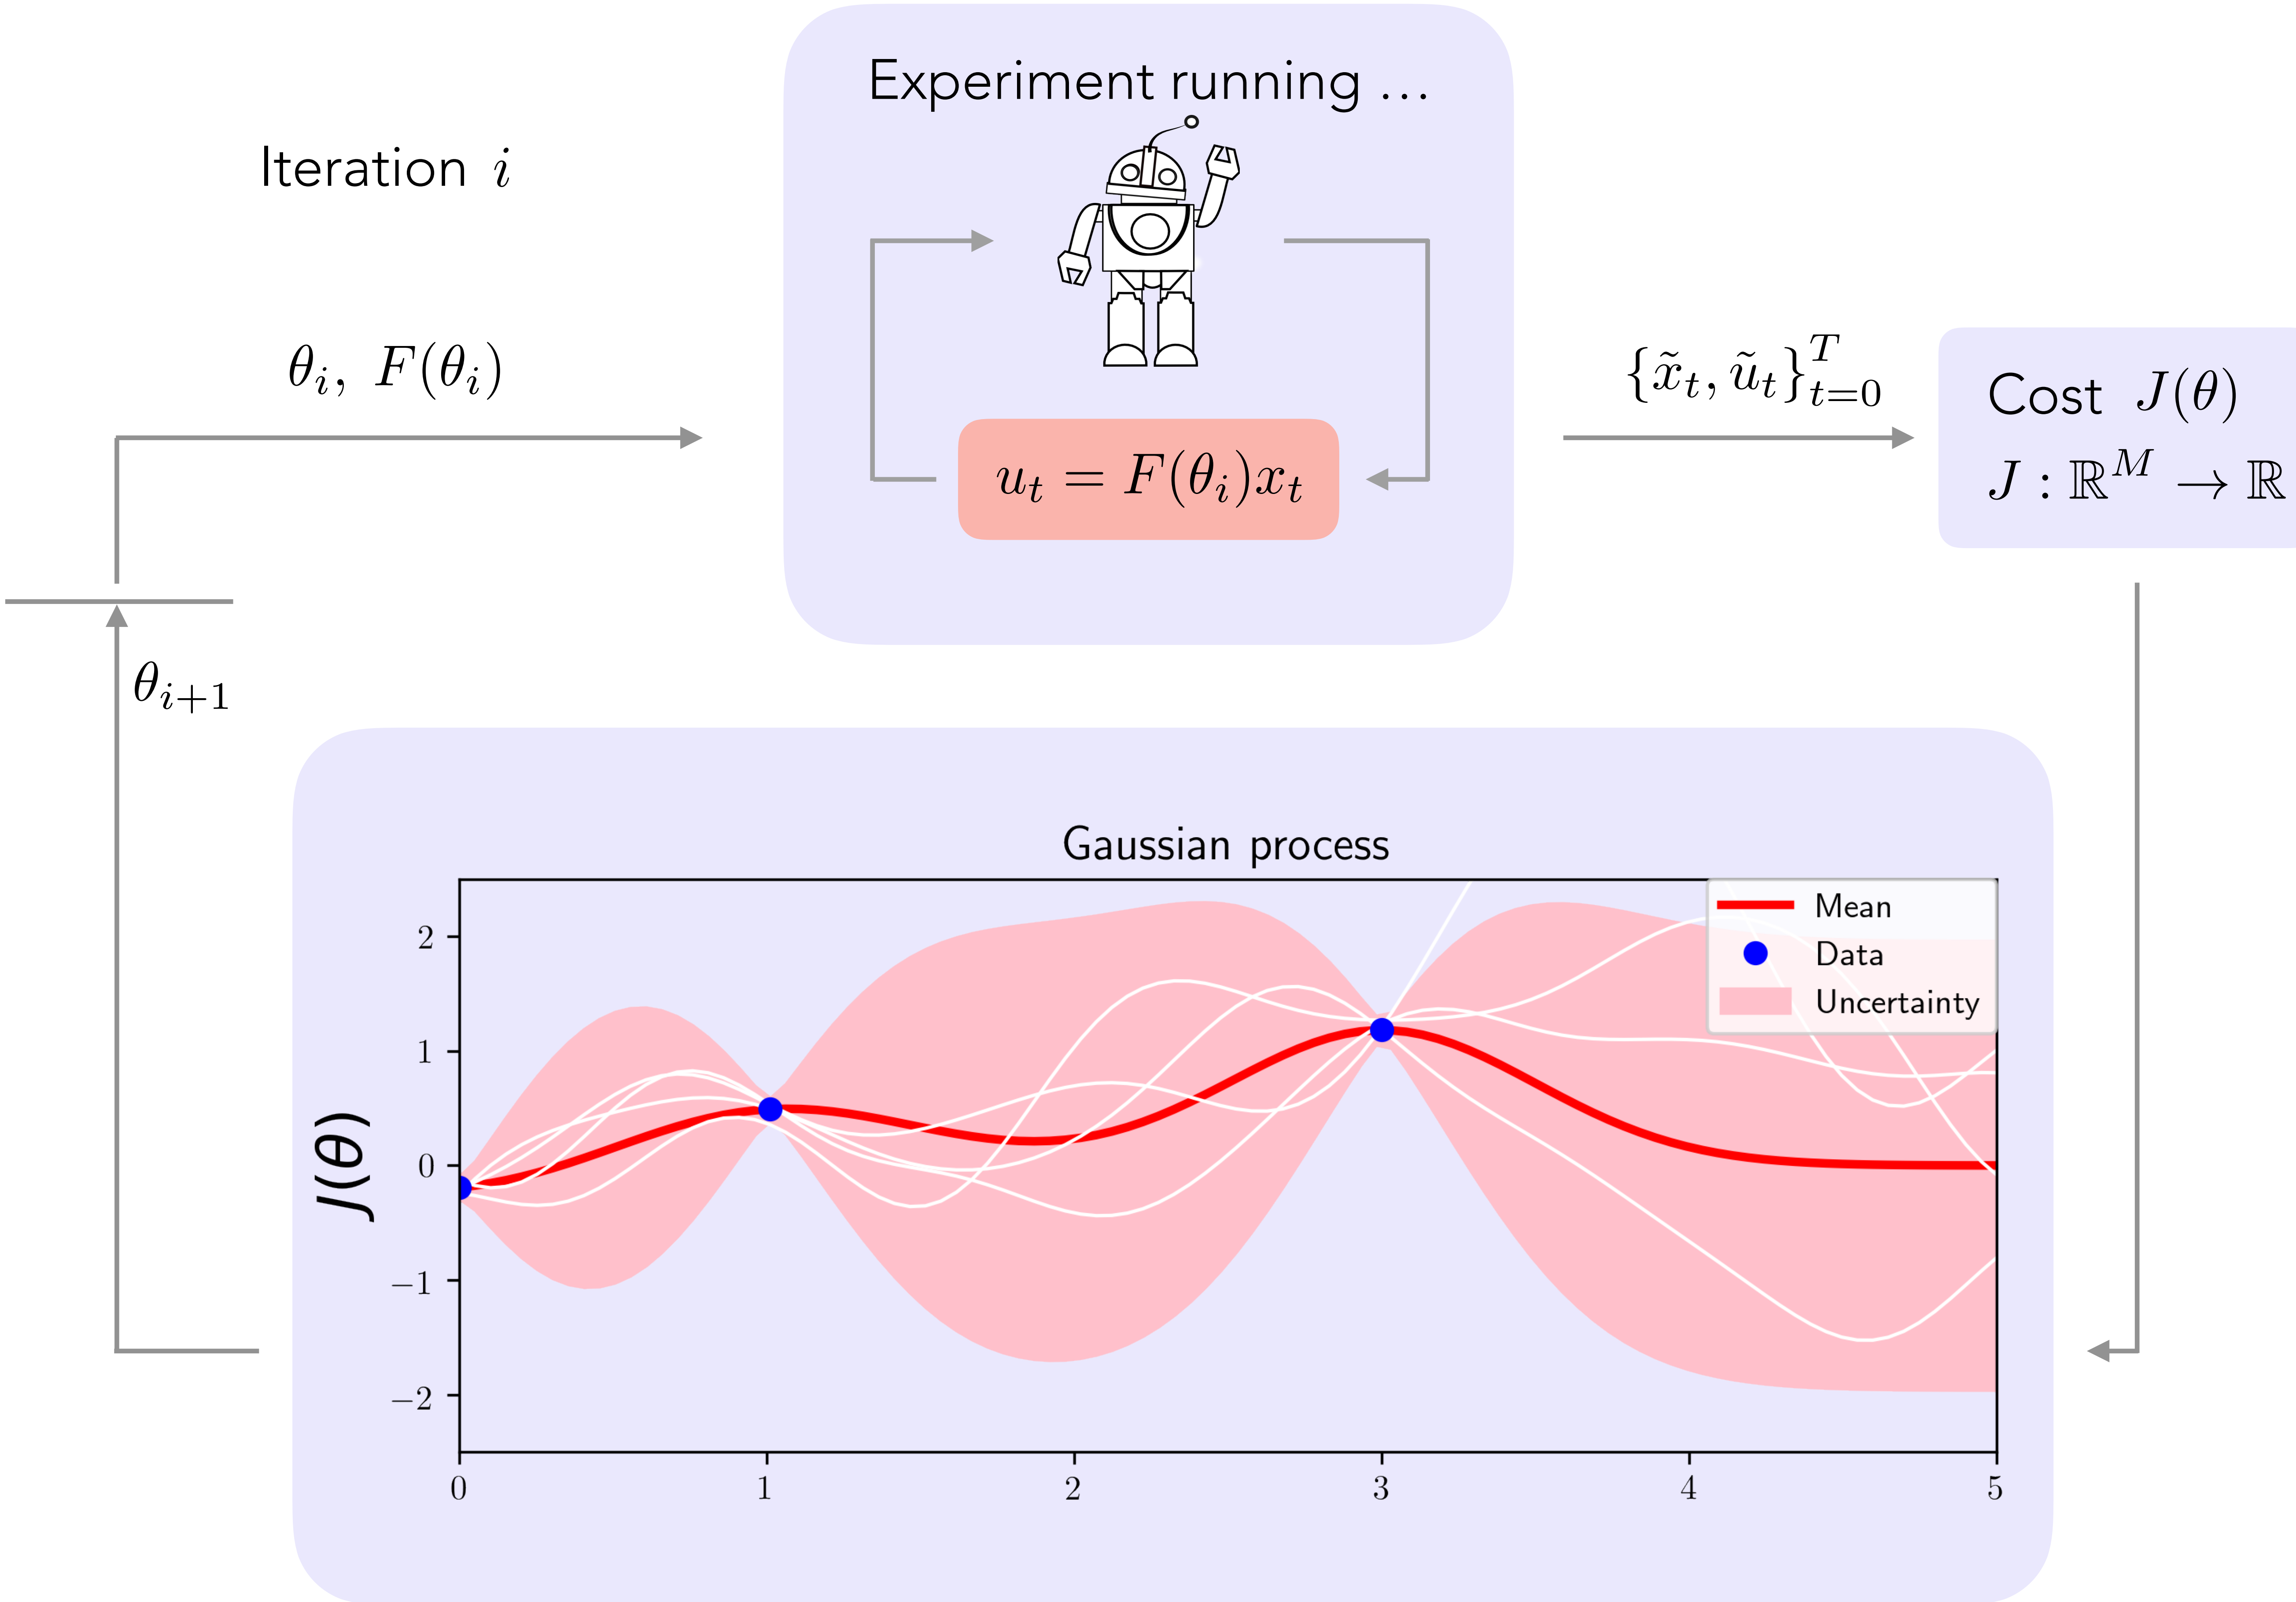
Experiment running …
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$

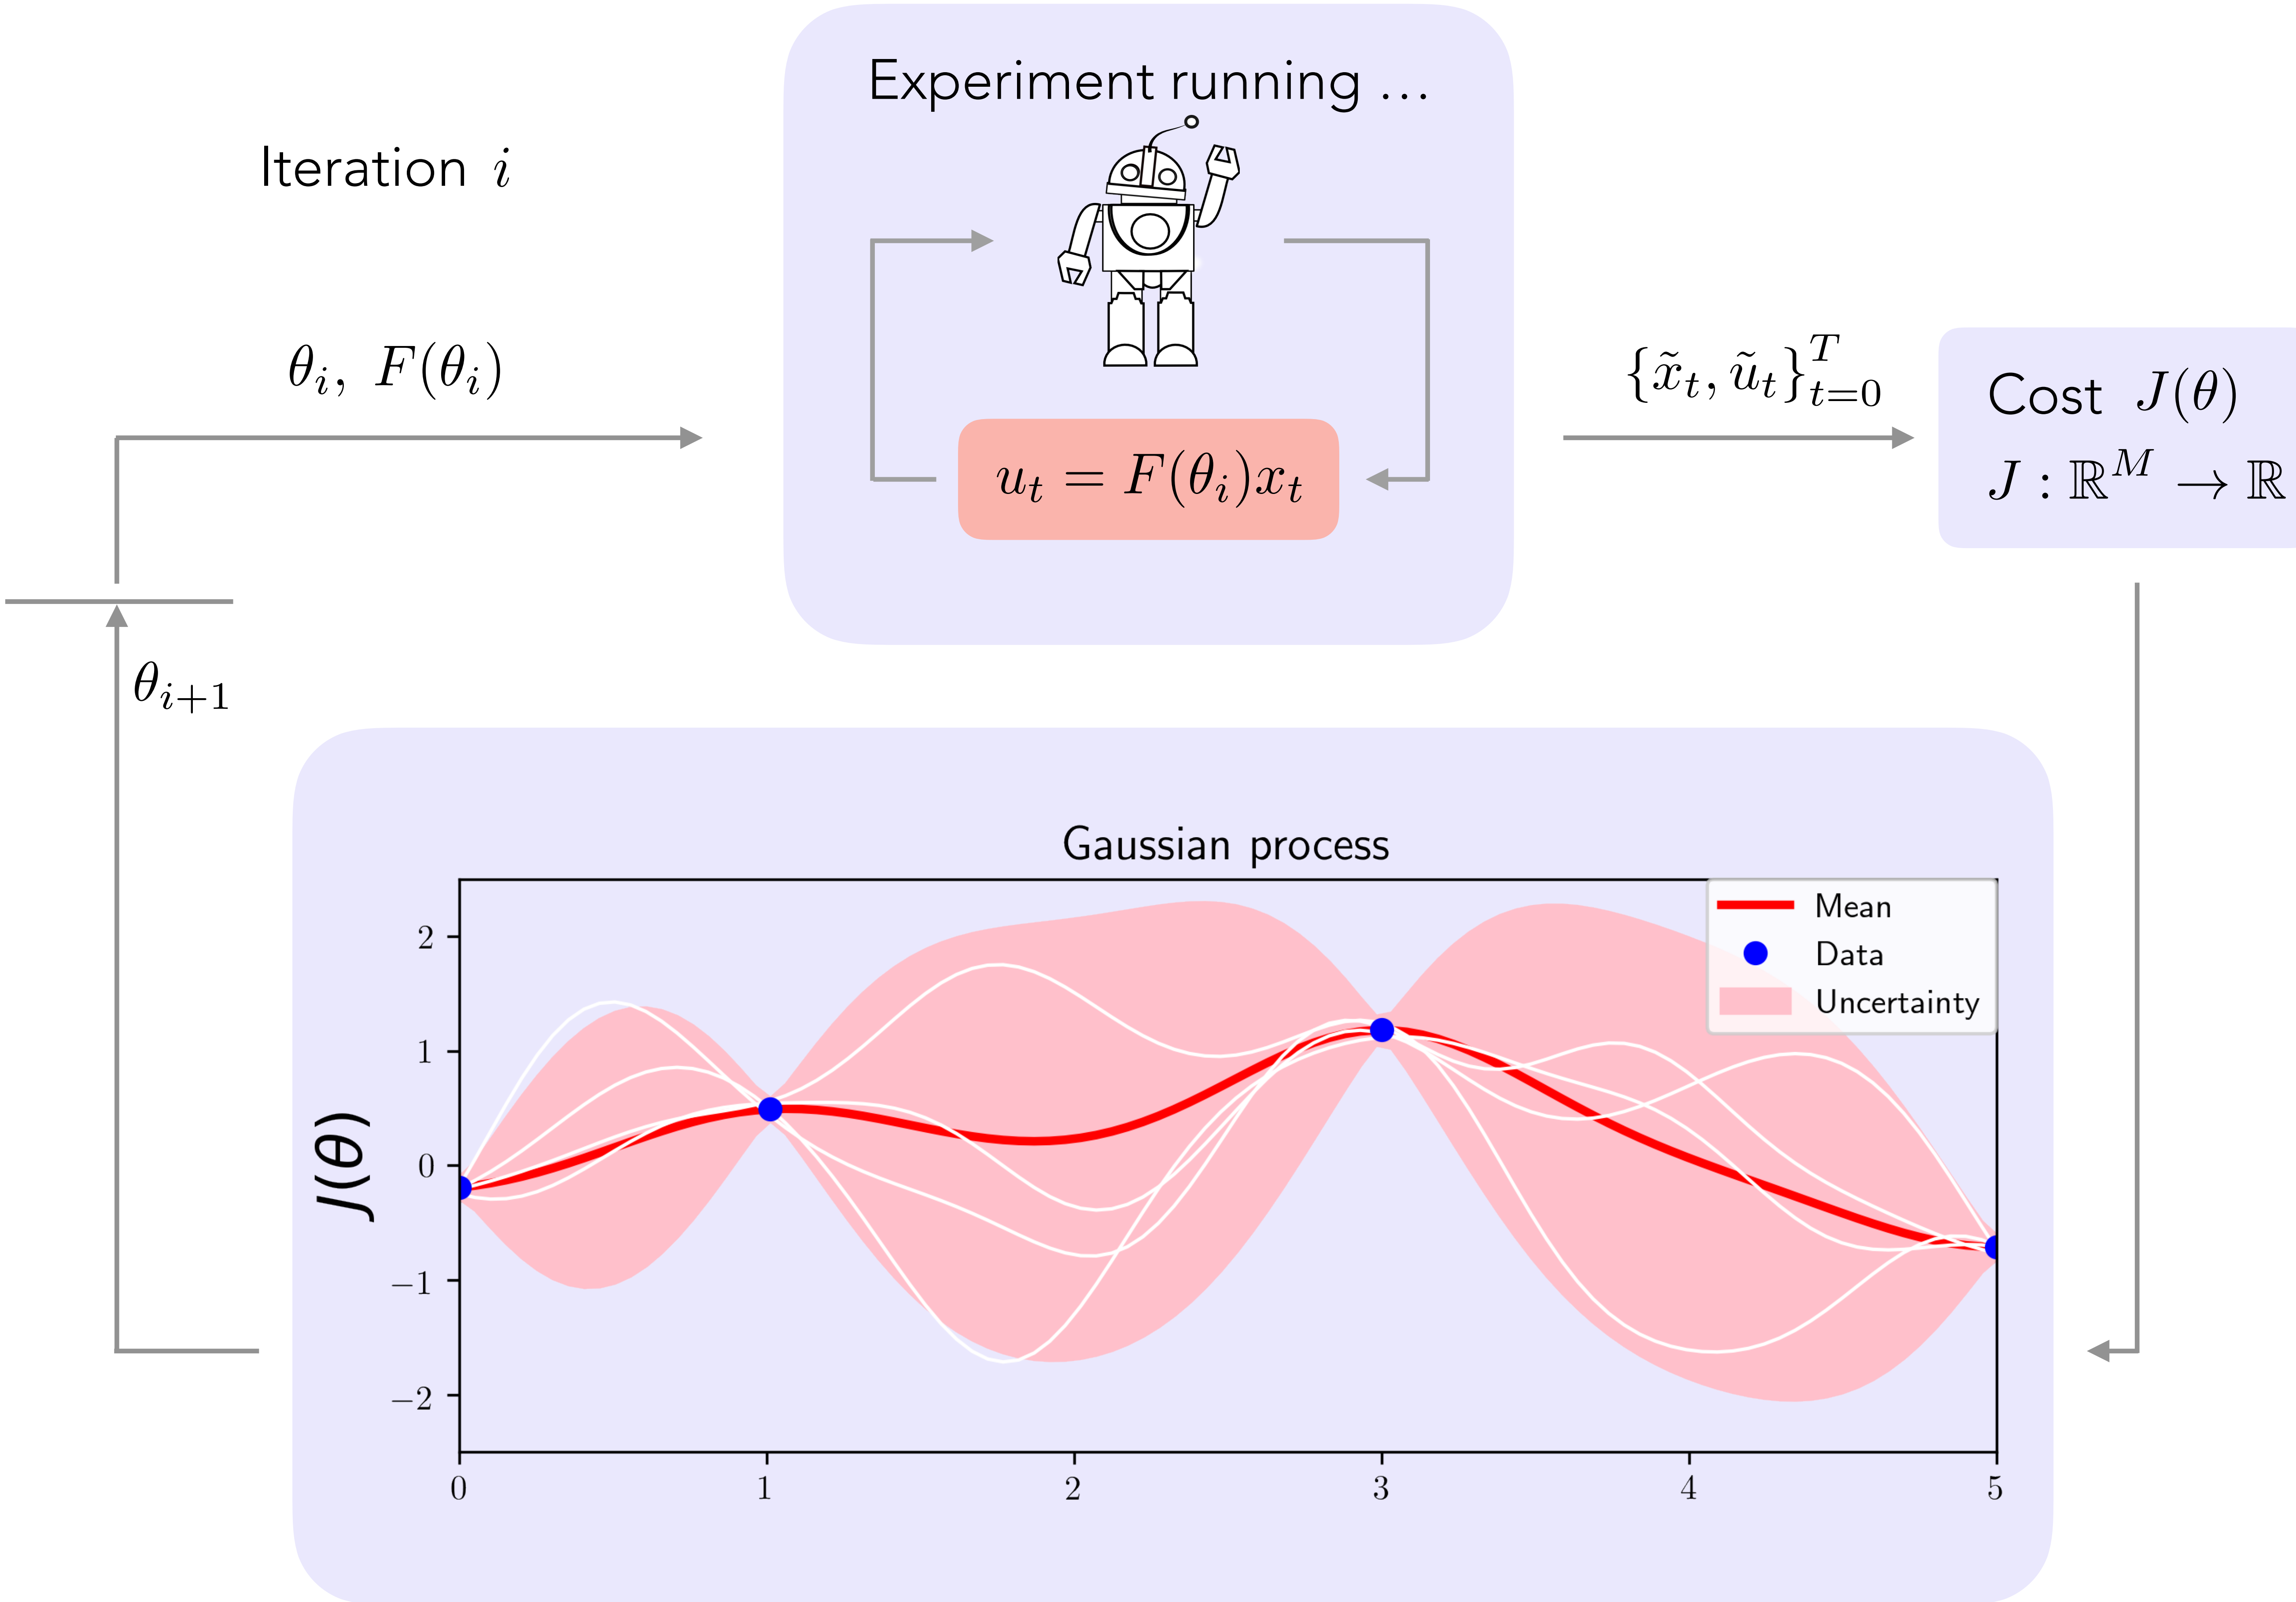
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5

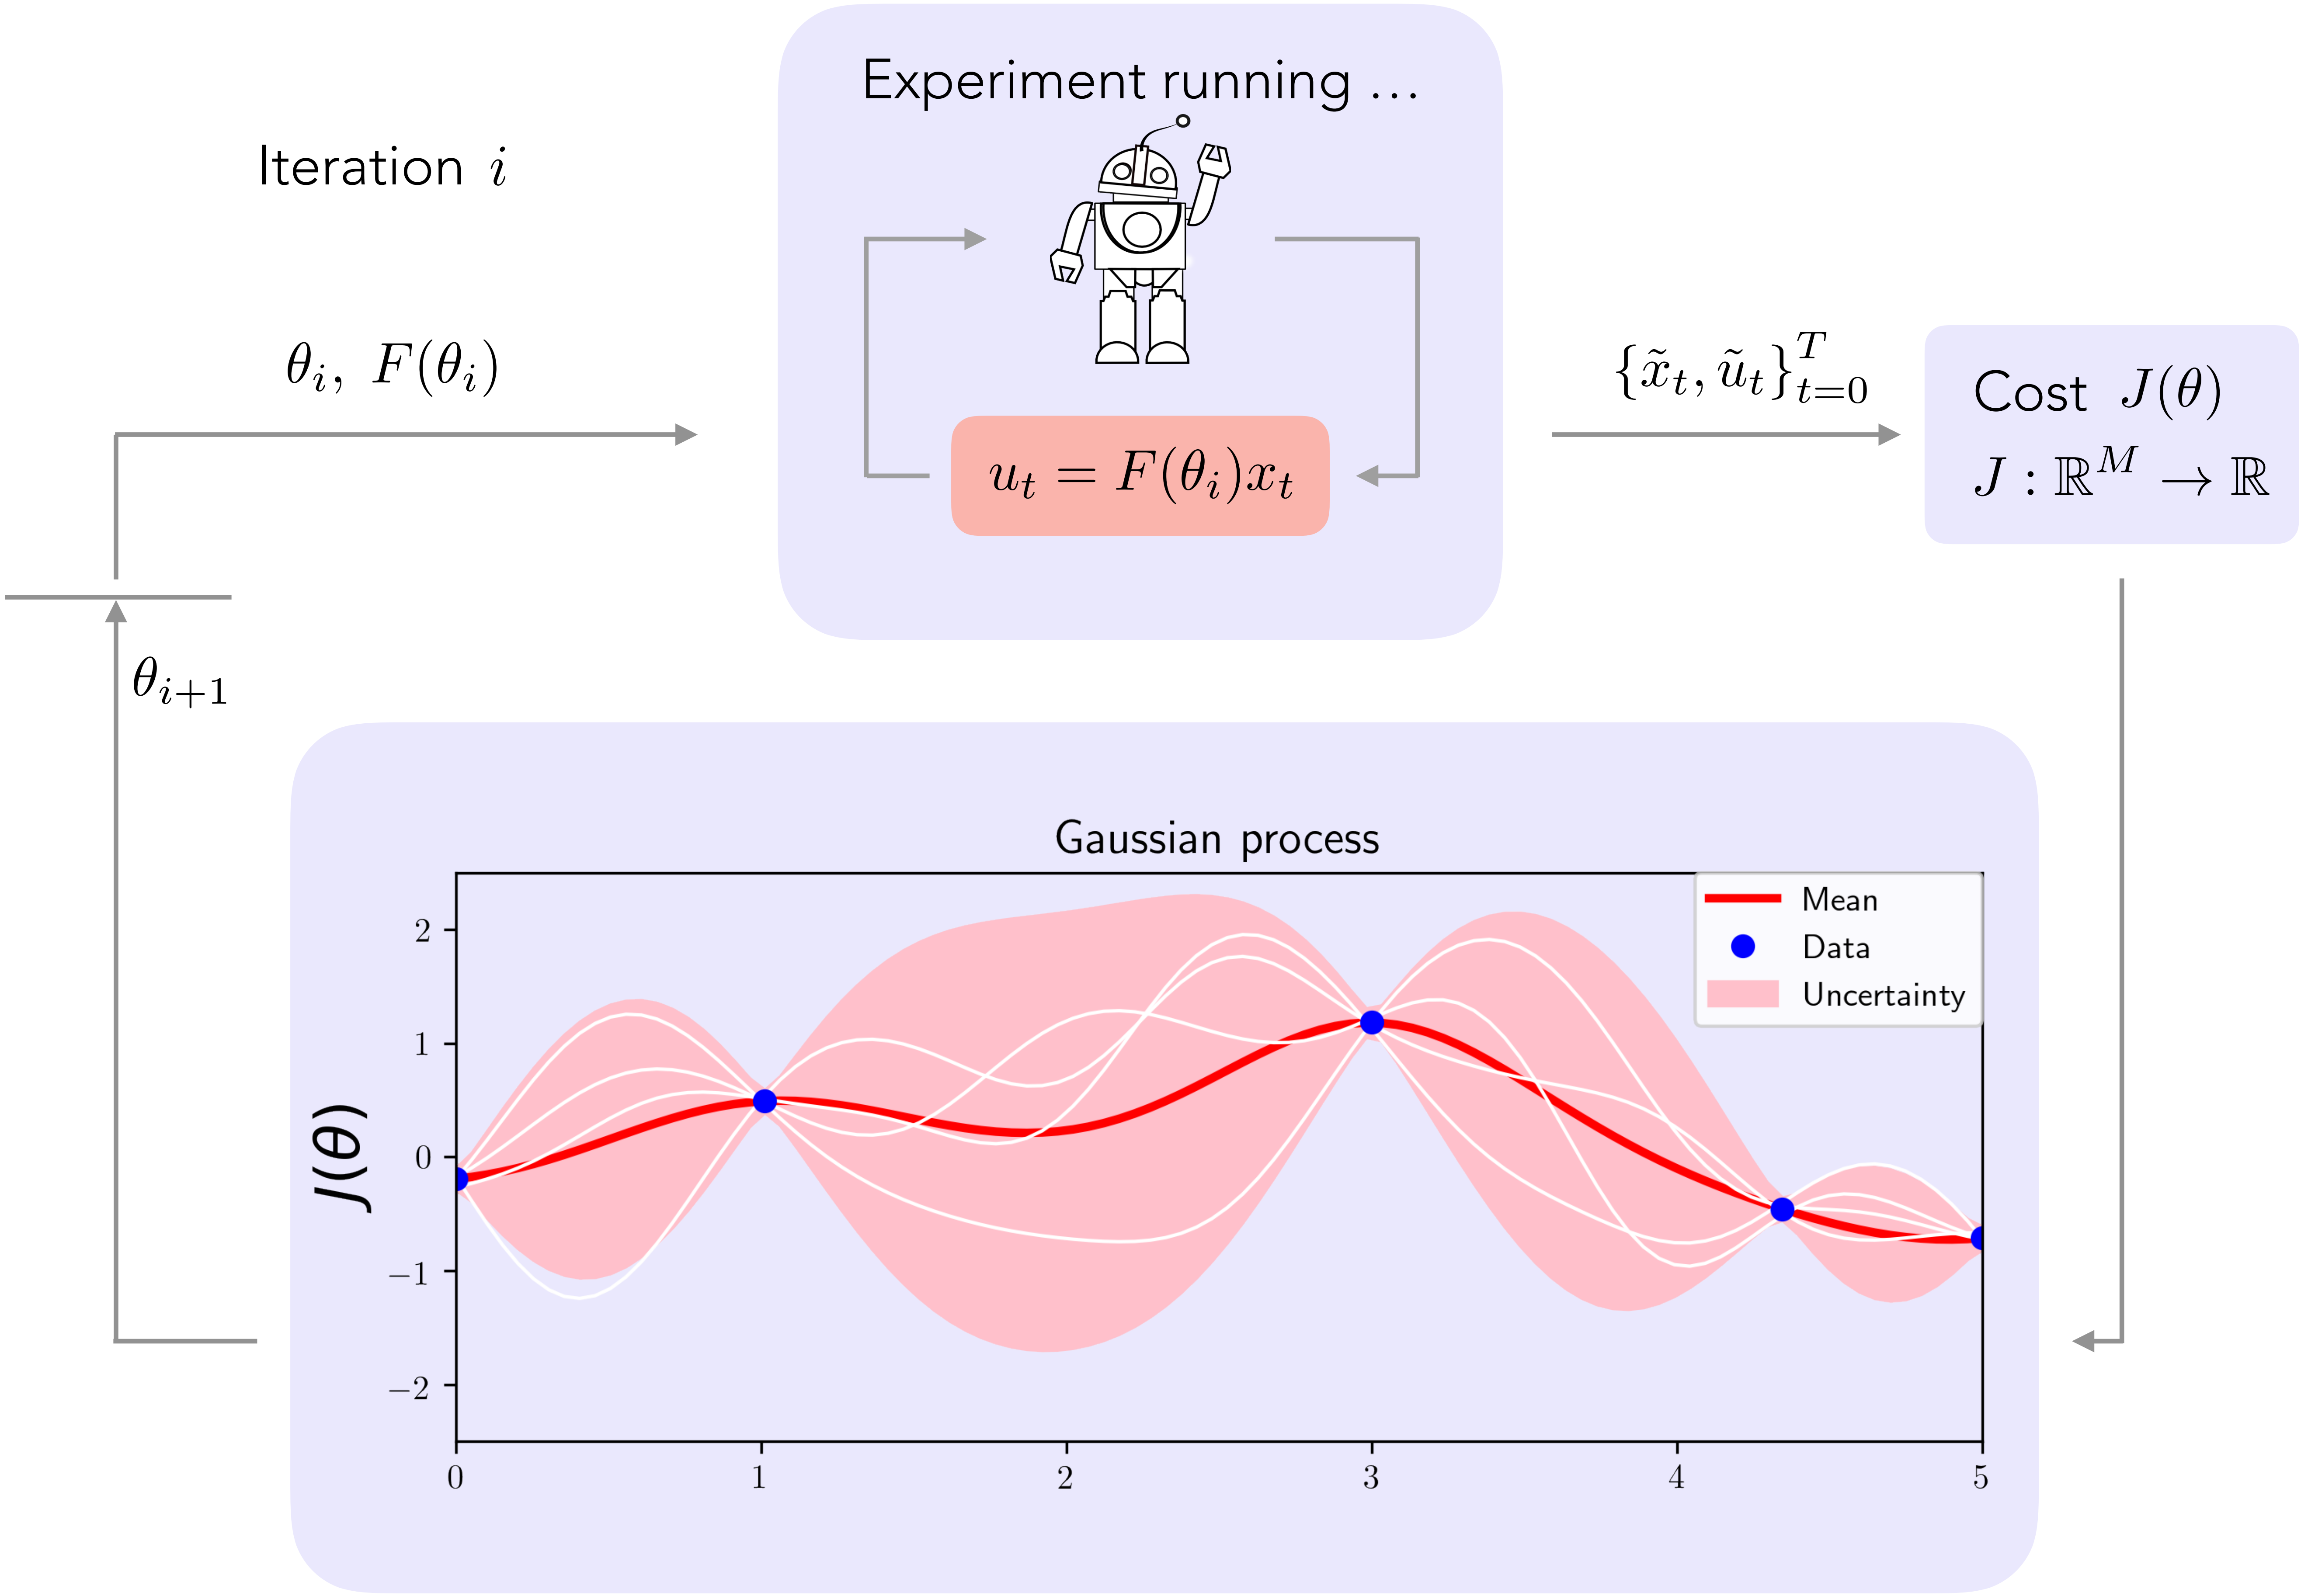
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
J(θ)
Mean
Data
Uncertainty
2
1
0
−1
−2
0
1
2
3
4
5

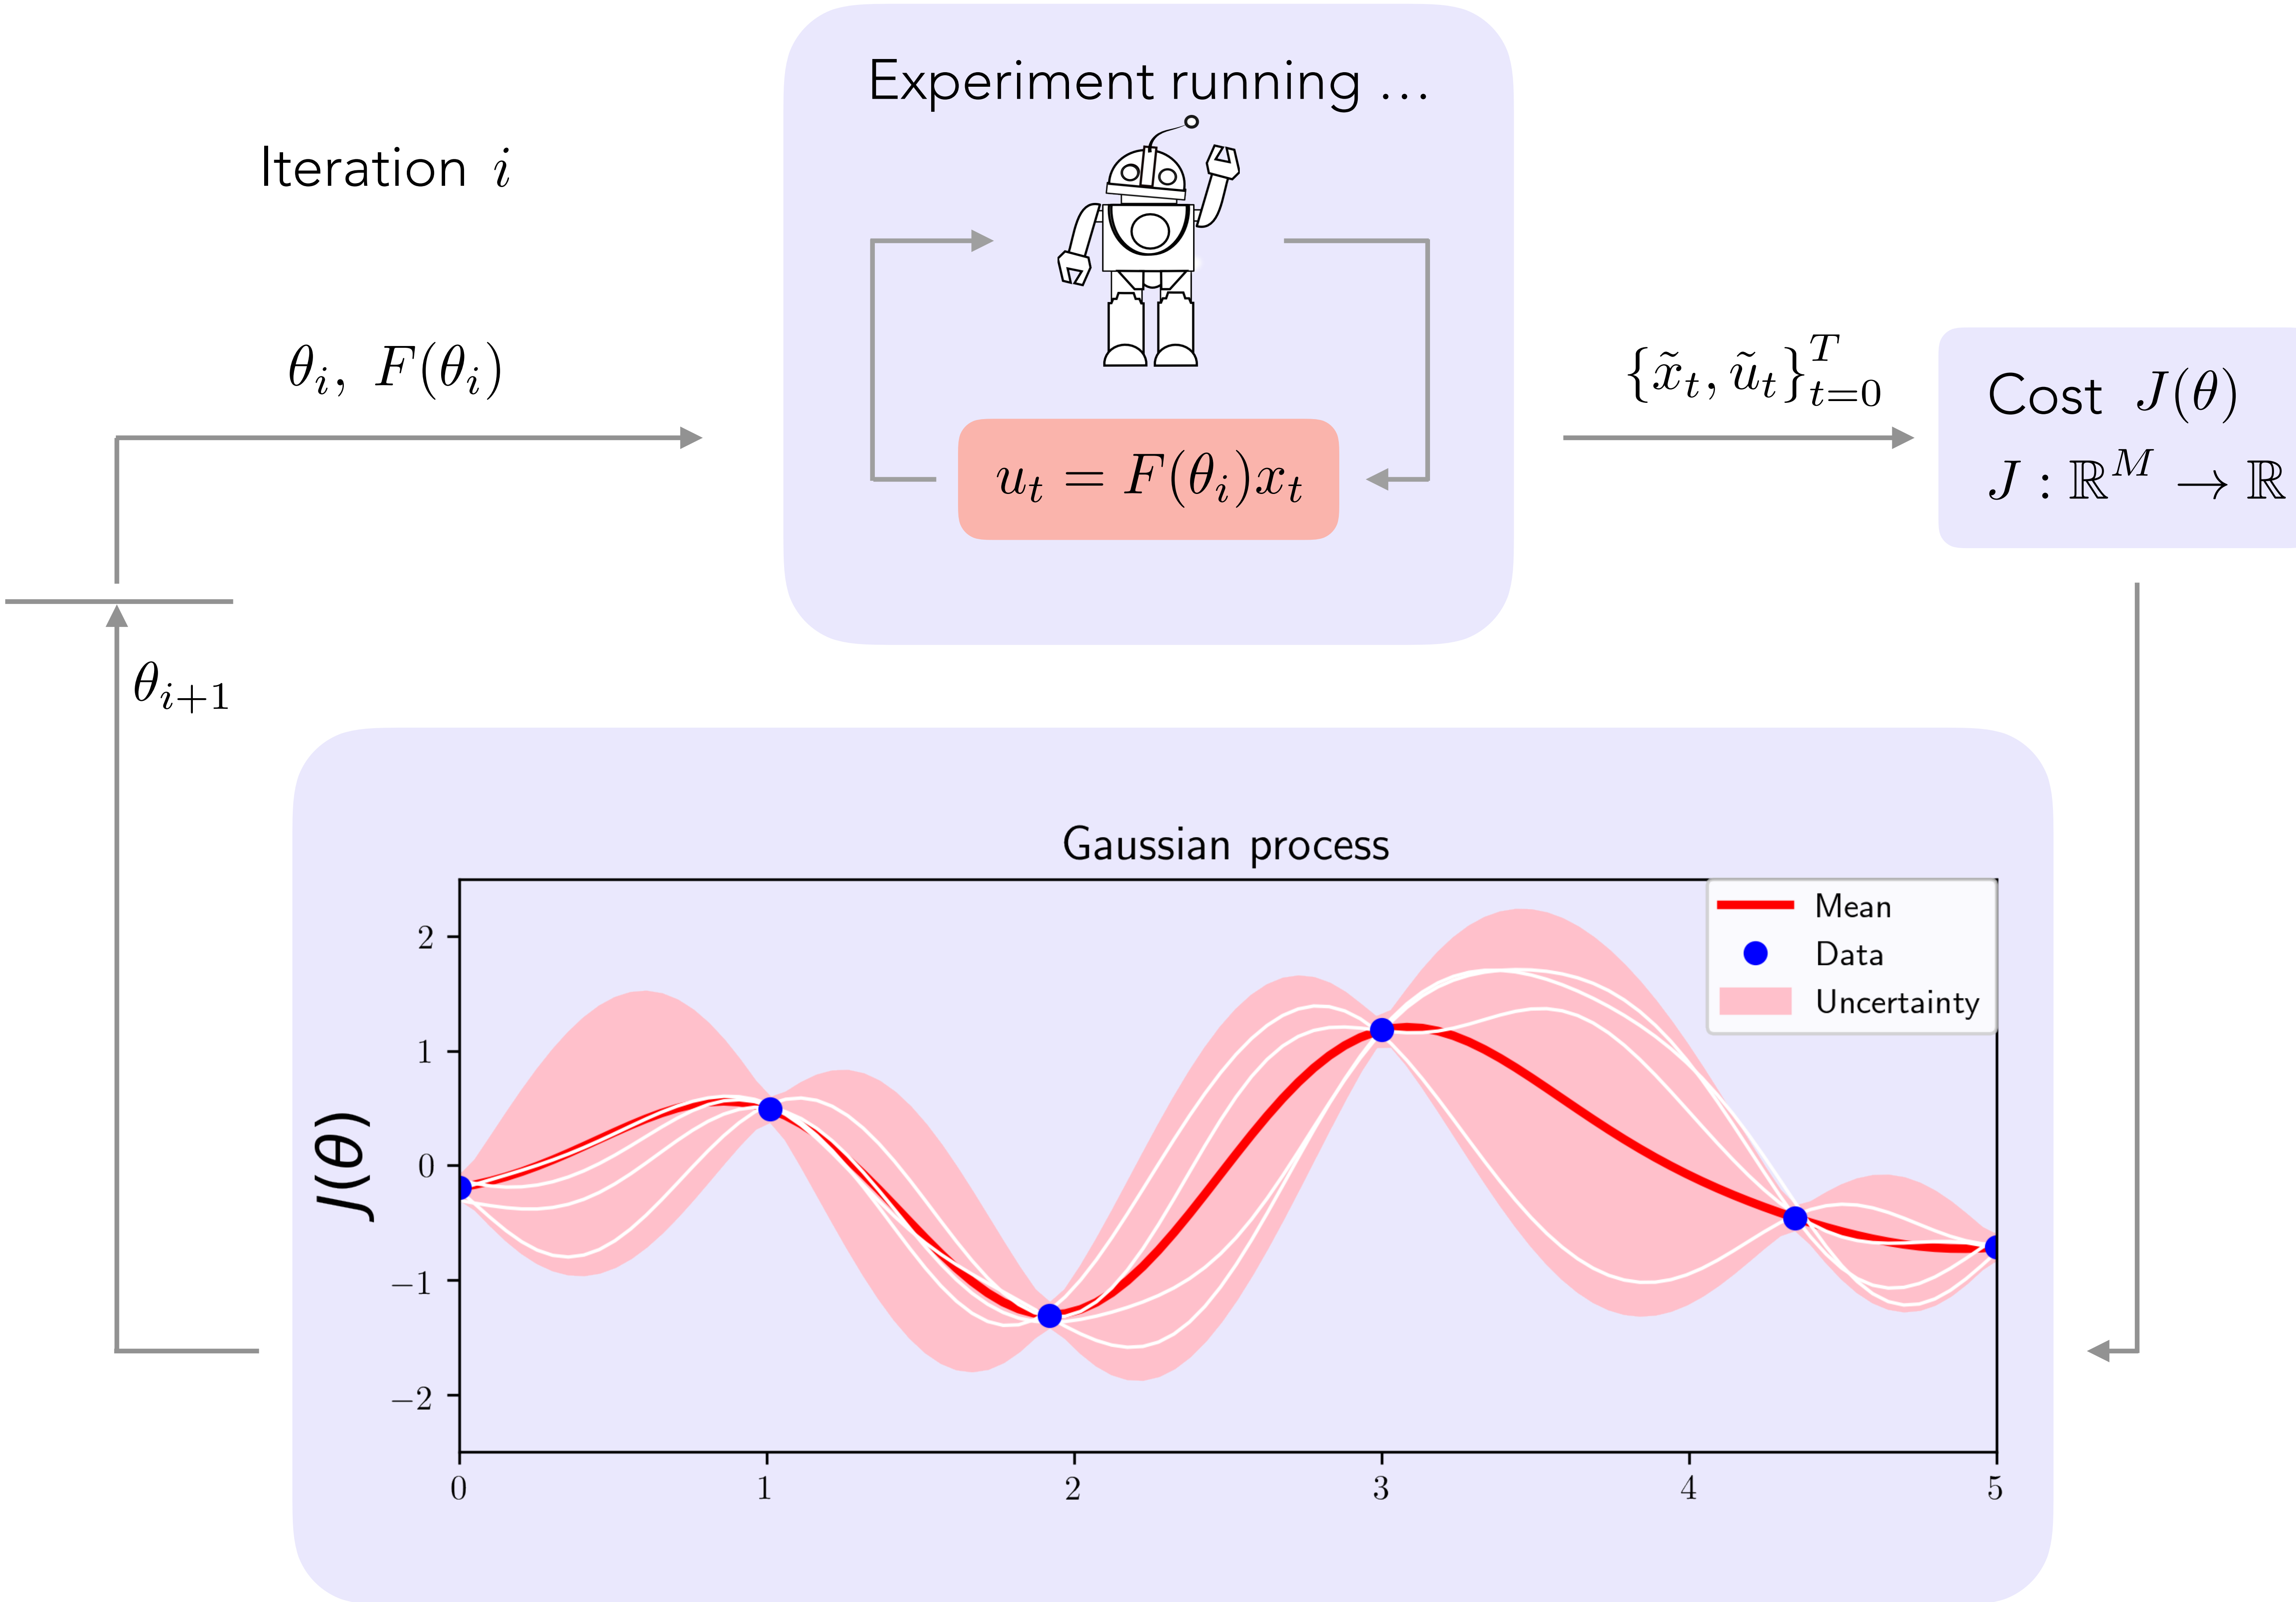
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5

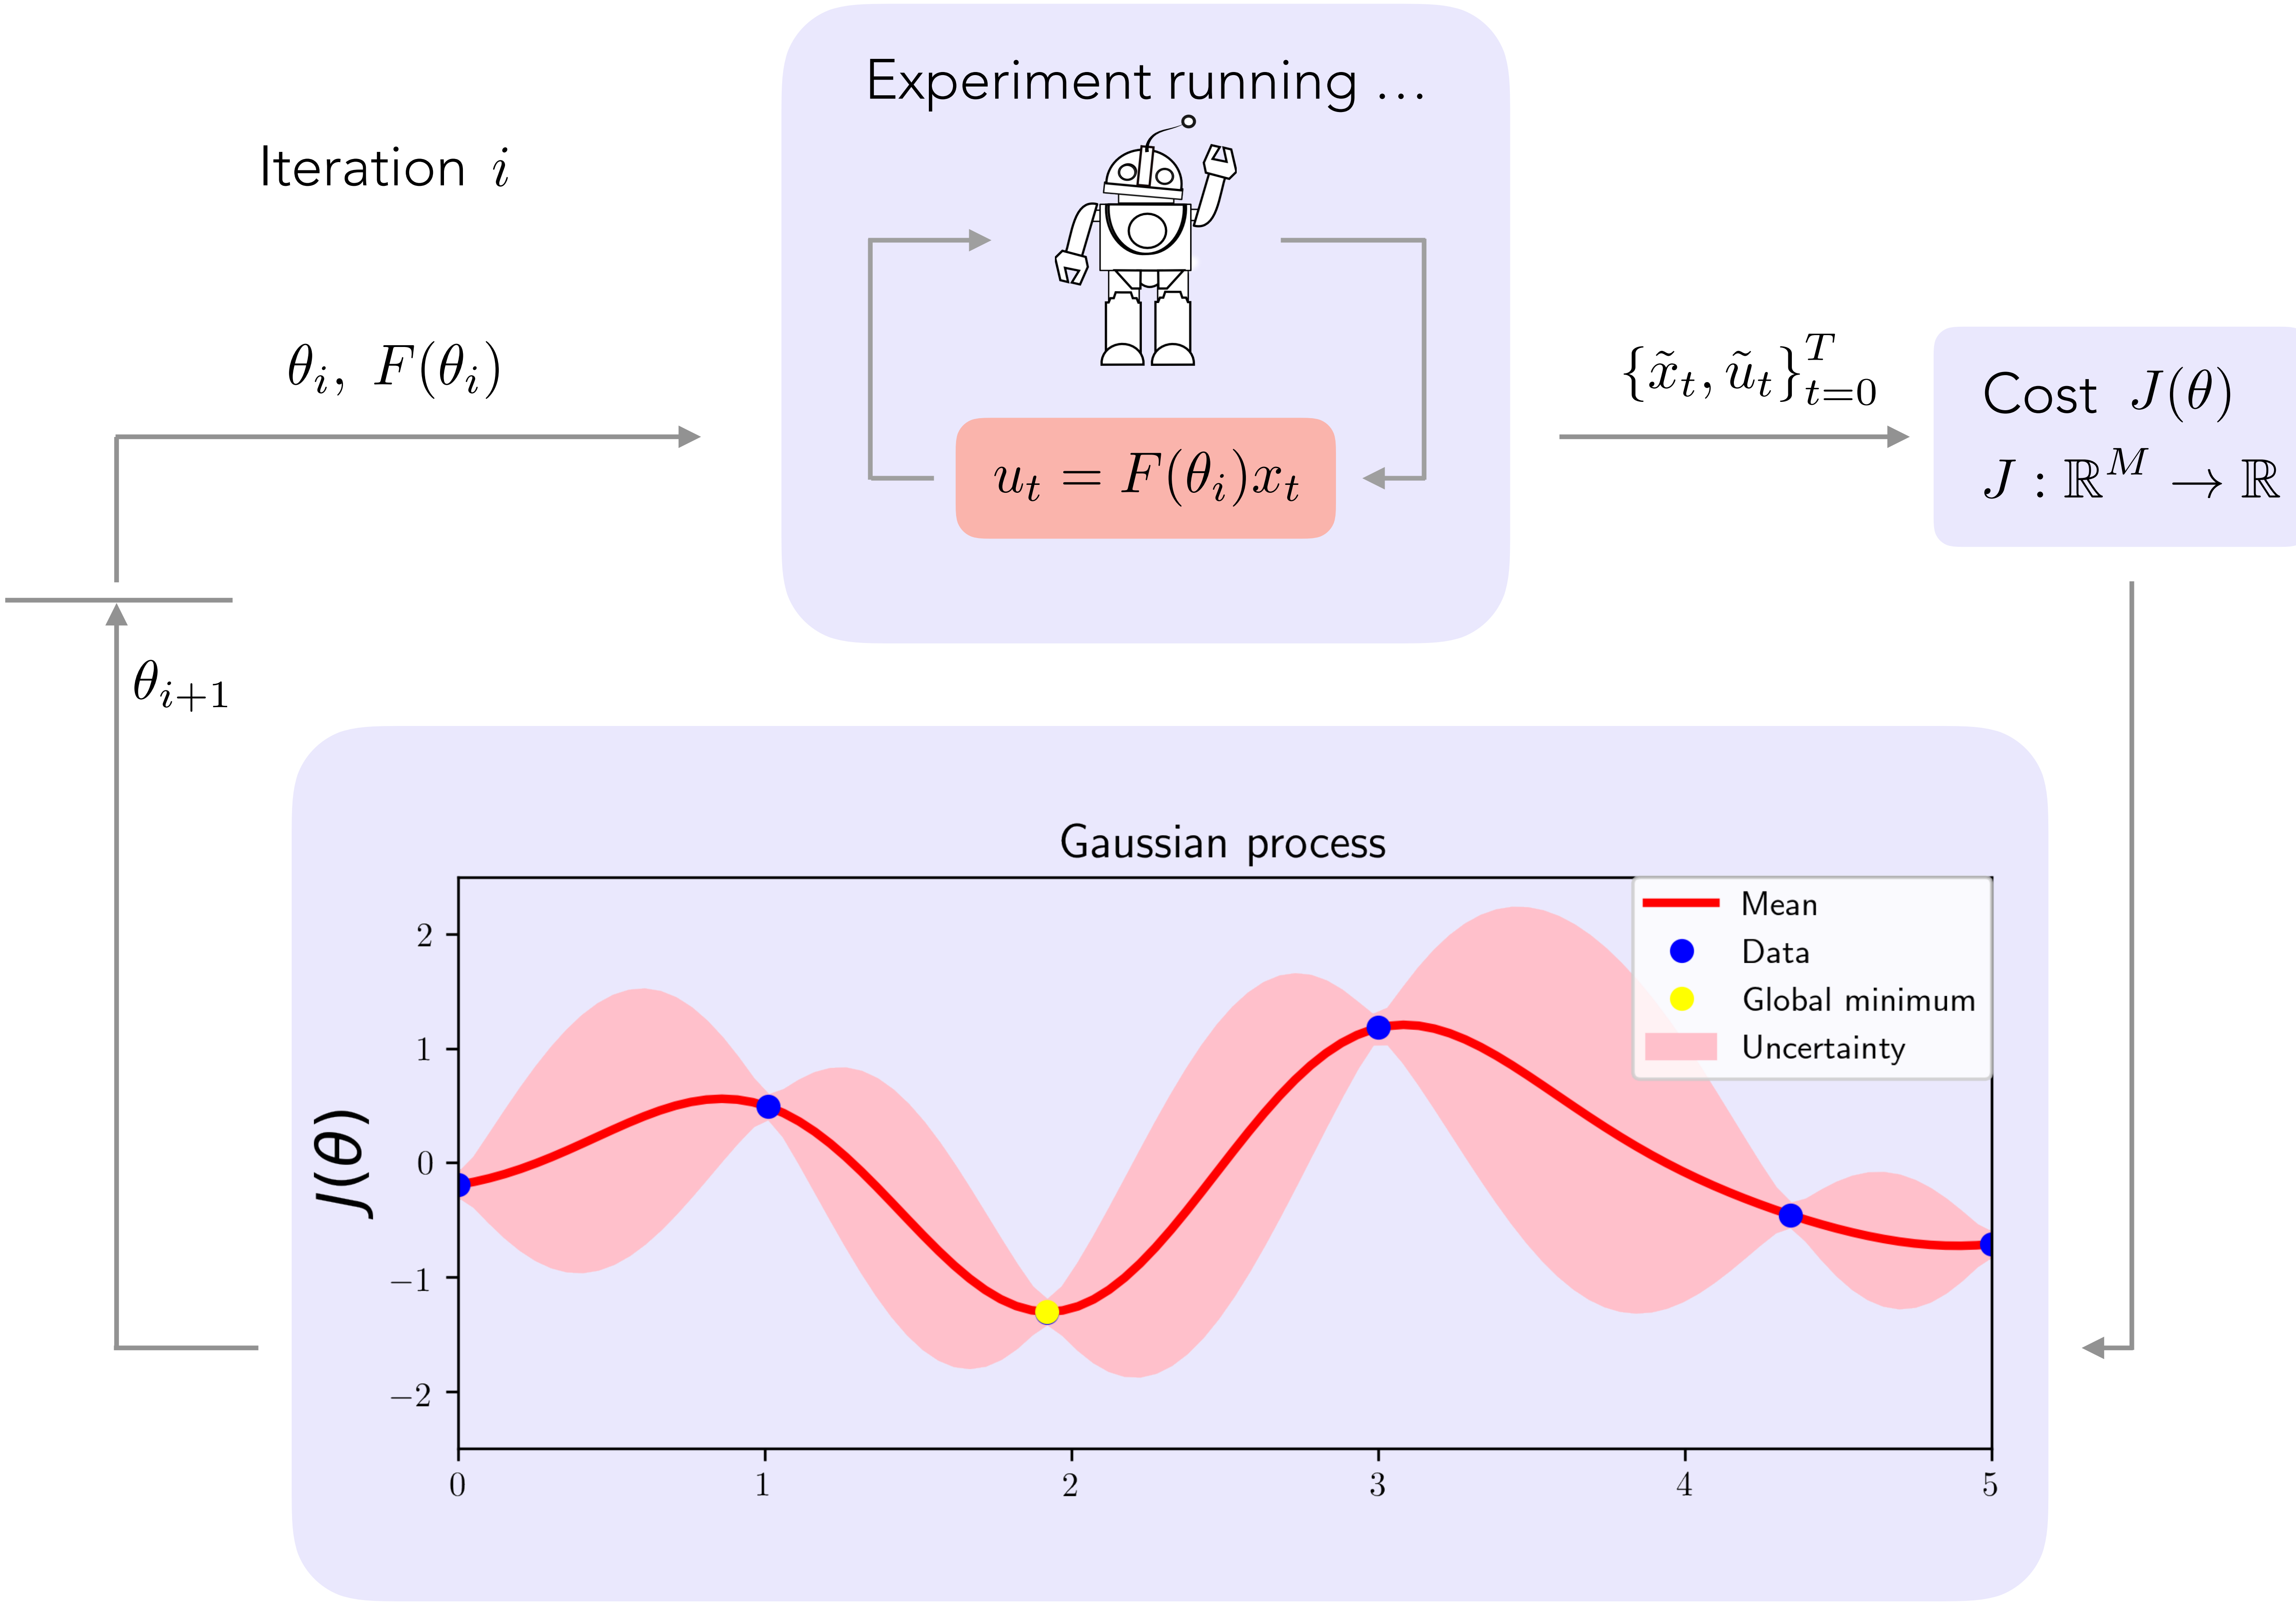
Iteration i
$\theta_i,\ F(\theta_i)$
Experiment running ...
$u_t = F(\theta_i) x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Global minimum
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5

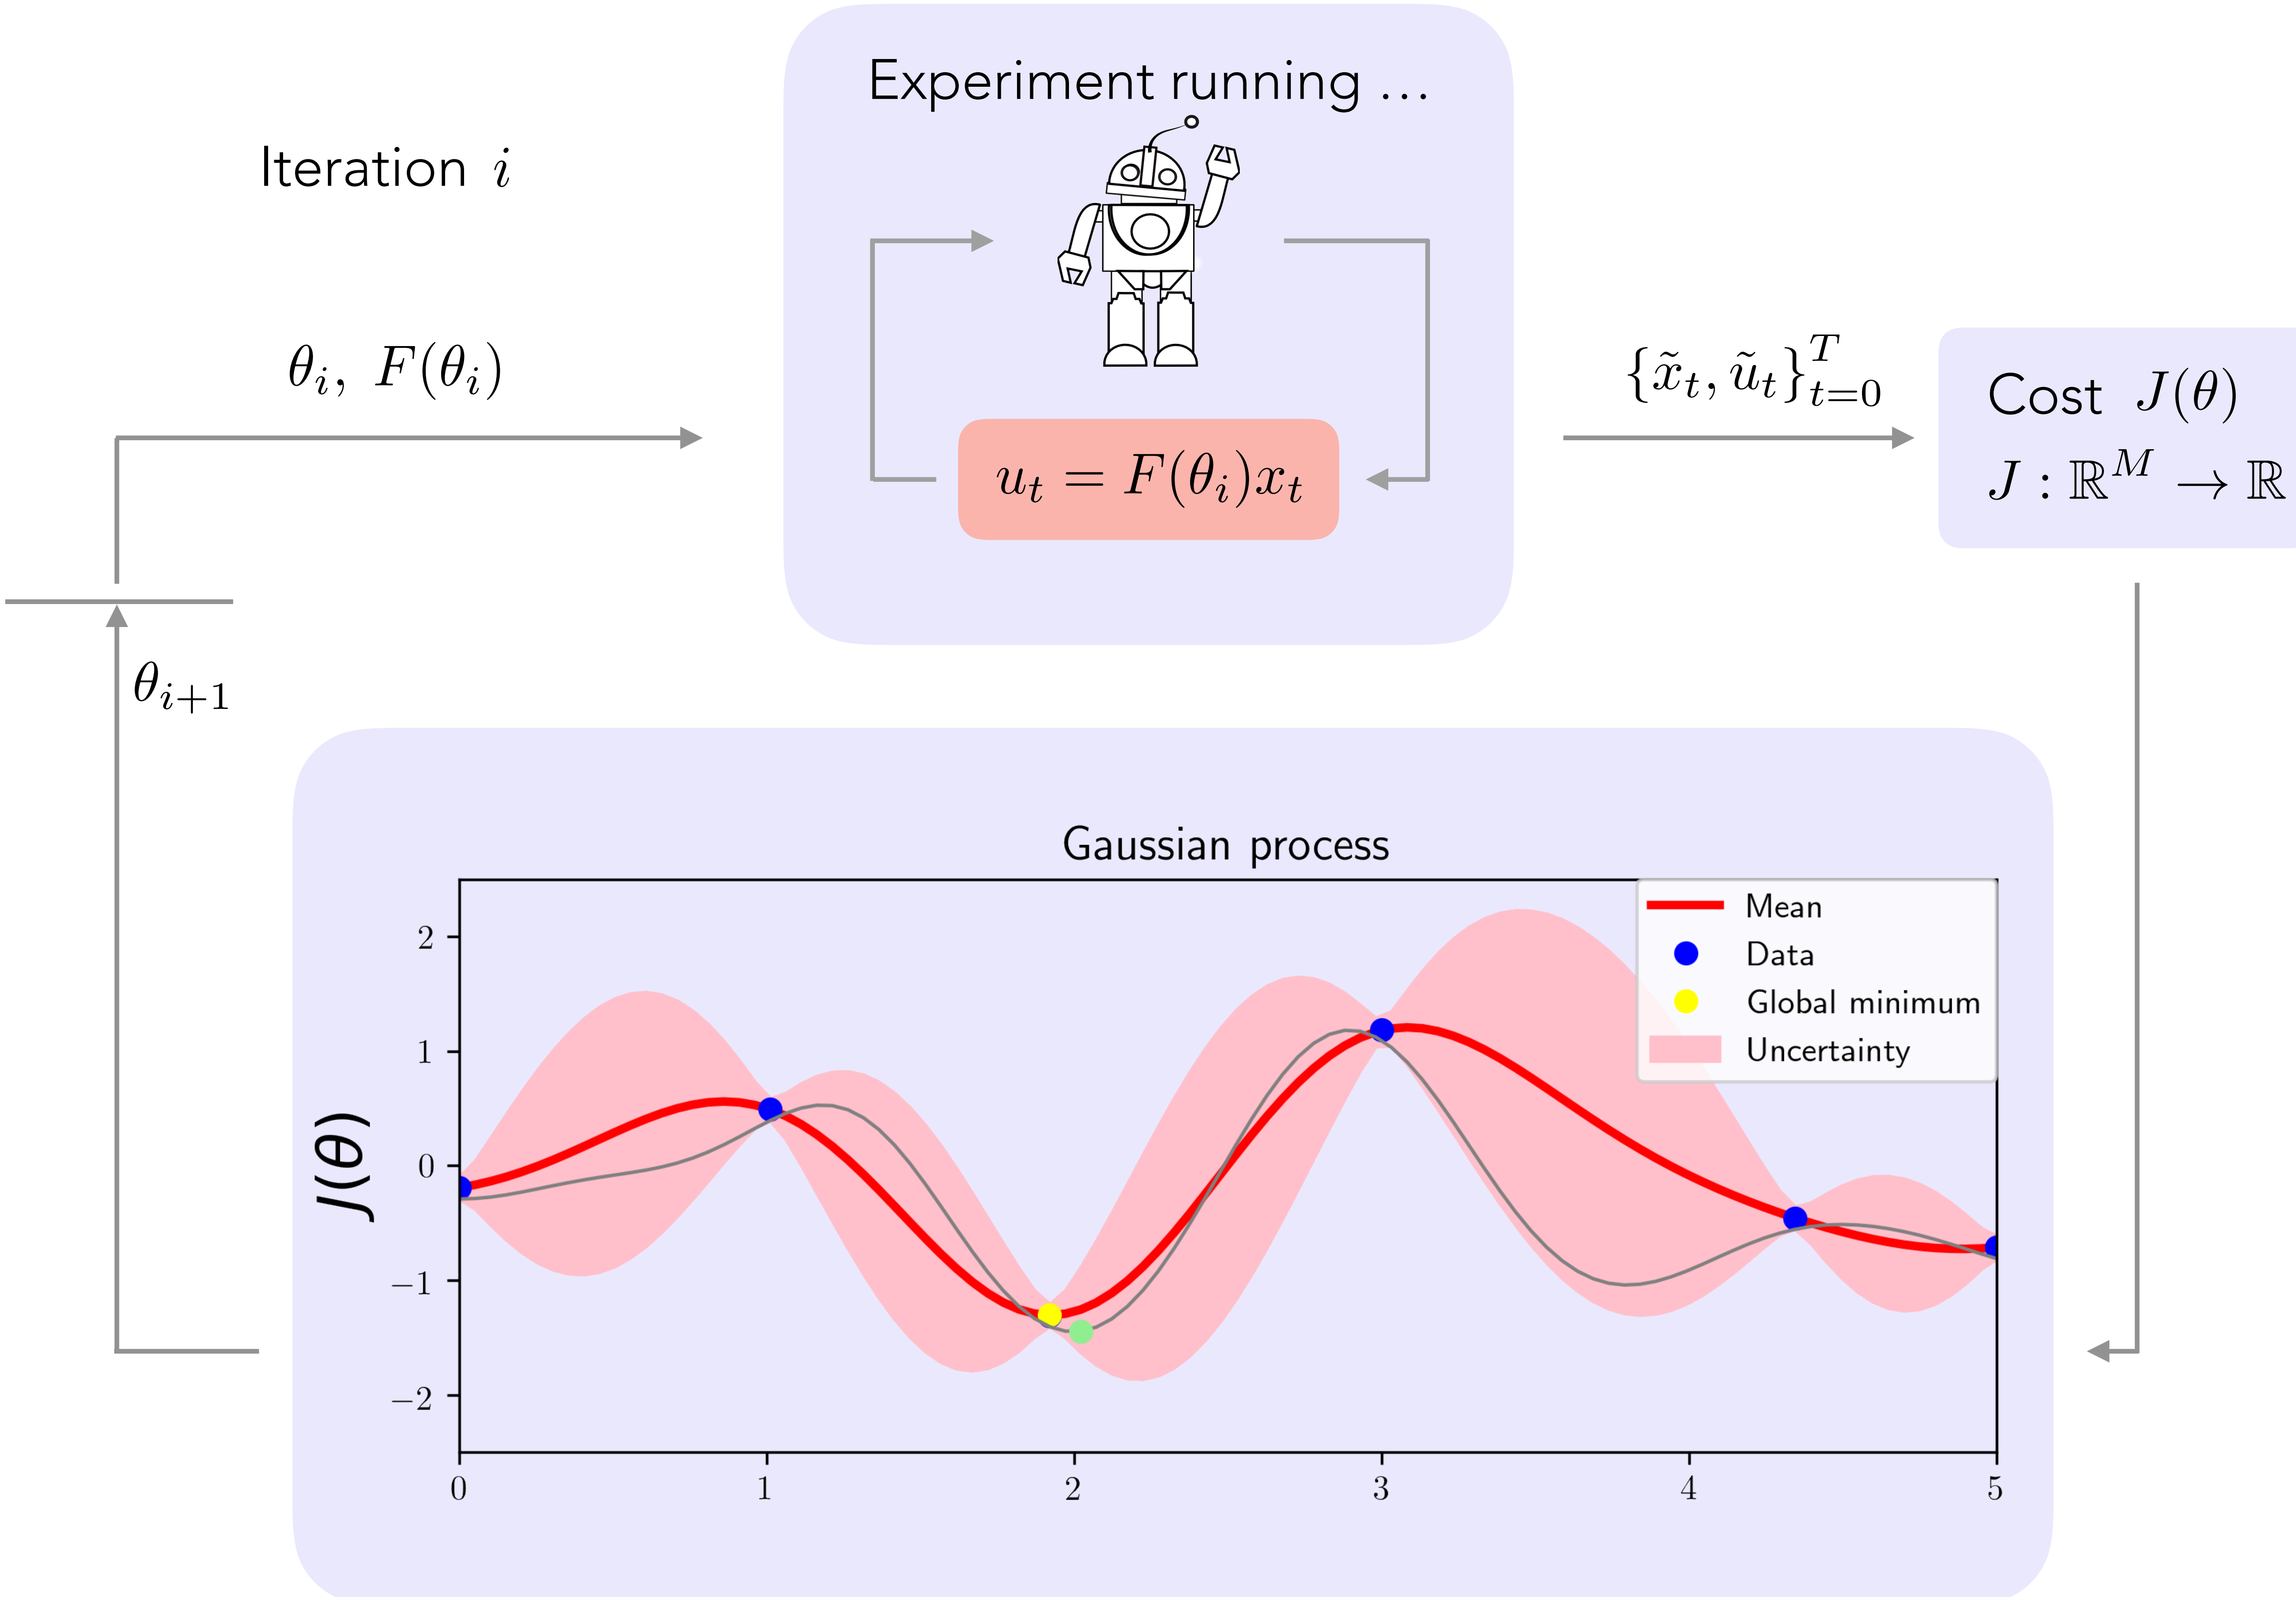
Experiment running ...
Iteration i
$\theta_i,\ F(\theta_i)$
$u_t = F(\theta_i)x_t$
$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$
Cost $J(\theta)$
$J : \mathbb{R}^M \to \mathbb{R}$
θ_{i+1}
Gaussian process
Mean
Data
Global minimum
Uncertainty
J(θ)
2
1
0
−1
−2
0
1
2
3
4
5

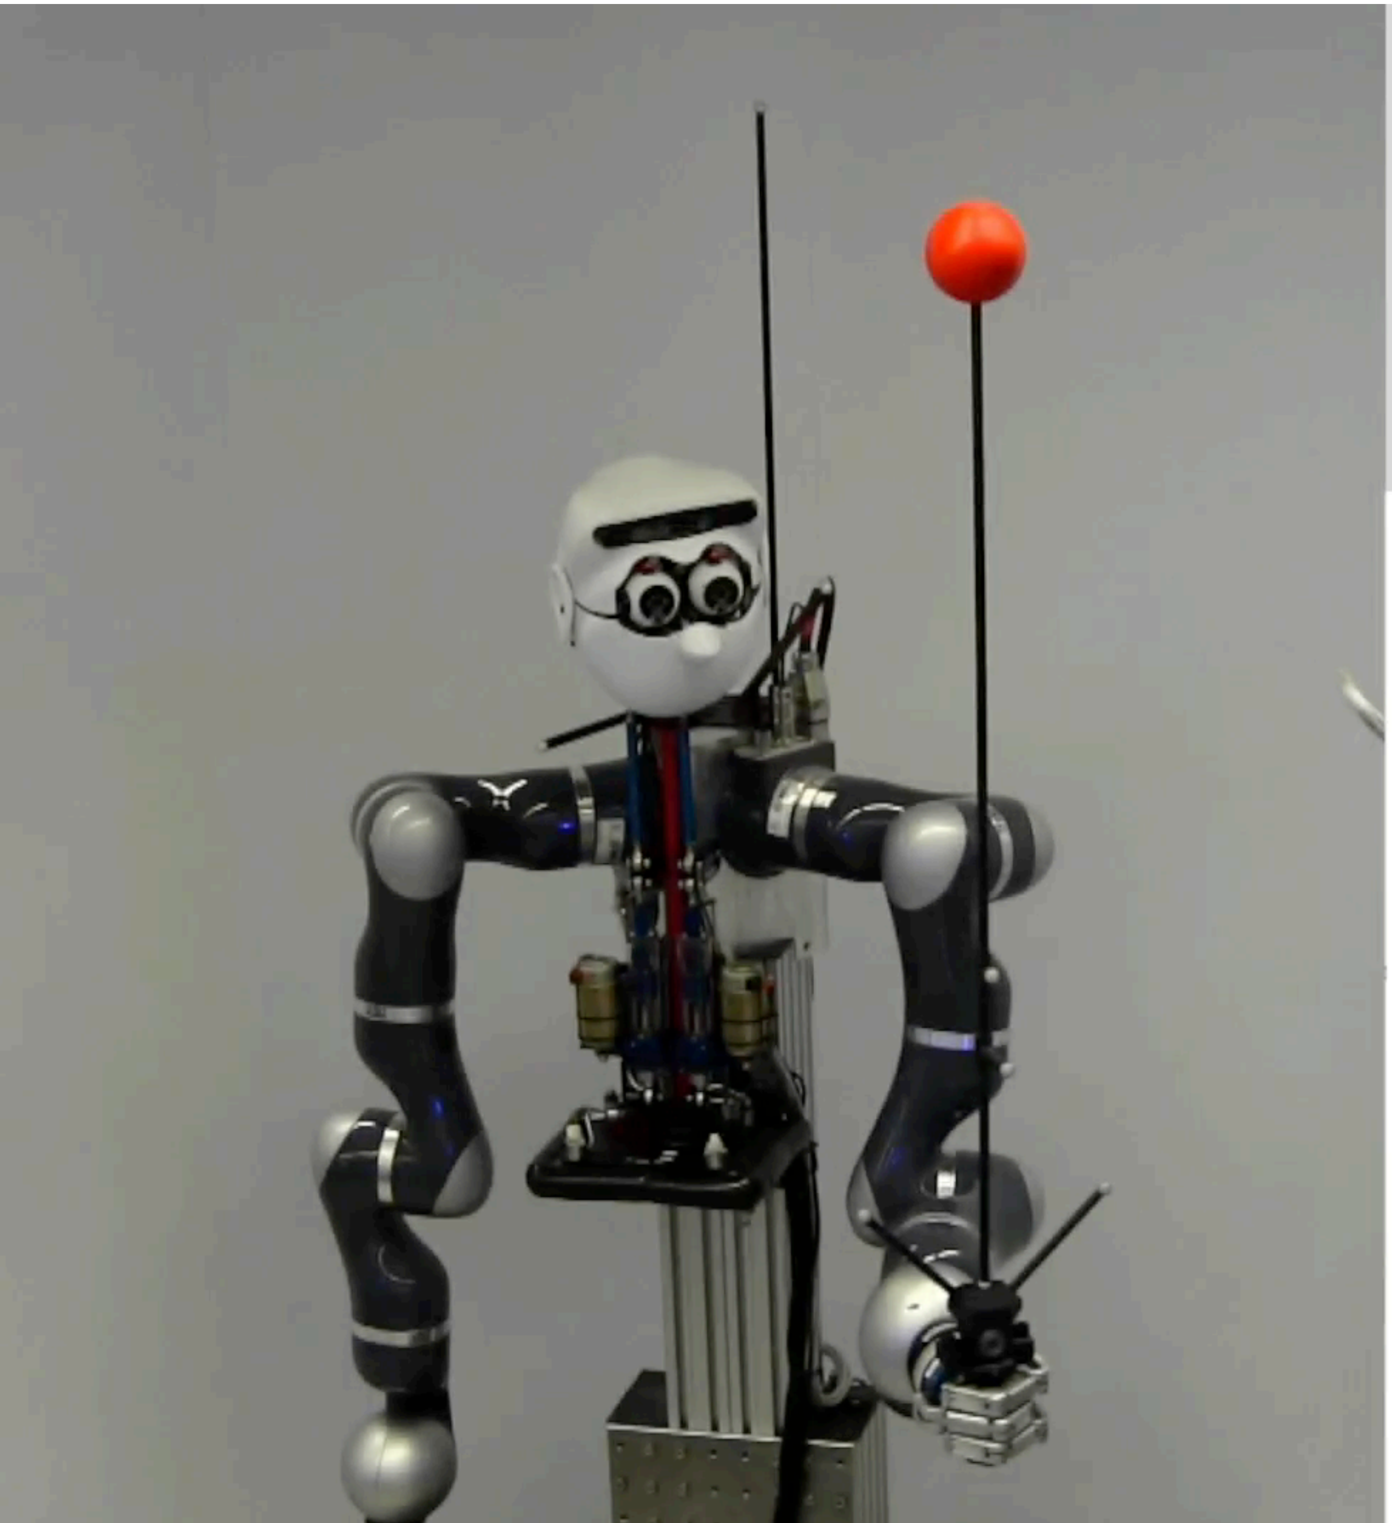

Experimental cost function captured by Gaussian process

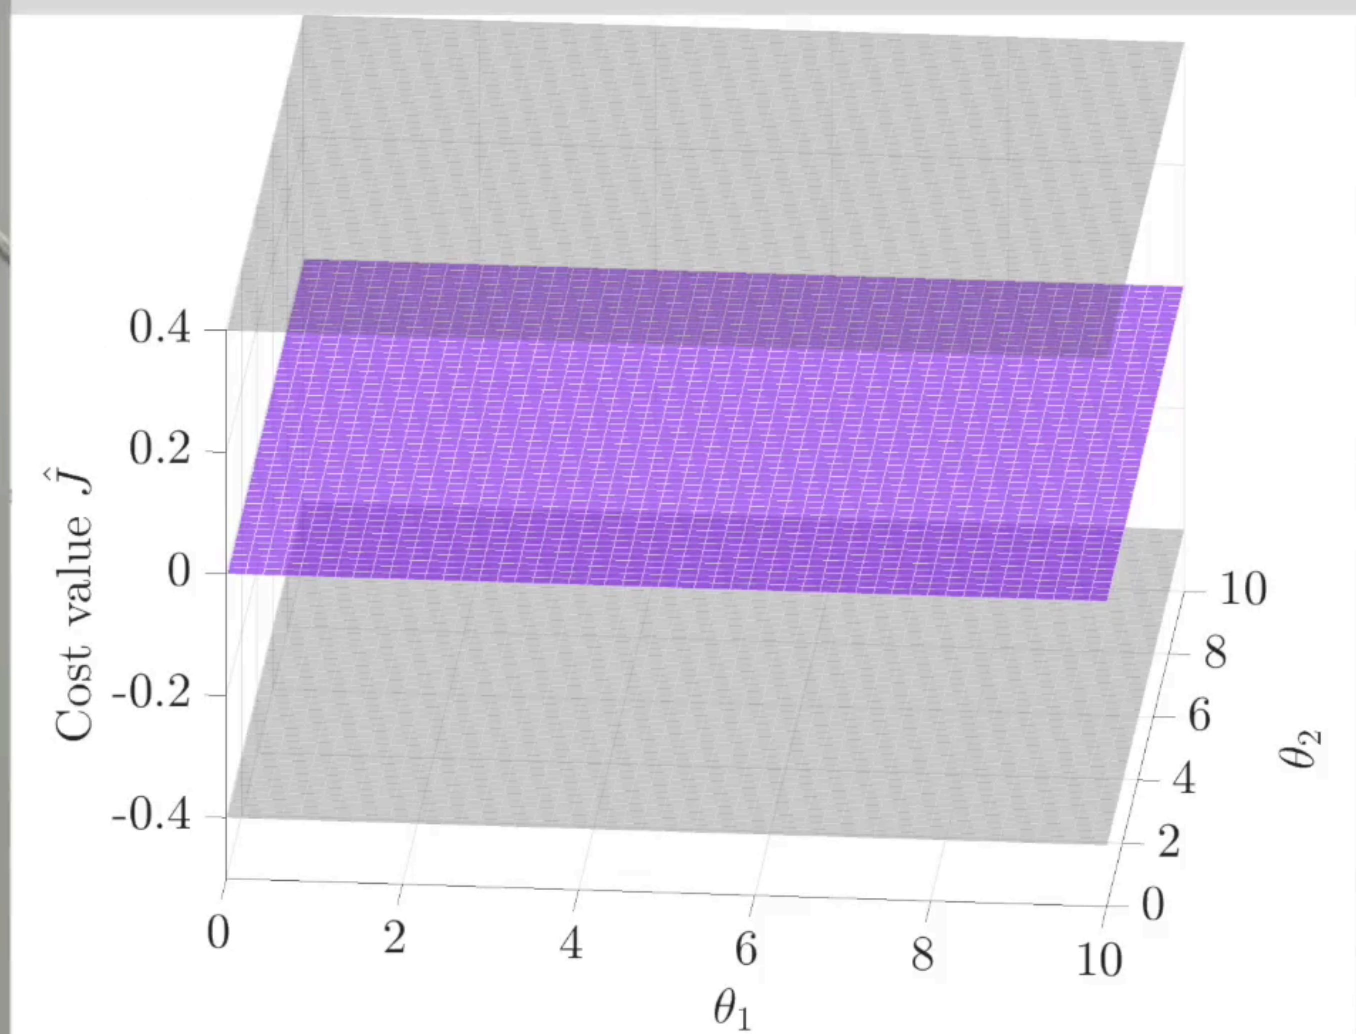

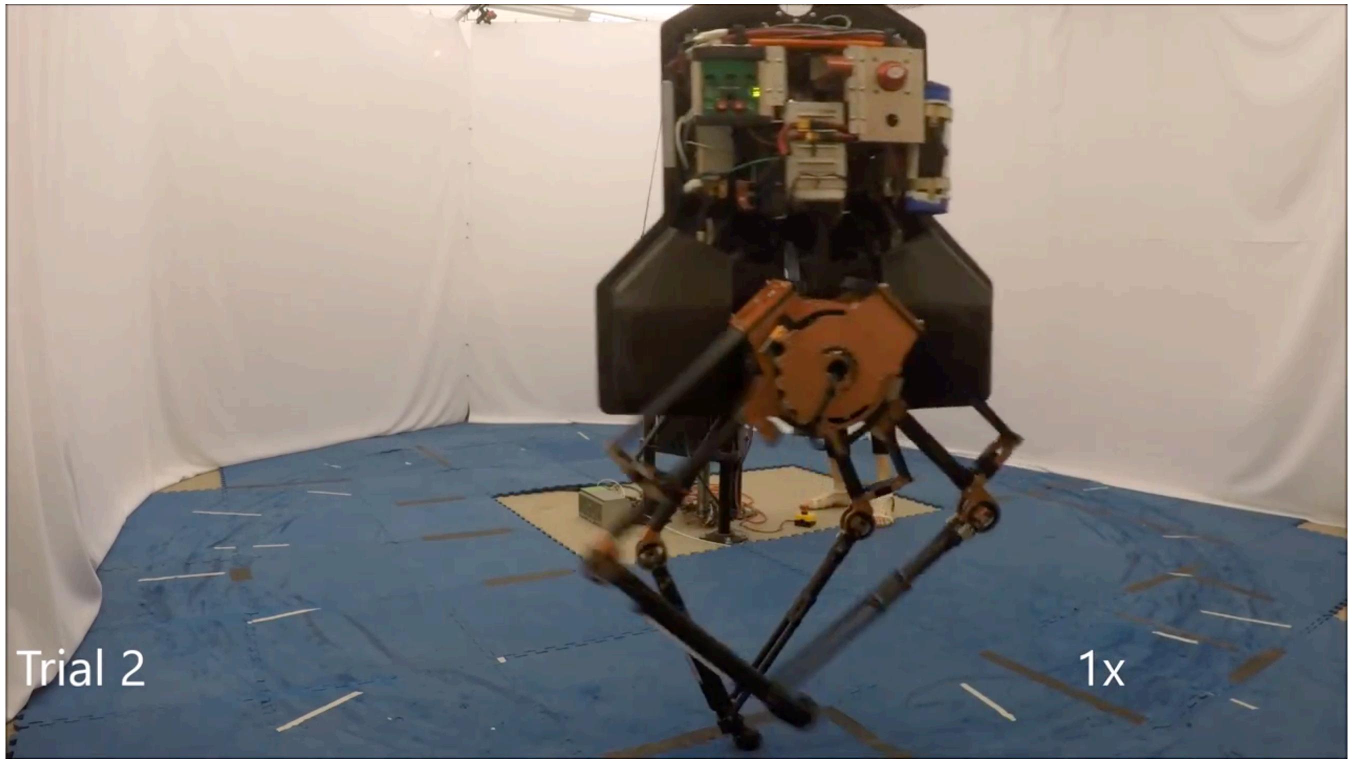

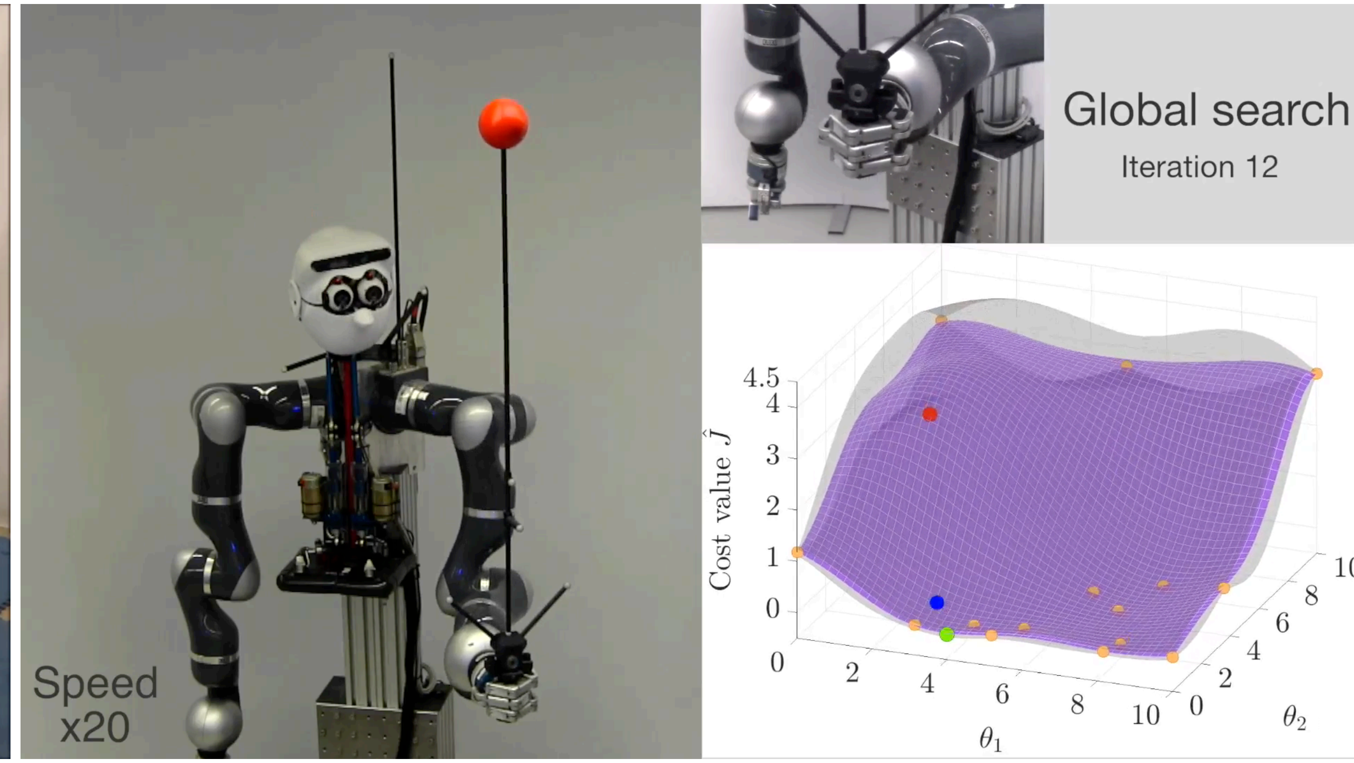

Antonova, R., et al., CoRL 2017

Marco. A, et al., ICRA 2016

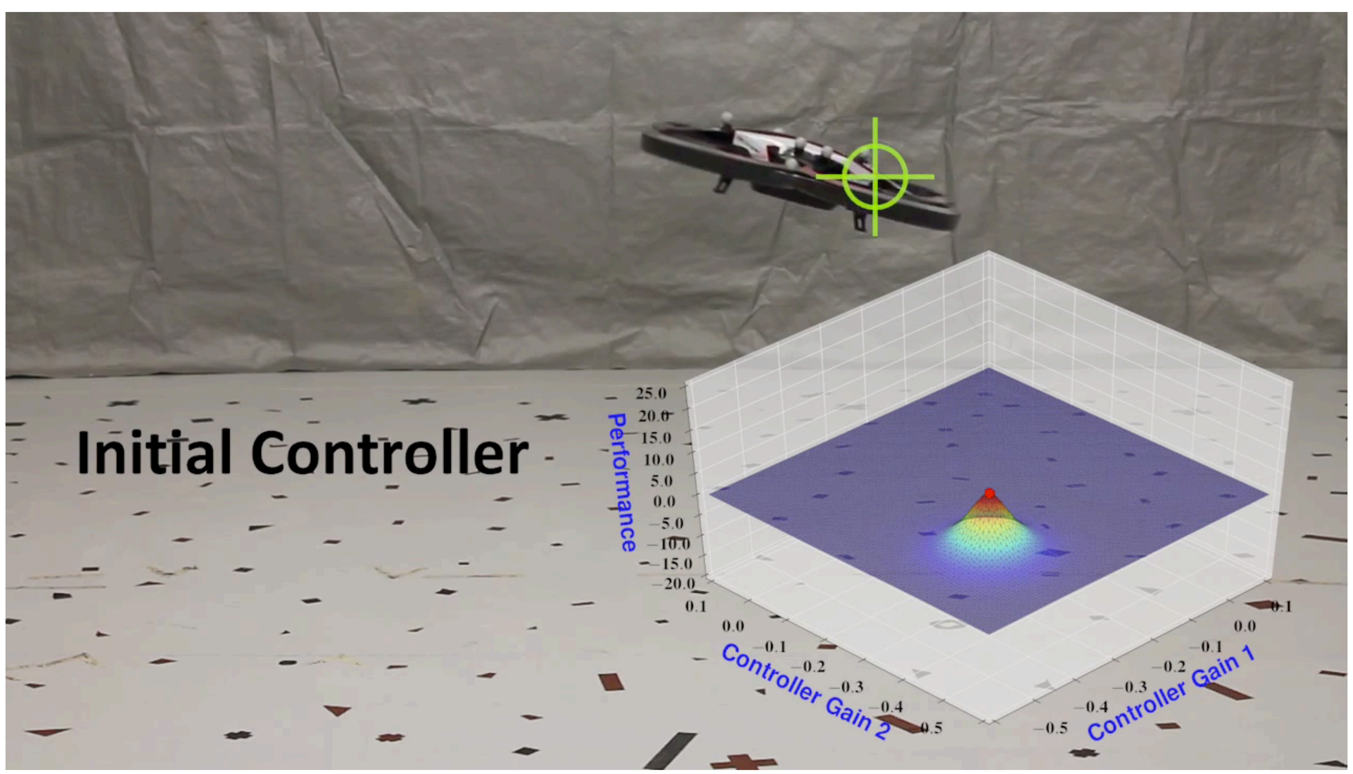

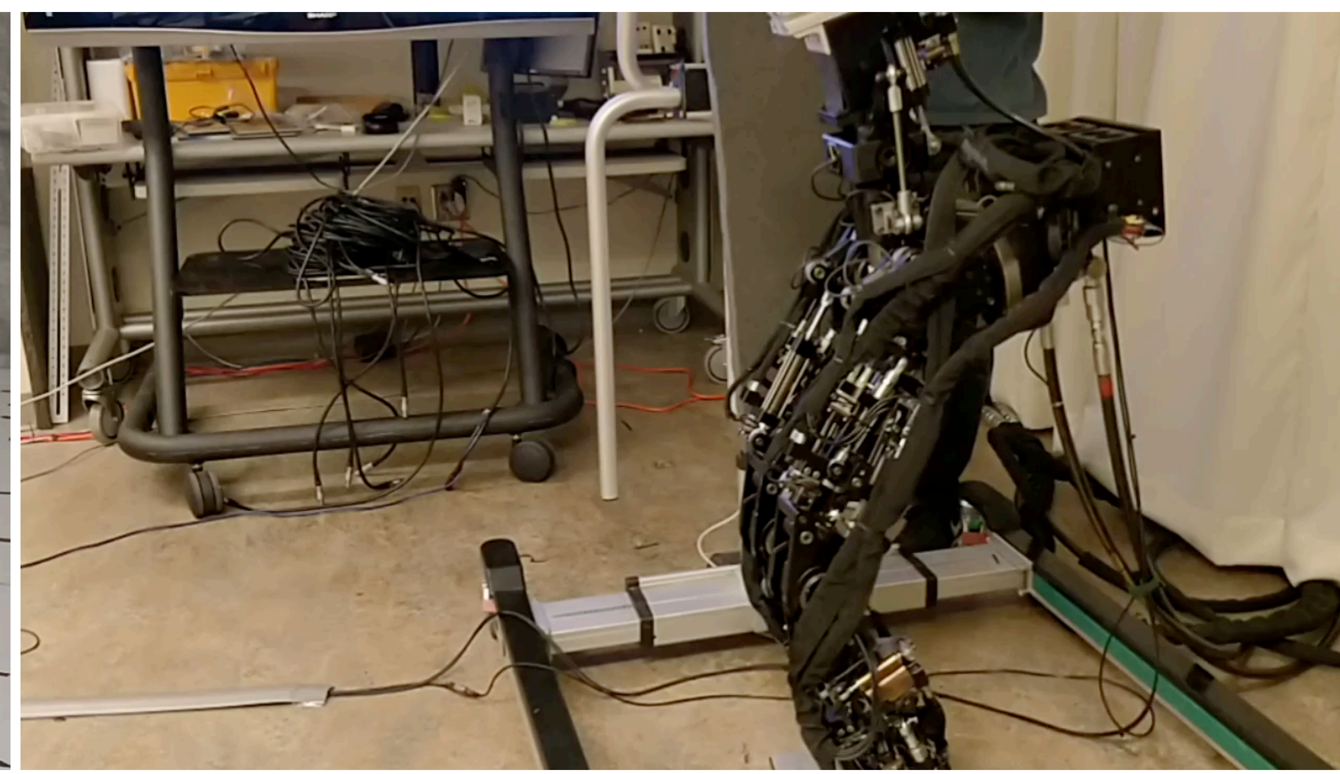

Berkenkamp, F., et al., ICRA 2017

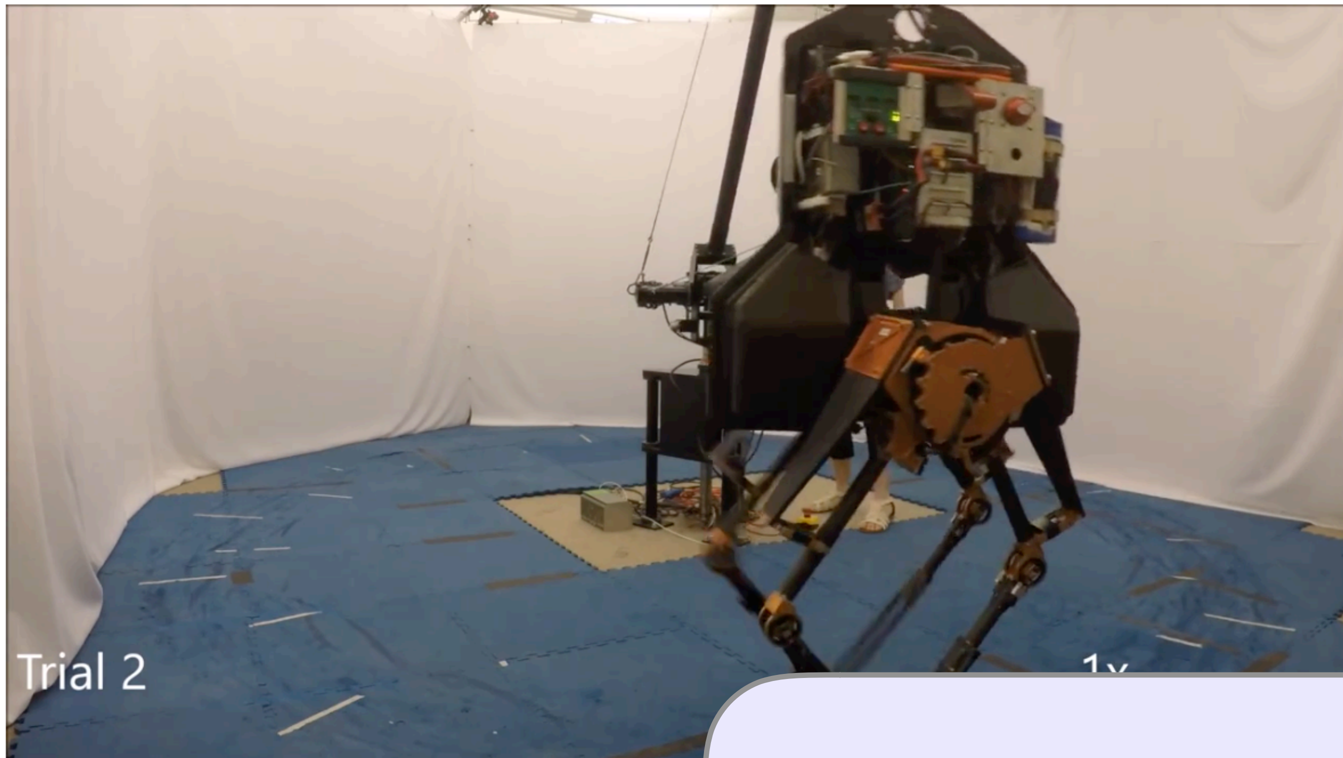

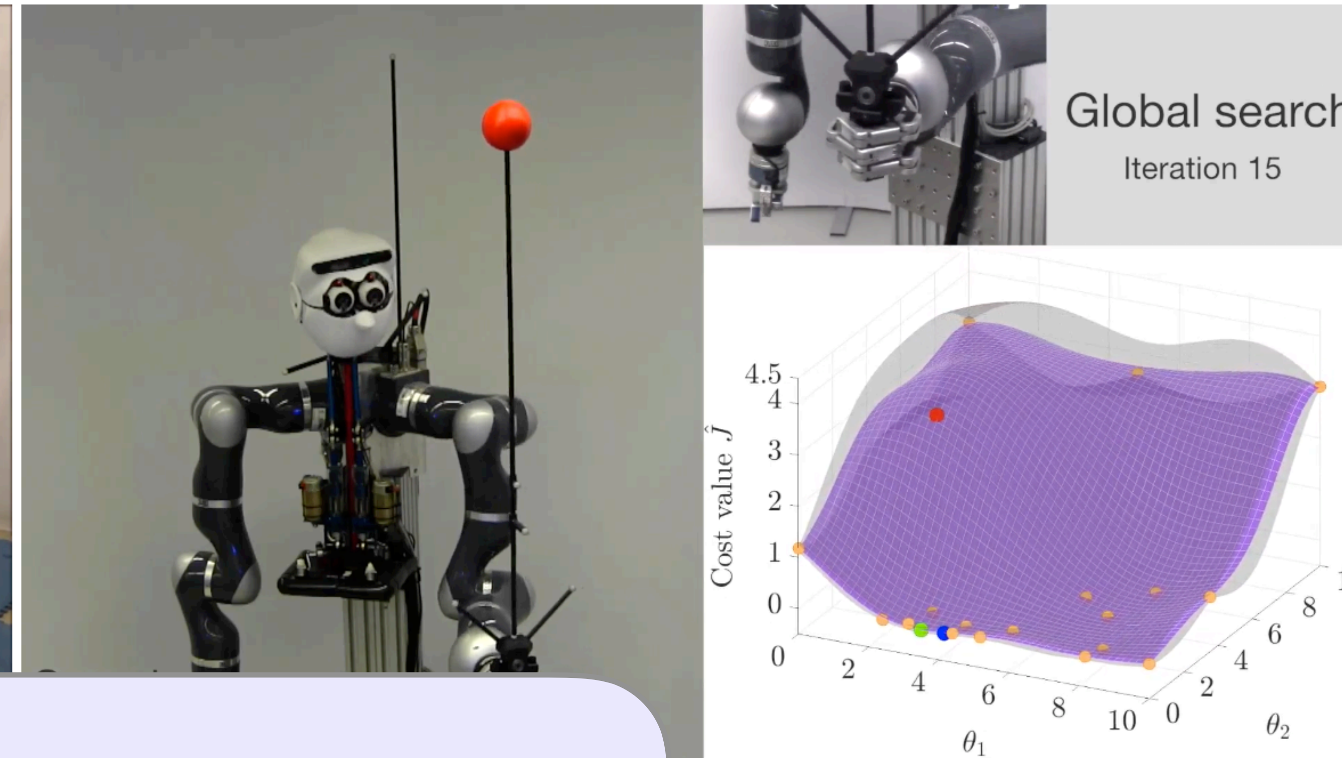

Antonova, R., et

et al., ICRA 2016

GP unaware of controller structure

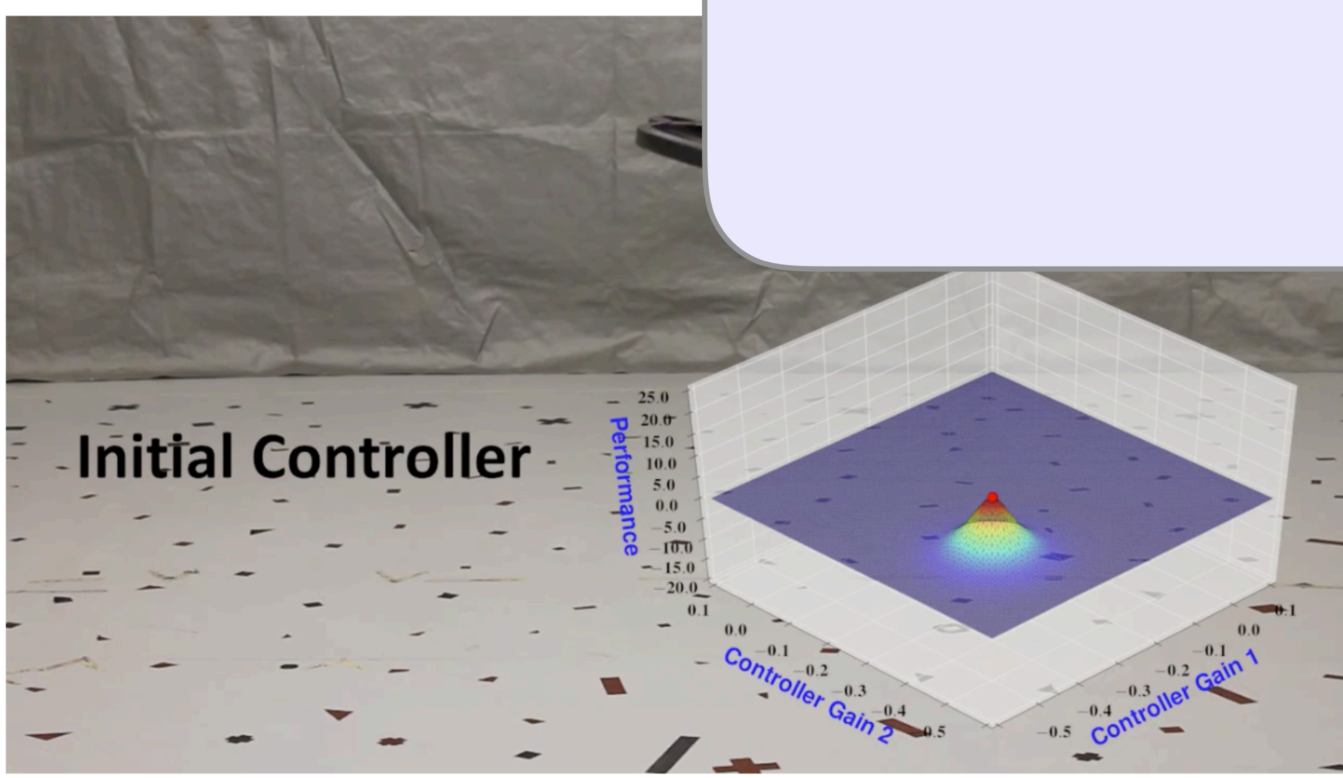

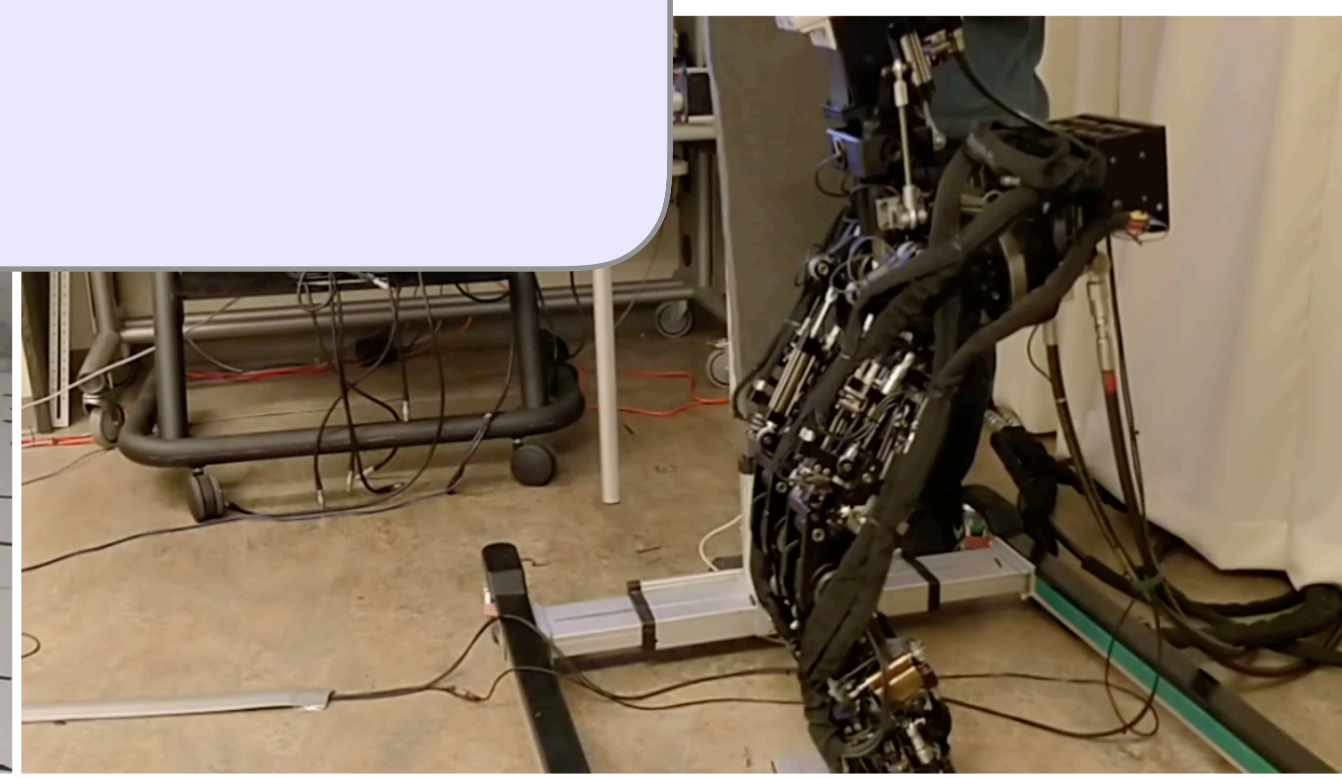

Berkenkamp, F., et al., ICRA 2017

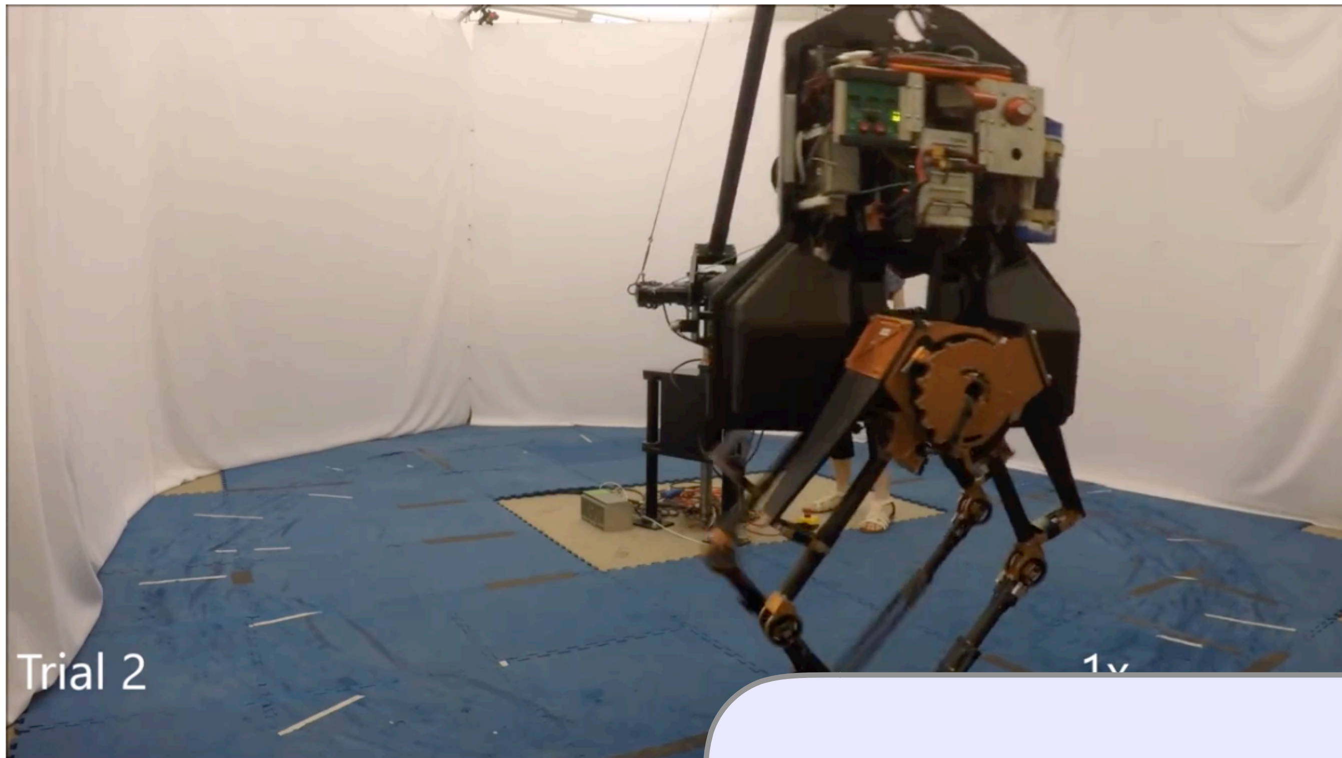

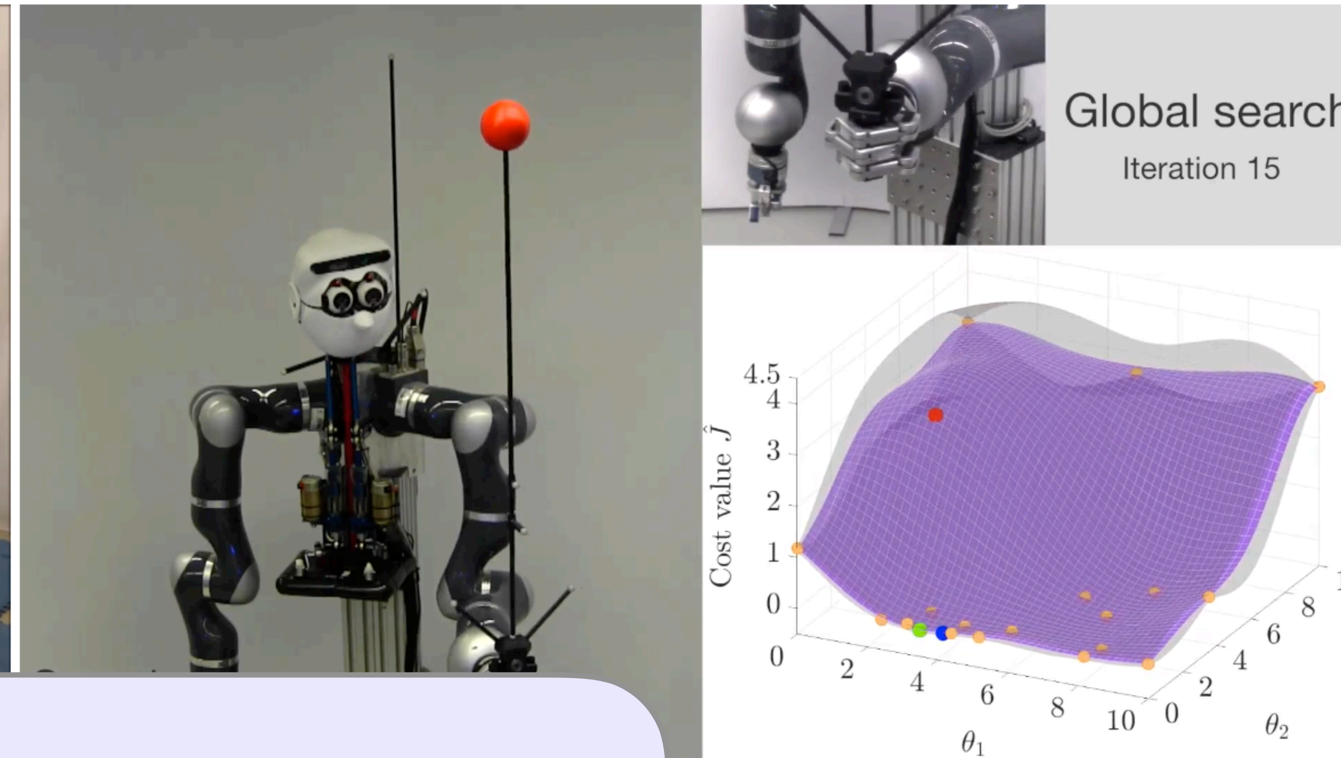

Antonova, R., et

et al., ICRA 2016

GP unaware of controller structure

Incorporate controller structure in GP

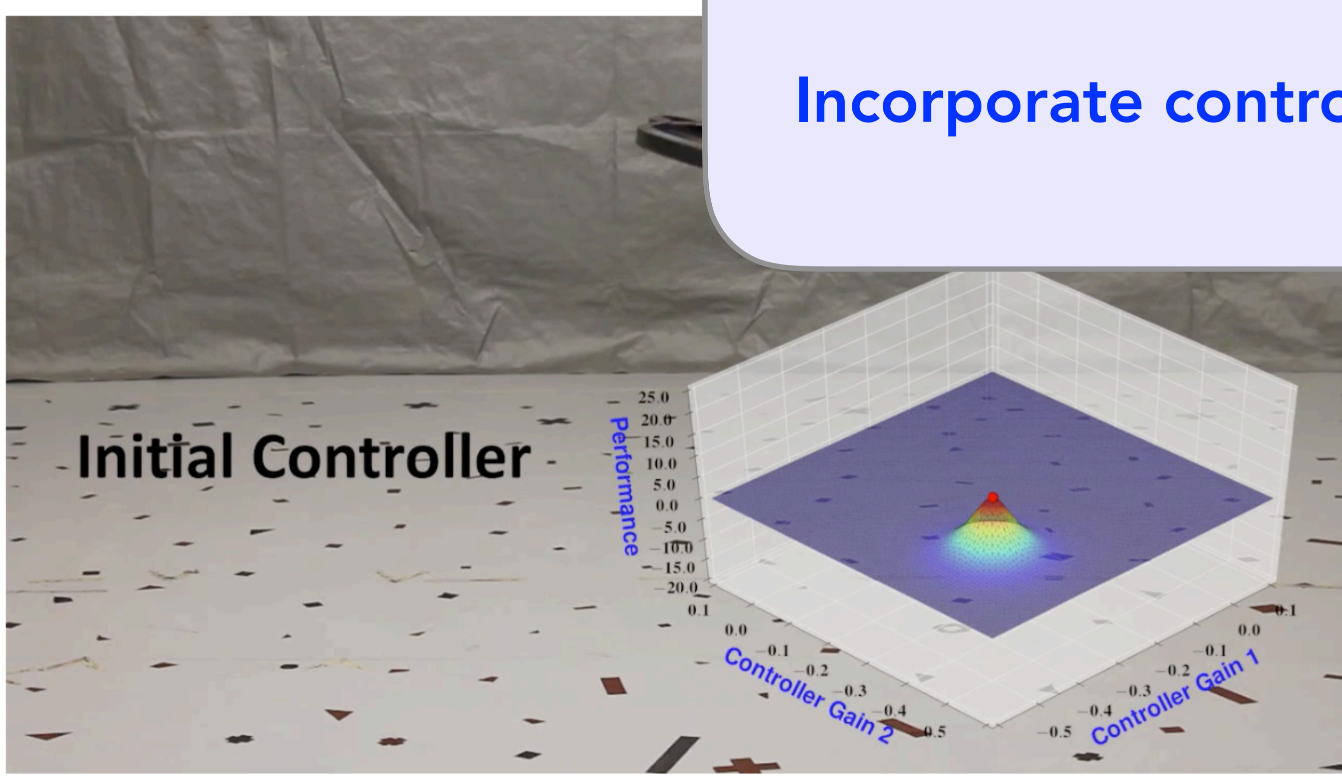

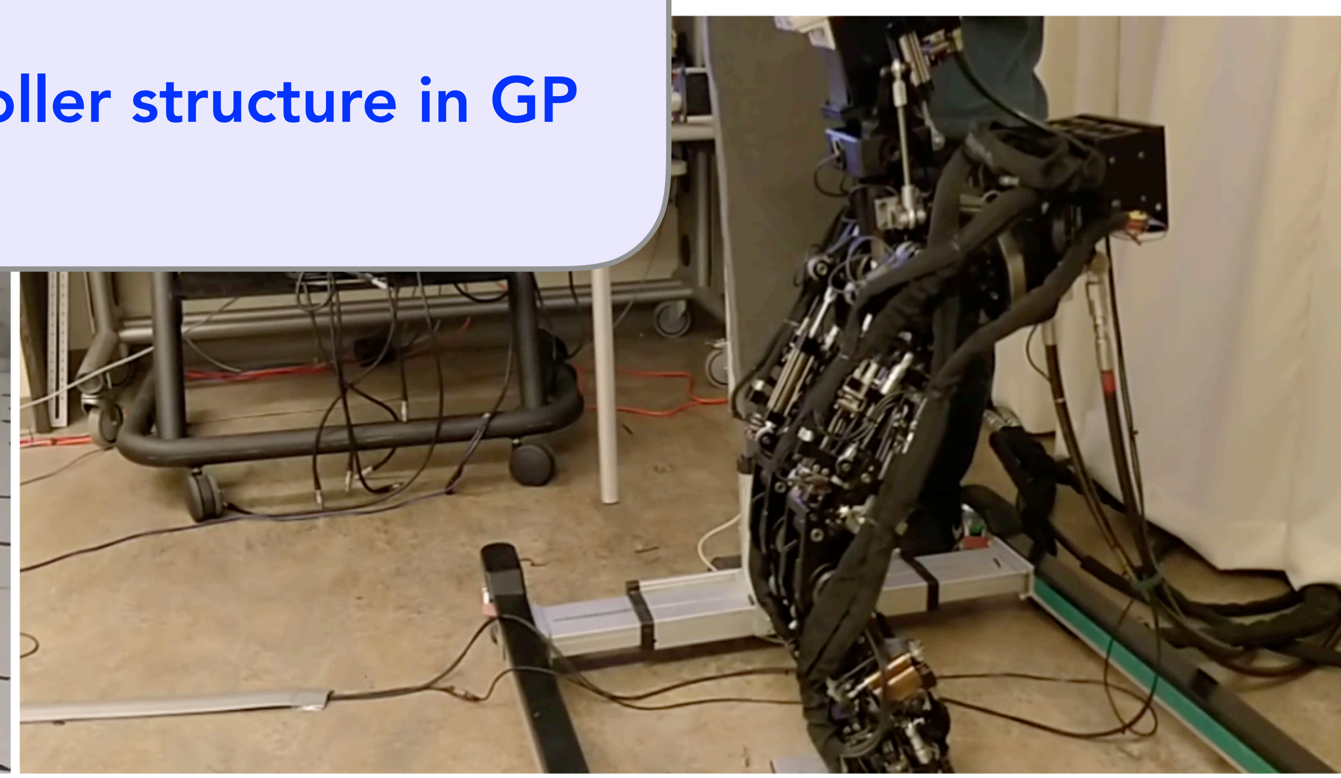

Berkenkamp, F., et al., ICRA 2017

✳Gaussian process: mean and kernel

$$J(\theta) \sim \mathcal{GP}(\mu(\theta), {\color{red}k(\theta,\theta')})$$

$$\mu(\theta) = \mathbb{E}\left[J(\theta)\right]$$

$${\color{red}k(\theta,\theta')} = \mathbb{C}\mathrm{ov}\left[J(\theta), J(\theta')\right]$$

$\theta_i,\ F(\theta_i)$

$u_t = F(\theta_i)x_t$

$\{\tilde{x}_t, \tilde{u}_t\}_{t=0}^{T}$

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

θ_{i+1}

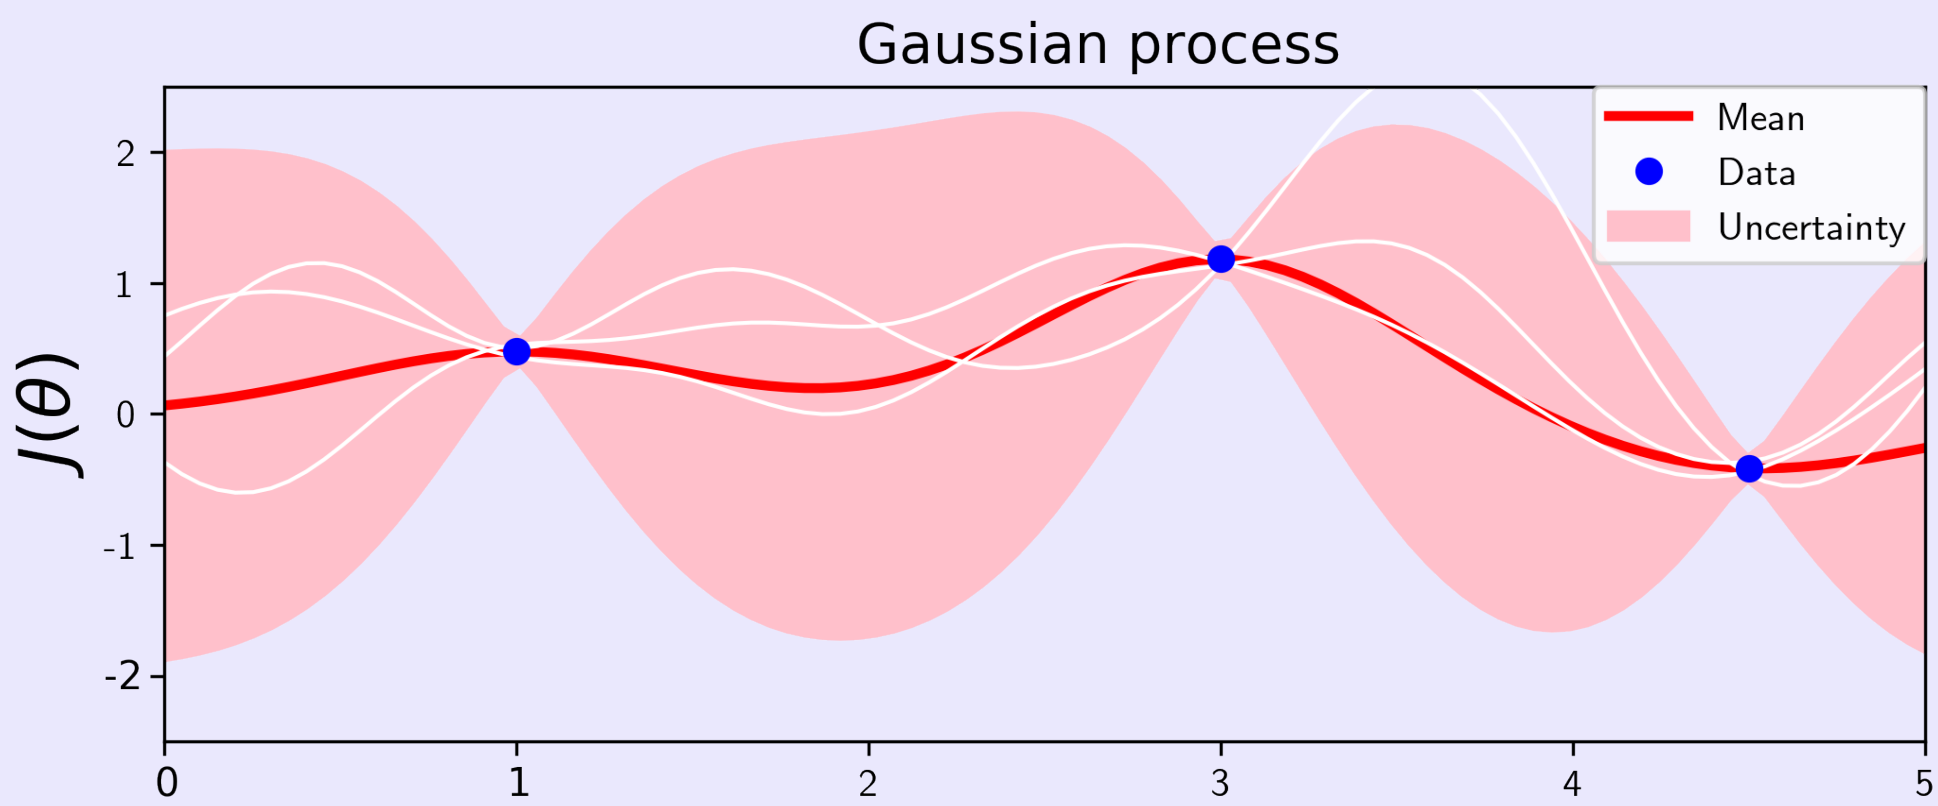

Alonso Marco | Max Planck Institute for Intelligent Systems

✴kernel encodes the **type of functions** we expect

Smooth functions

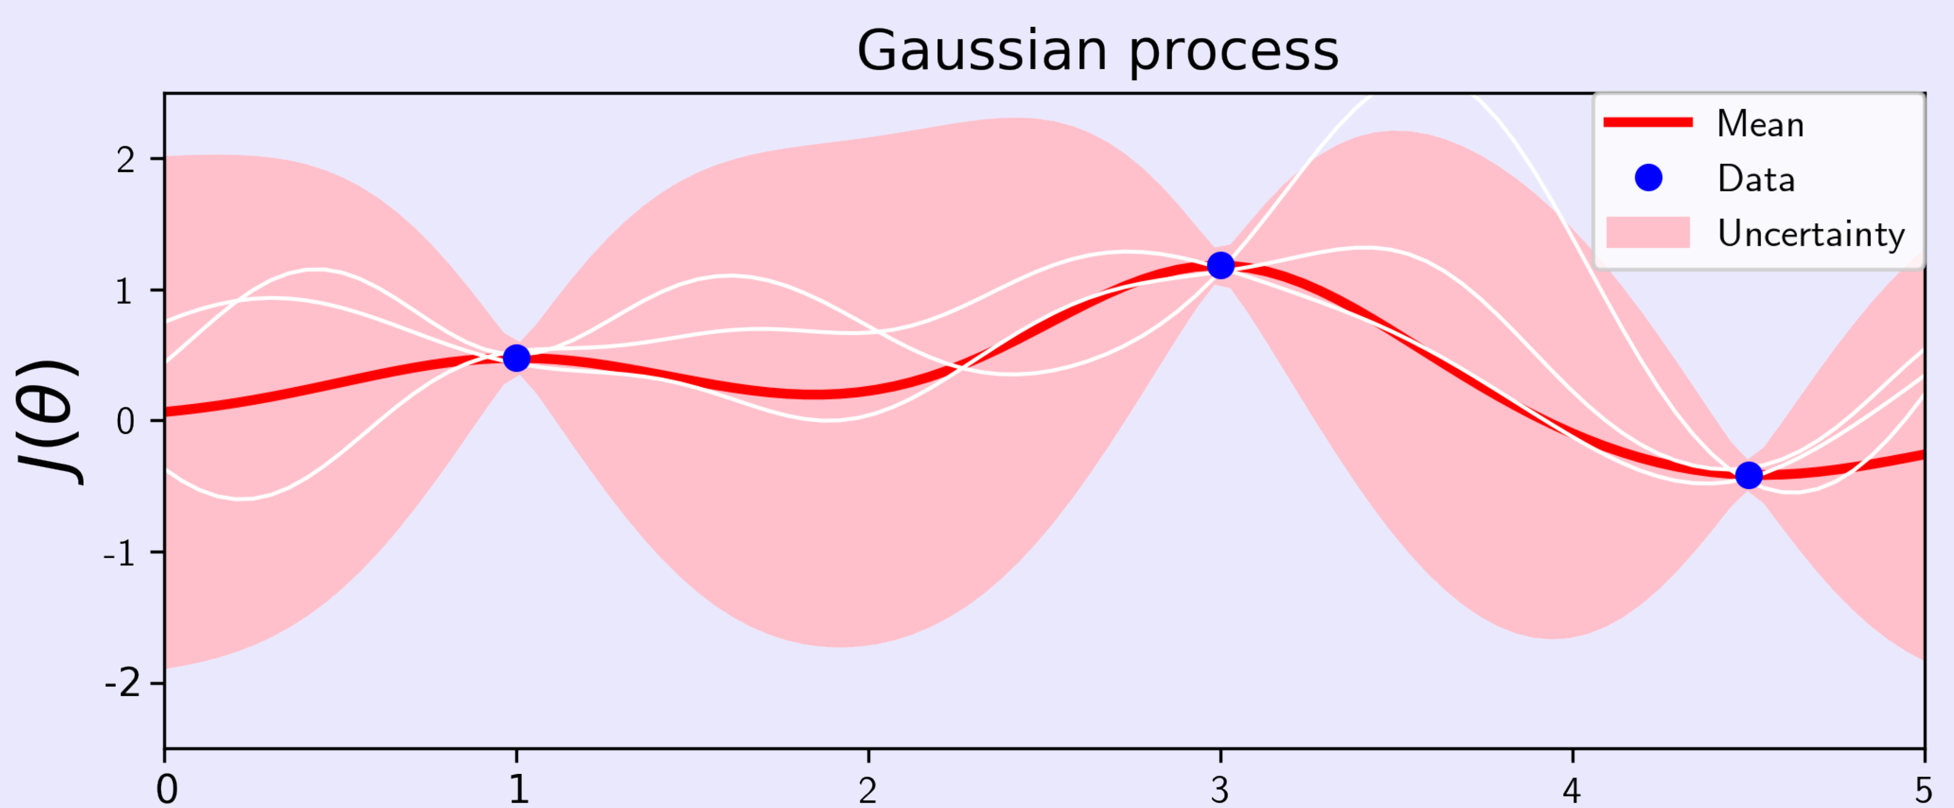

Alonso Marco | Max Planck Institute for Intelligent Systems

✴kernel encodes the **type of functions** we expect

Sharp functions

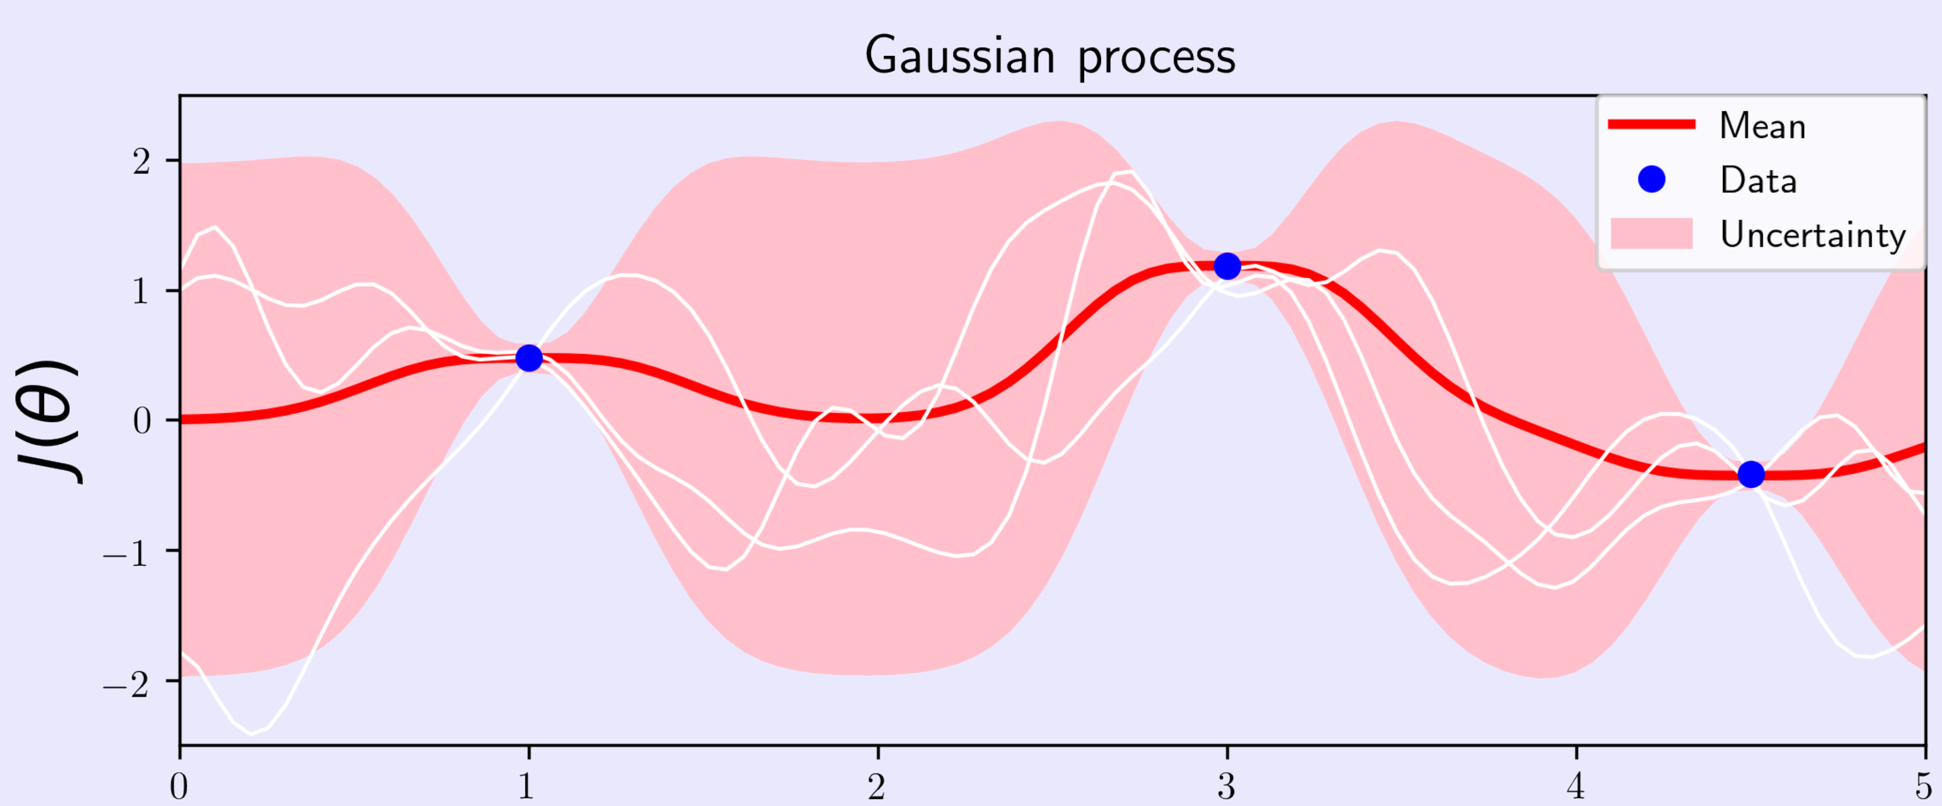

✴kernel encodes the **type of functions** we expect

Rapidly-changing functions

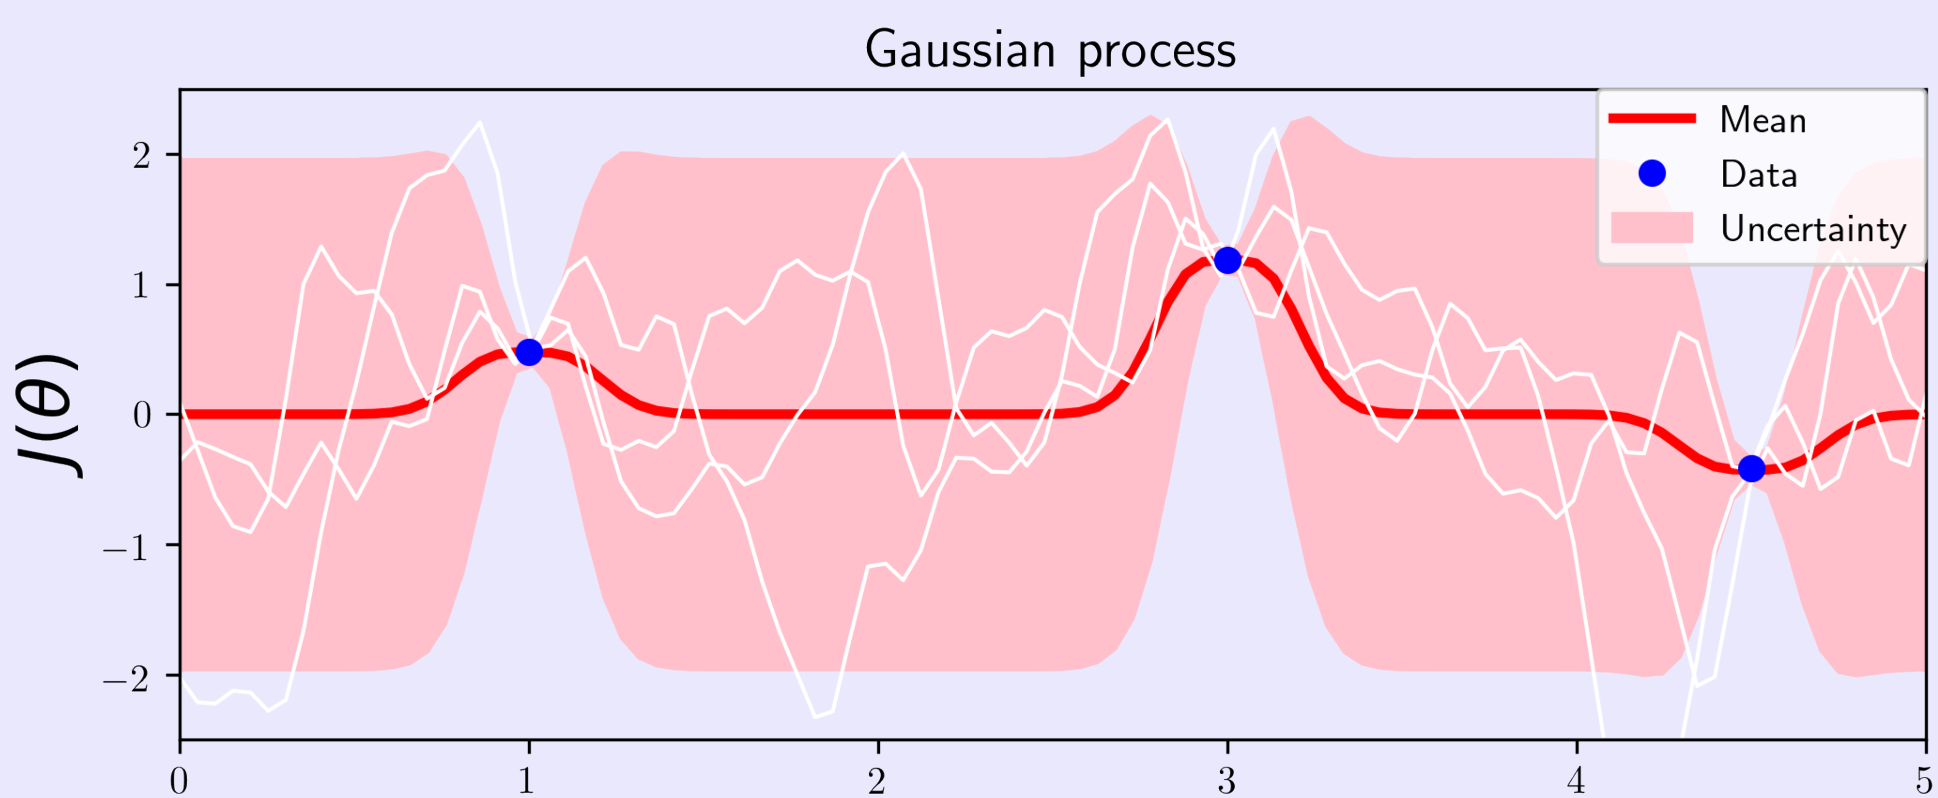

Alonso Marco | Max Planck Institute for Intelligent Systems

Which kernel describes my data better?

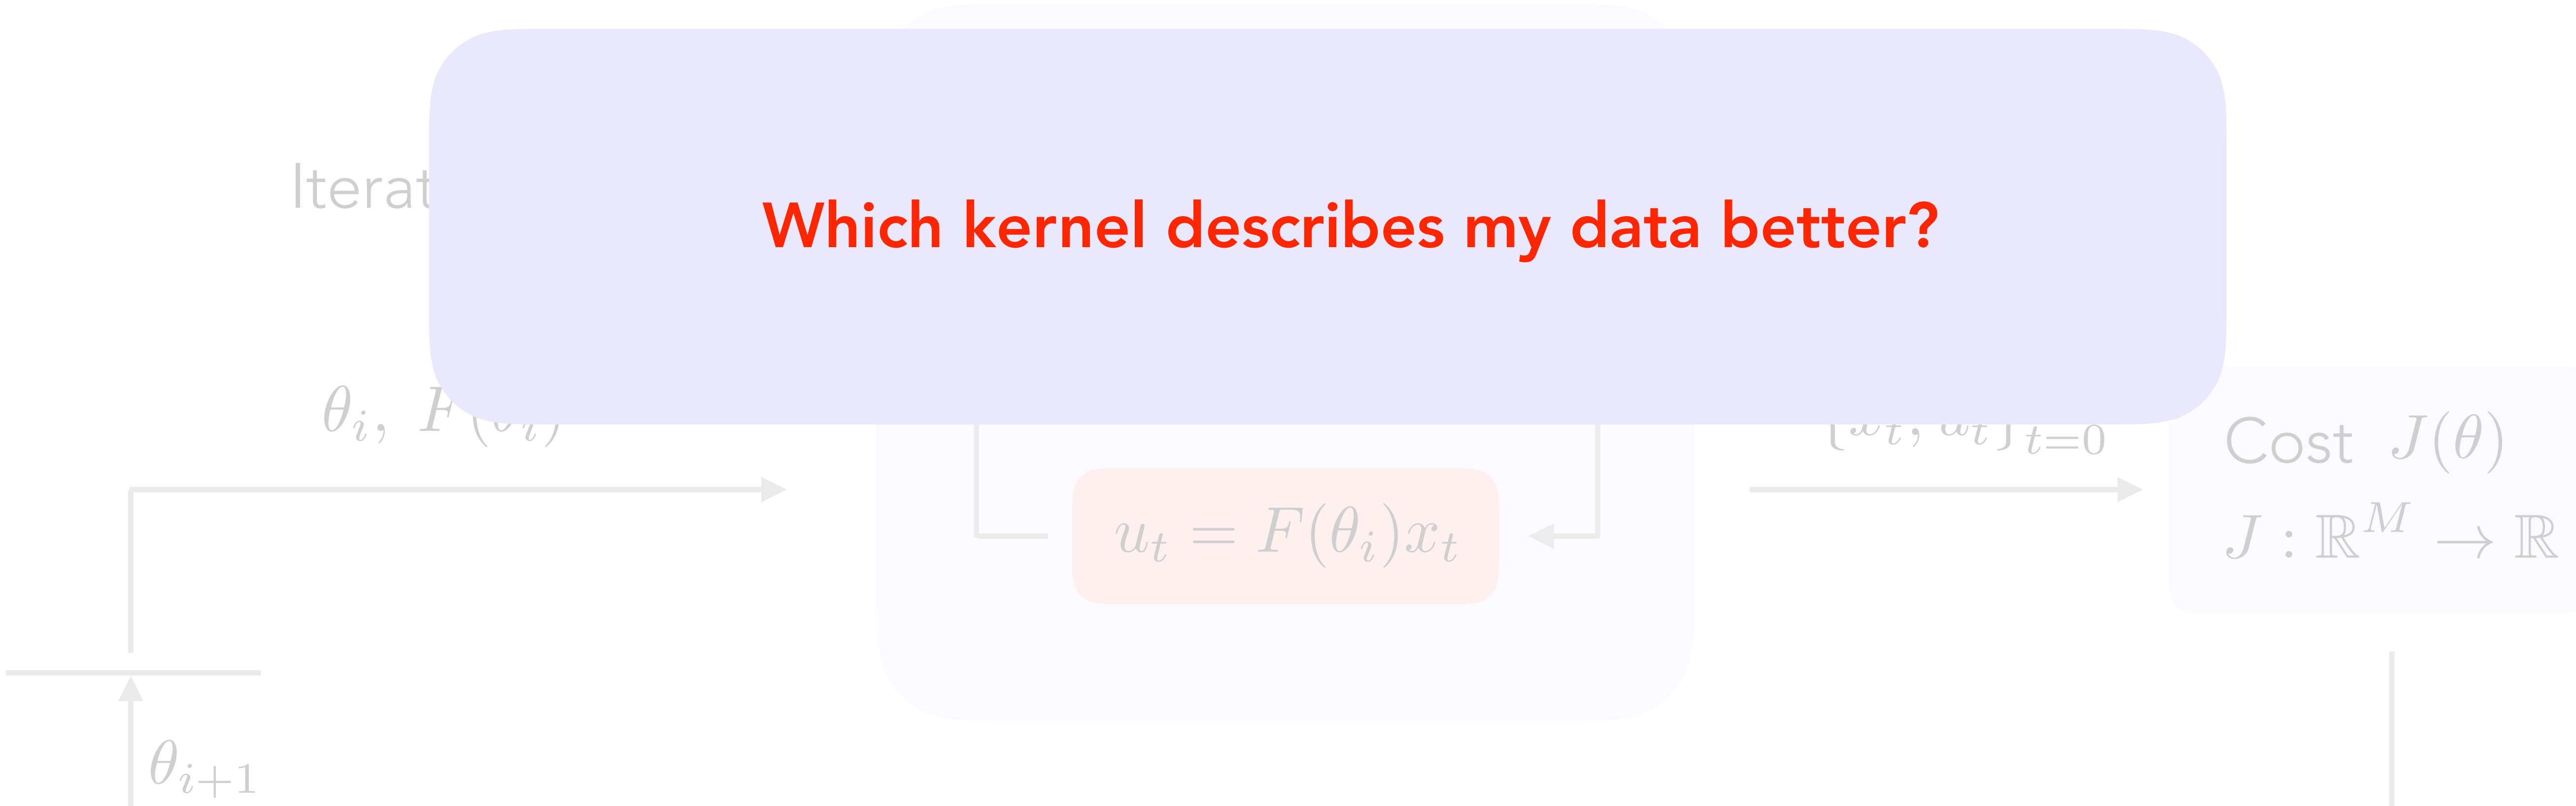

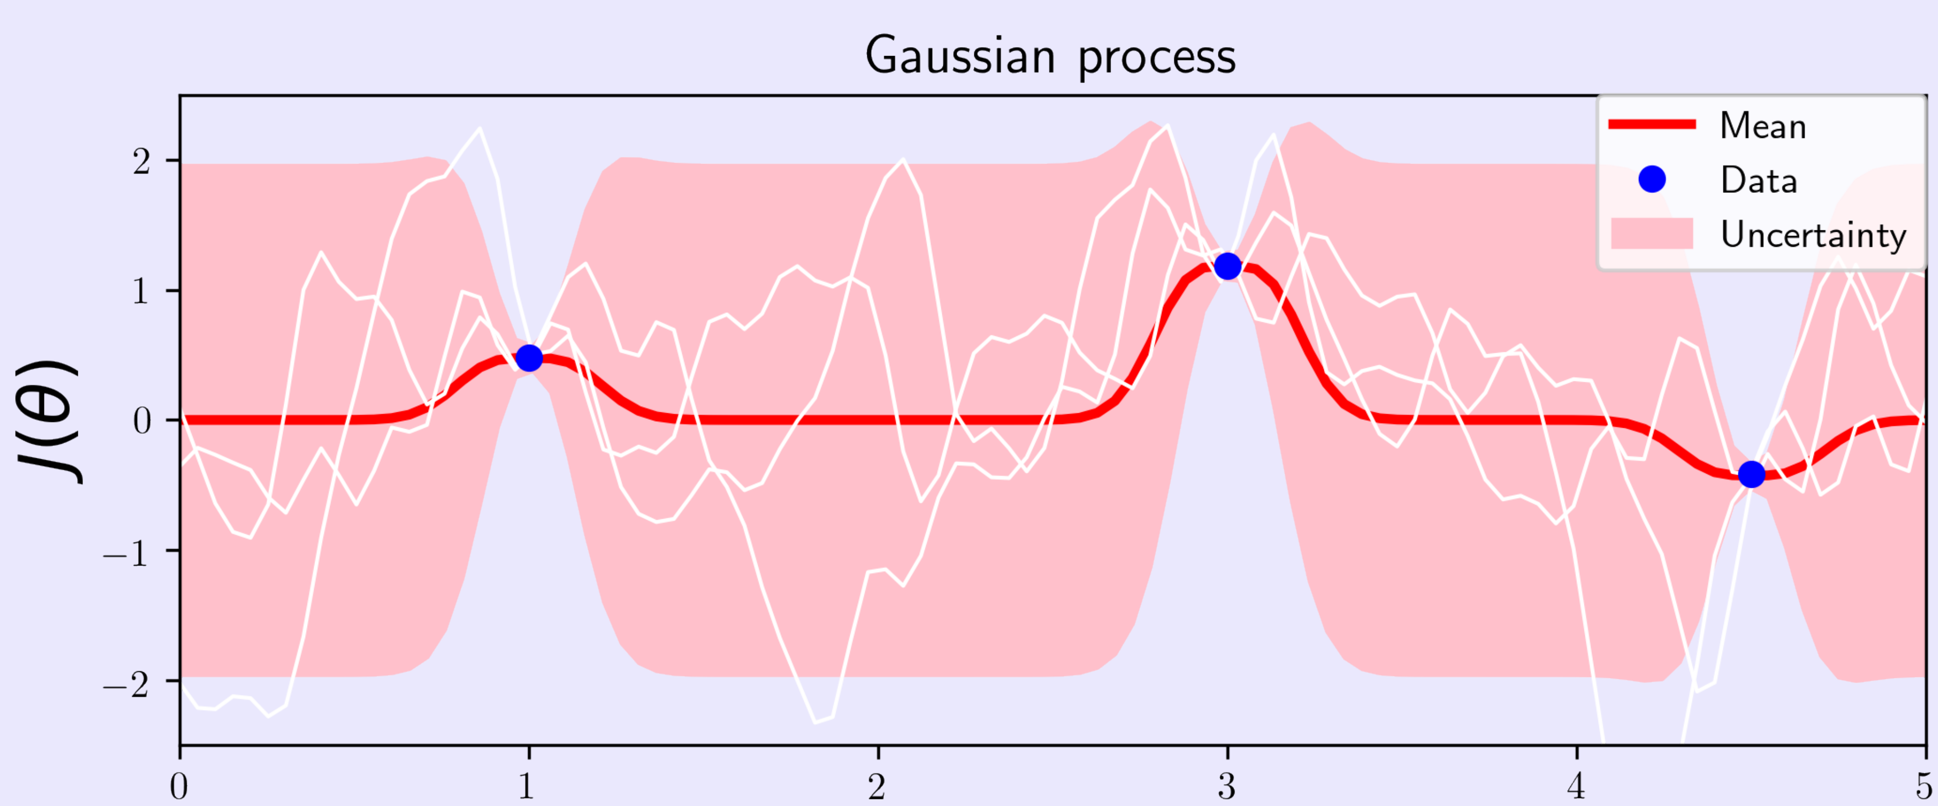

Which kernel describes my data better?

Controller → 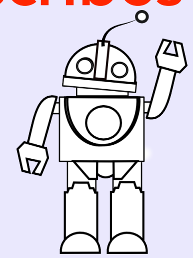 → Data

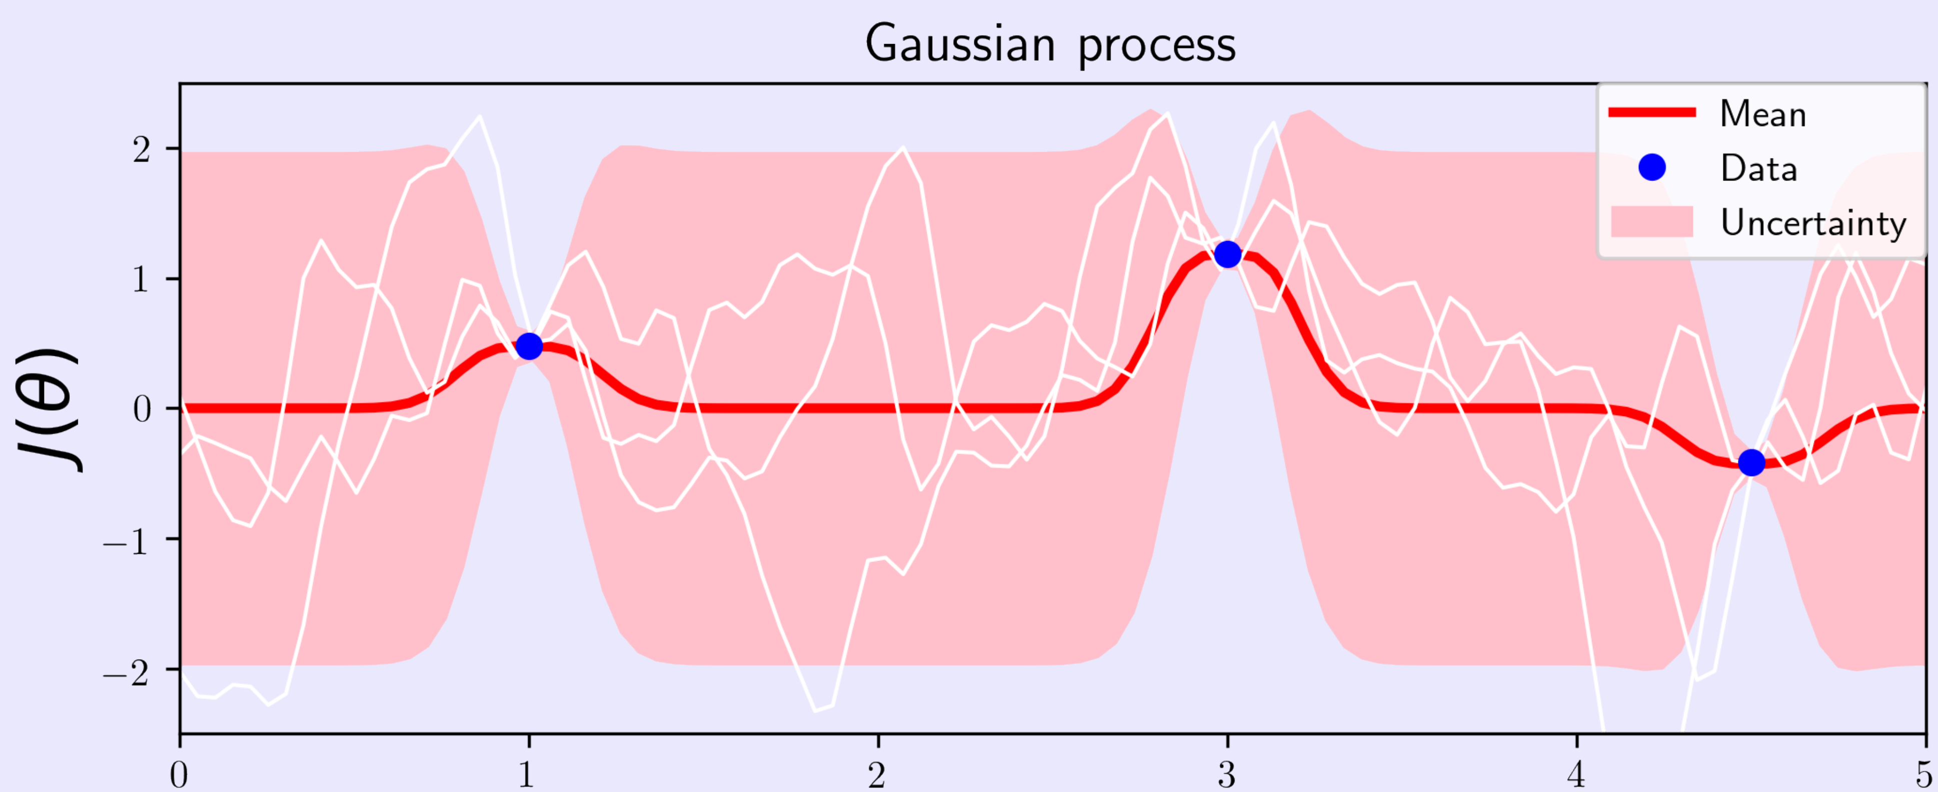

Which kernel describes my data better?

Controller → 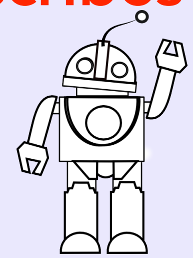 → Data

Controller structure into kernel $k(\theta, \theta')$

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

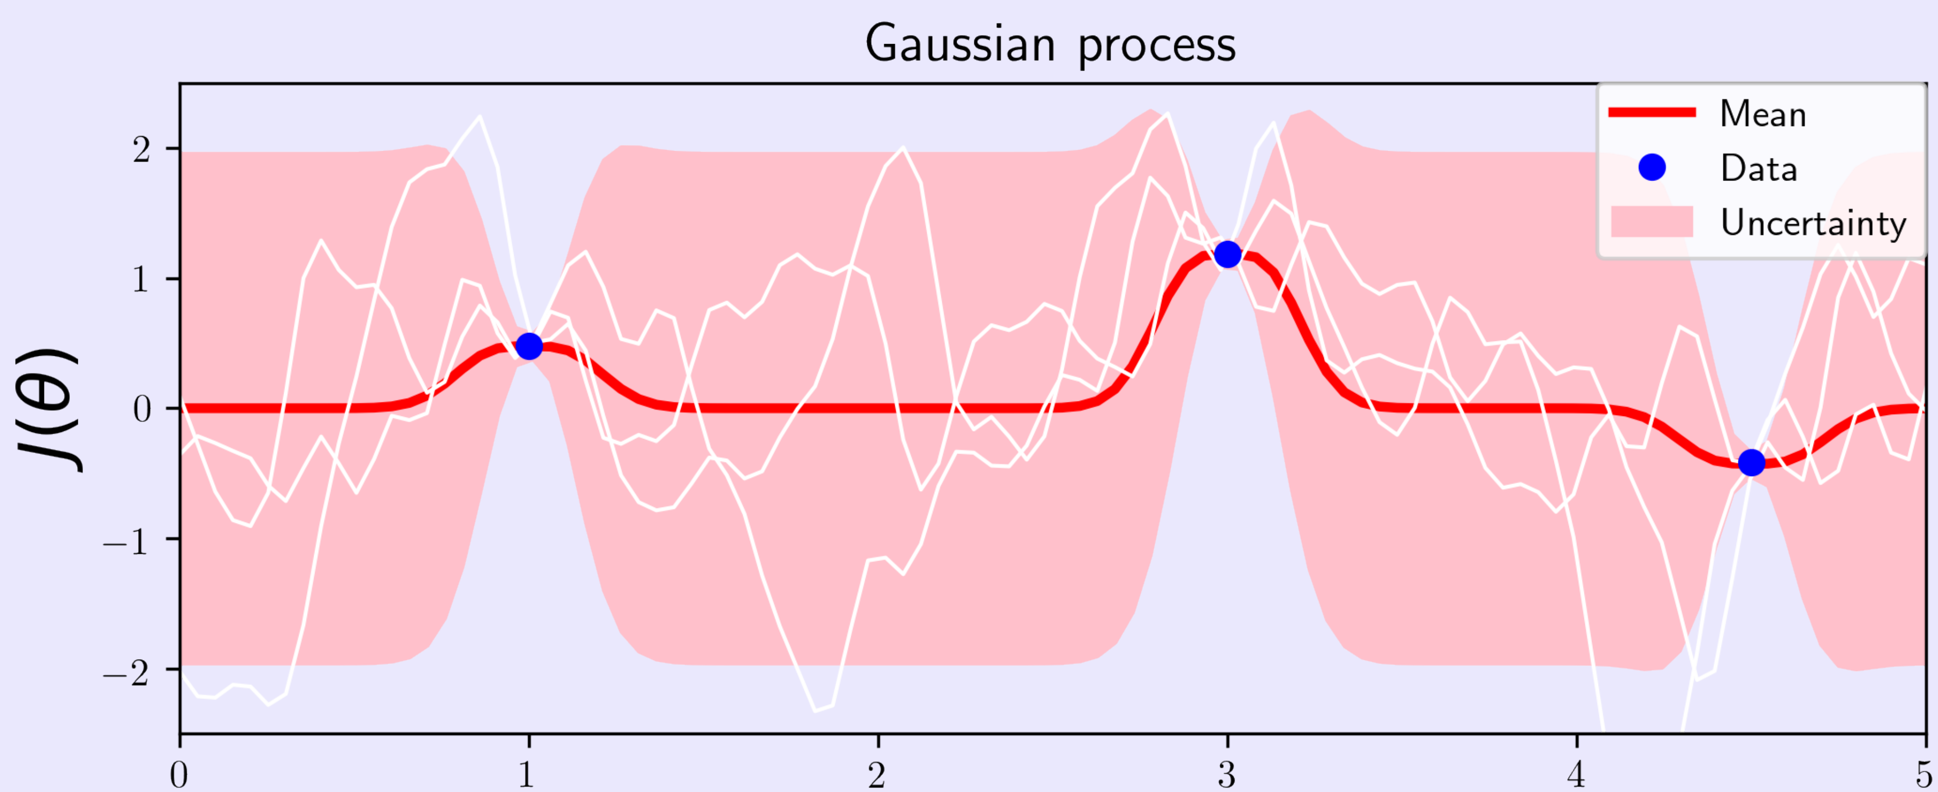

Alonso Marco | Max Planck Institute for Intelligent Systems

Which kernel describes my data better?

LQR Controller → 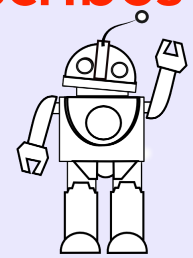 → **Data**

LQR kernel

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

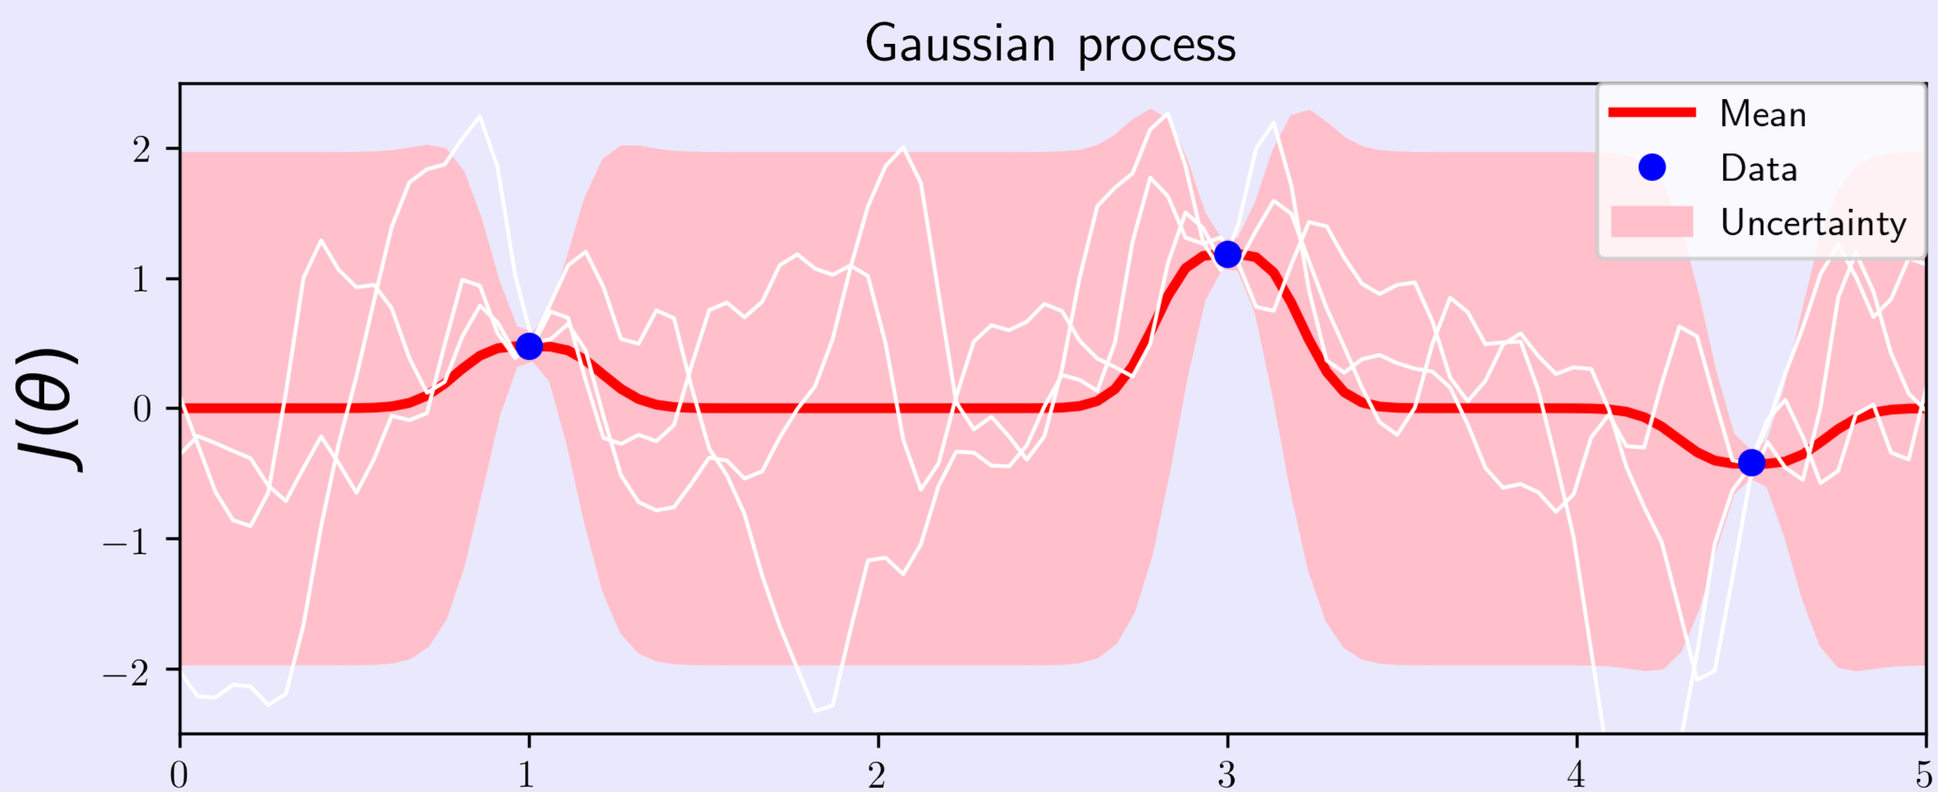

Alonso Marco | Max Planck Institute for Intelligent Systems

Goal

✴Incorporate LQR controller structure into kernel

Goal

✴Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Goal

✳Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Steps

✳Parametric LQR kernel

✳Non-parametric LQR kernel

✳Simulation results

Parametric LQR kernel

Consider a scalar LTI system $(\hat{a}, \hat{b})$

$$x_{t+1} = \hat{a} x_t + \hat{b} u_t + v_t, \; v_t \sim \mathcal{N}(0, v)$$

Model available (a, b)

$$x_{t+1} = a x_t + b u_t + v_t, \;\; v_t \sim \mathcal{N}(0, v)$$

State feedback controller

$$u_t = f x_t$$

Quadratic cost function

$$J = \lim_{T \to \infty} \frac{1}{T} \mathbb{E} \Big[\sum_{t=0}^{T-1} q x_t^2 + r u_t^2 \Big] \longrightarrow J(f) = v \frac{q + r f^2}{1 - (a + b f)^2}$$

Parametric LQR kernel

Deterministic cost function

$$J(f) = v\frac{q + rf^2}{1 - (a + bf)^2}$$

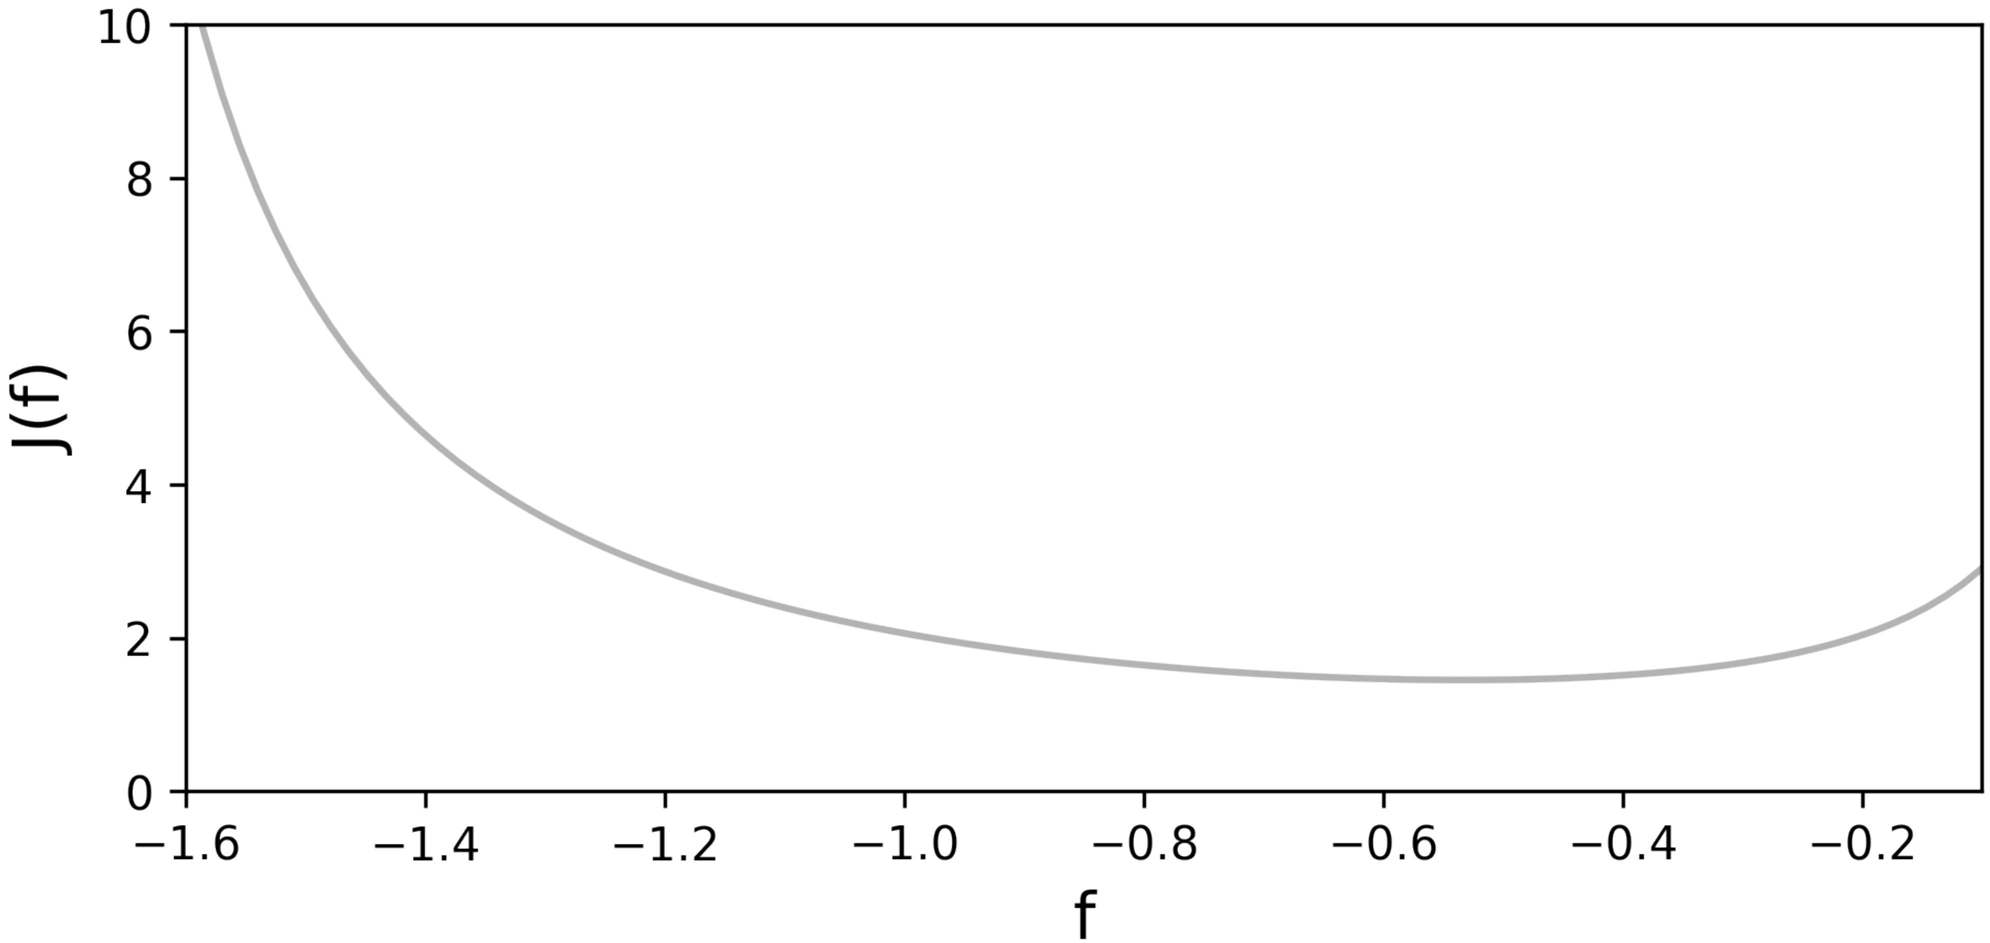

Parametric LQR kernel

Deterministic cost function

$$J(f) = v\frac{q + rf^2}{1 - (a + bf)^2} =: \underline{\phi_{(a,b)}(f)}$$

Stochastic cost function

$$J_{\mathrm{LQR}}(f) = w\,\underline{\phi_{(a,b)}(f)}, \quad w \sim \mathcal{N}(0, \sigma_w^2)$$

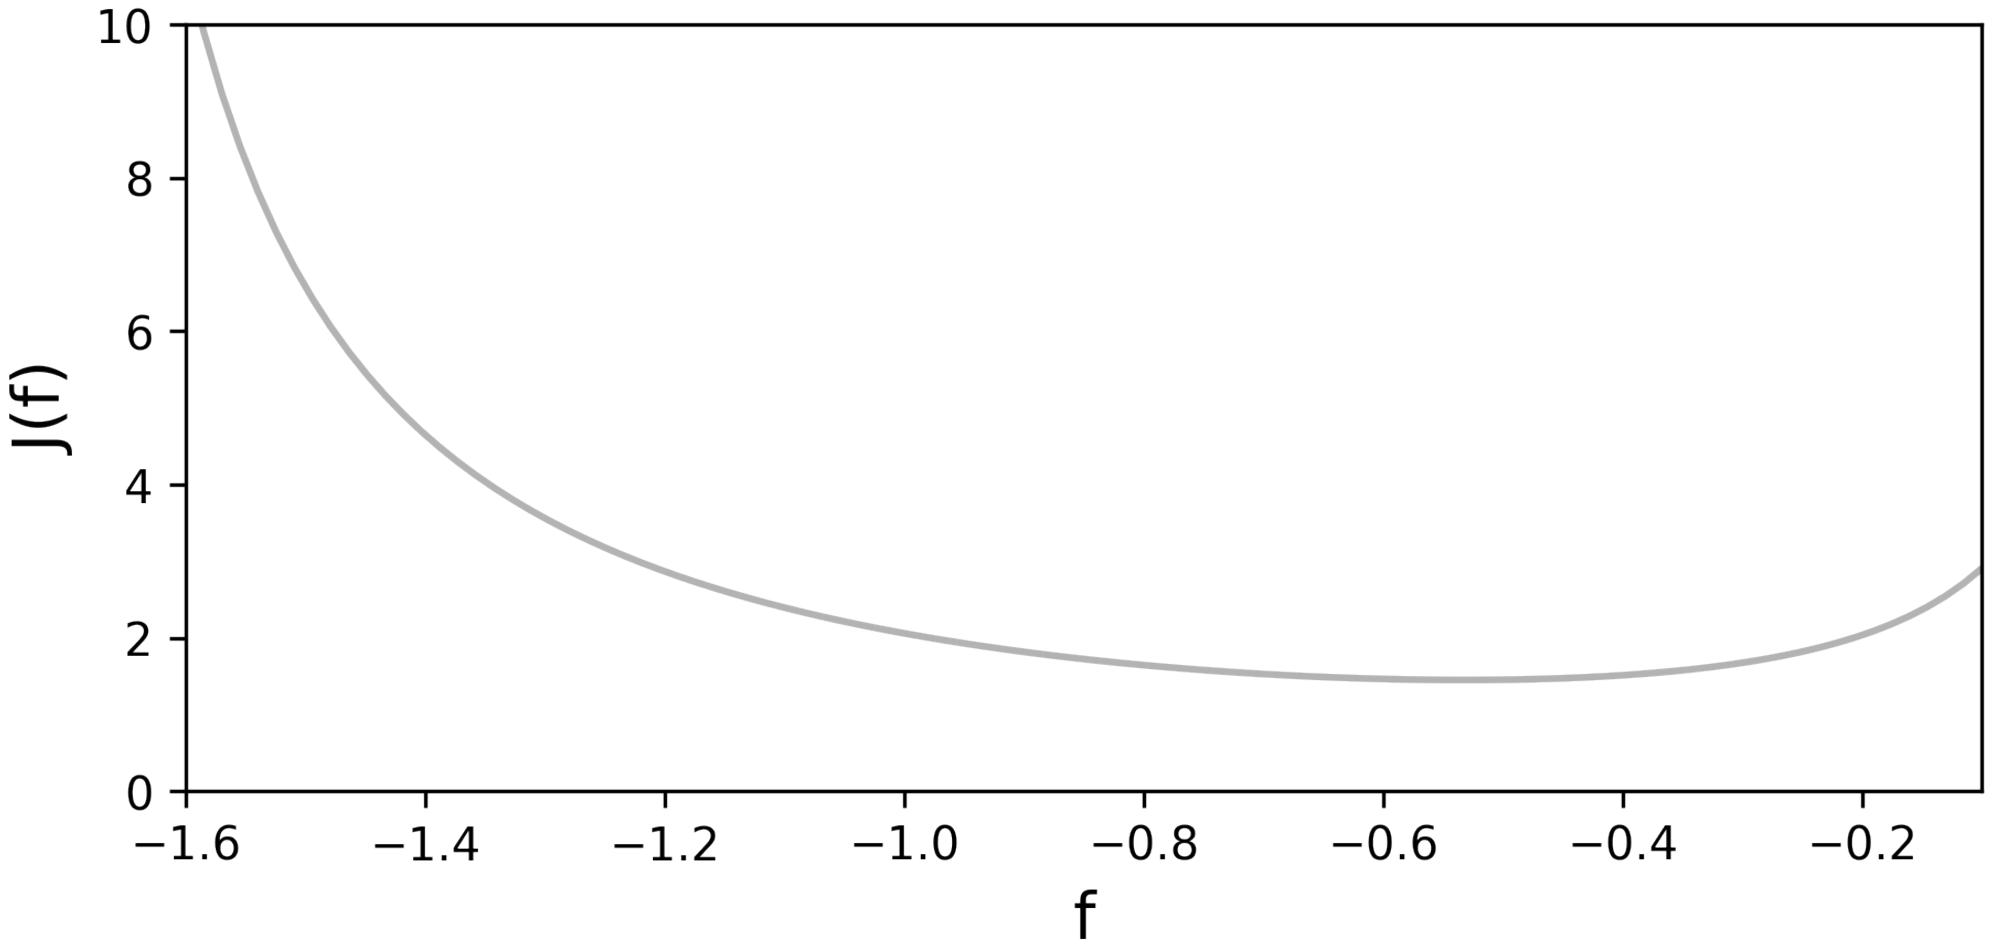

Parametric LQR kernel

Deterministic cost function

$$J(f) = v\frac{q + rf^2}{1-(a+bf)^2} =: \underline{\phi_{(a,b)}(f)}$$

Stochastic cost function

$$J_{\mathrm{LQR}}(f) = w\,\underline{\phi_{(a,b)}(f)}, \quad w \sim \mathcal{N}(0, \sigma_w^2)$$

Expected value $\quad \mathbb{E}\left[J_{\mathrm{LQR}}(f)\right] = 0$

Covariance

$$\begin{aligned}\mathbb{C}\mathrm{ov}\left[J_{\mathrm{LQR}}(f), J_{\mathrm{LQR}}(f')\right] &= \mathbb{E}\left[J_{\mathrm{LQR}}(f)J_{\mathrm{LQR}}(f')\right] \\ &= \mathbb{E}\left[w^2\right]\phi_{(\bar{a},\bar{b})}(f)\phi_{(\bar{a},\bar{b})}(f') \\ &= \sigma_w^2\phi_{(\bar{a},\bar{b})}(f)\phi_{(\bar{a},\bar{b})}(f') \\ &= k_{\mathrm{LQR}}(f,f')\end{aligned}$$

Parametric LQR kernel

$$k_{\mathrm{pLQR}}(f,f') = \sigma_w^2\frac{v^2(q+rf^2)(q+rf'^2)}{(1-(\bar{a}+\bar{b}f)^2)(1-(\bar{a}+\bar{b}f')^2)}$$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

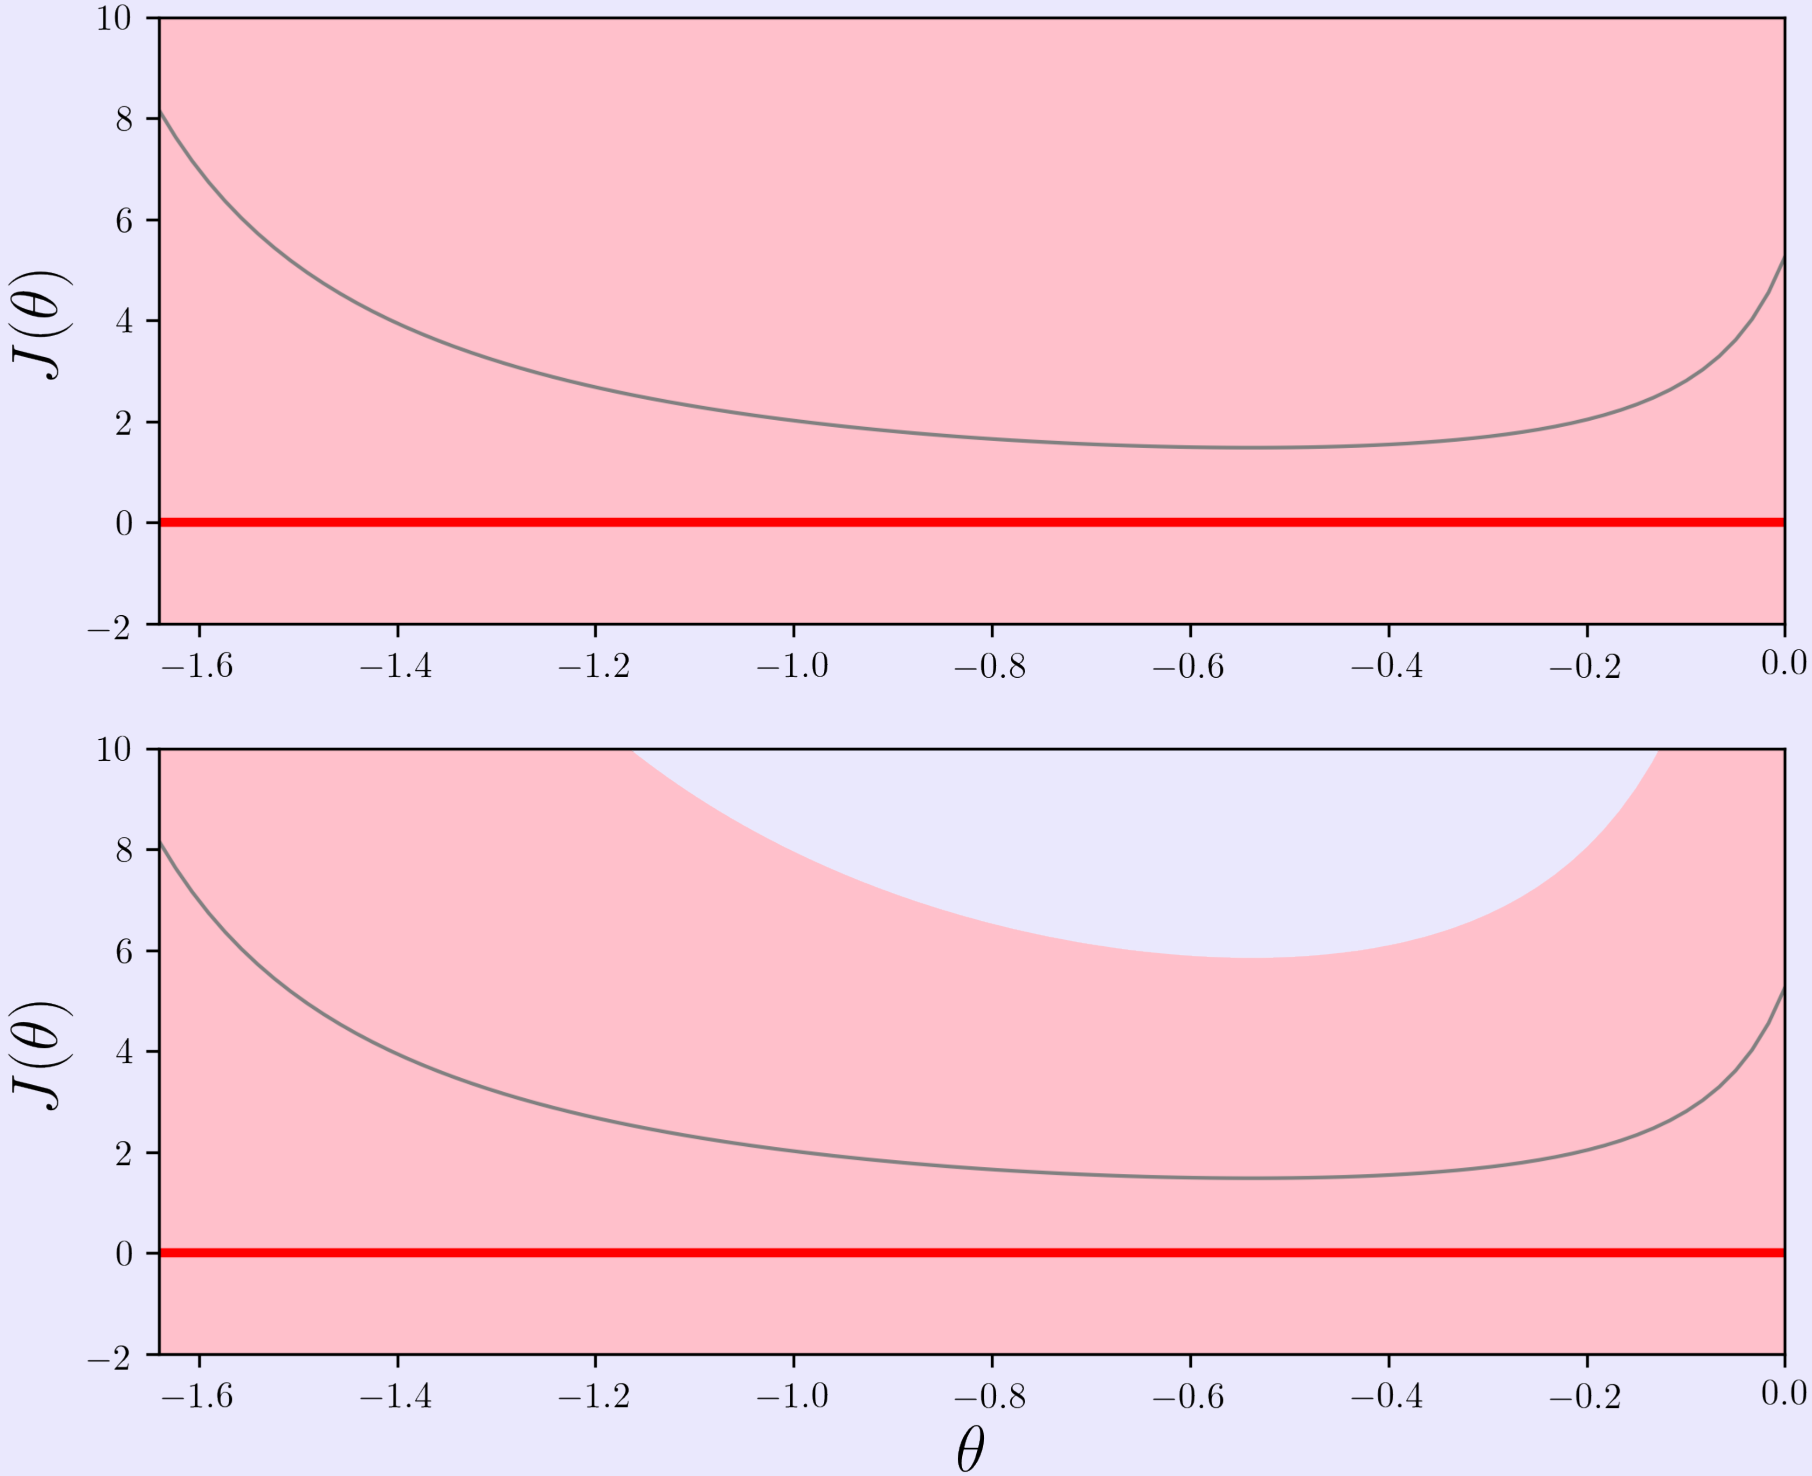

Standard kernel

Parametric LQR kernel
$a = 0.9,\ b = 1$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

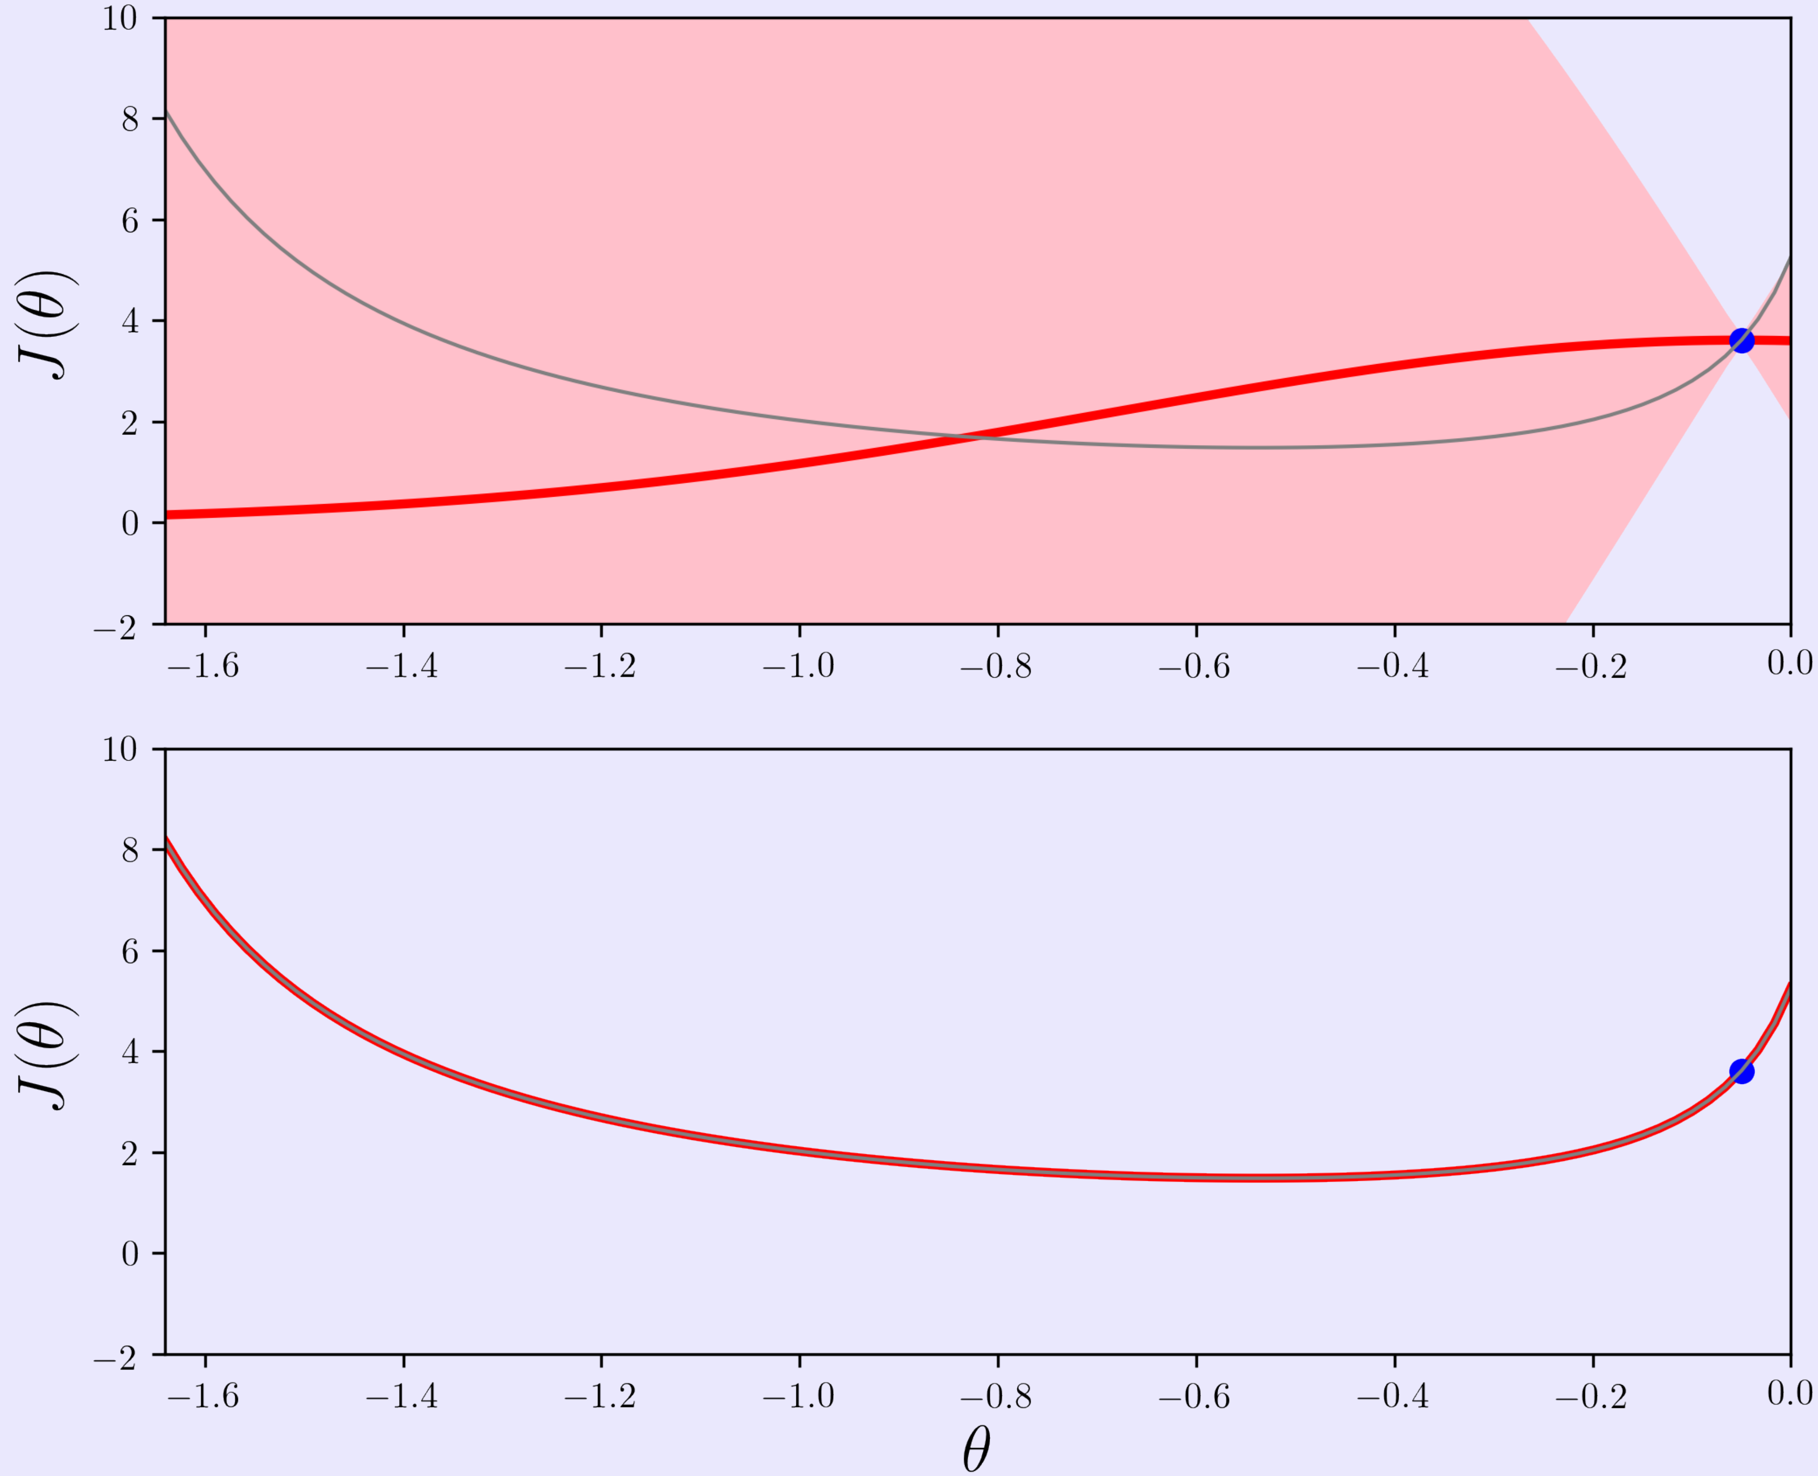

Standard kernel

Parametric LQR kernel

$a = 0.9,\ b = 1$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

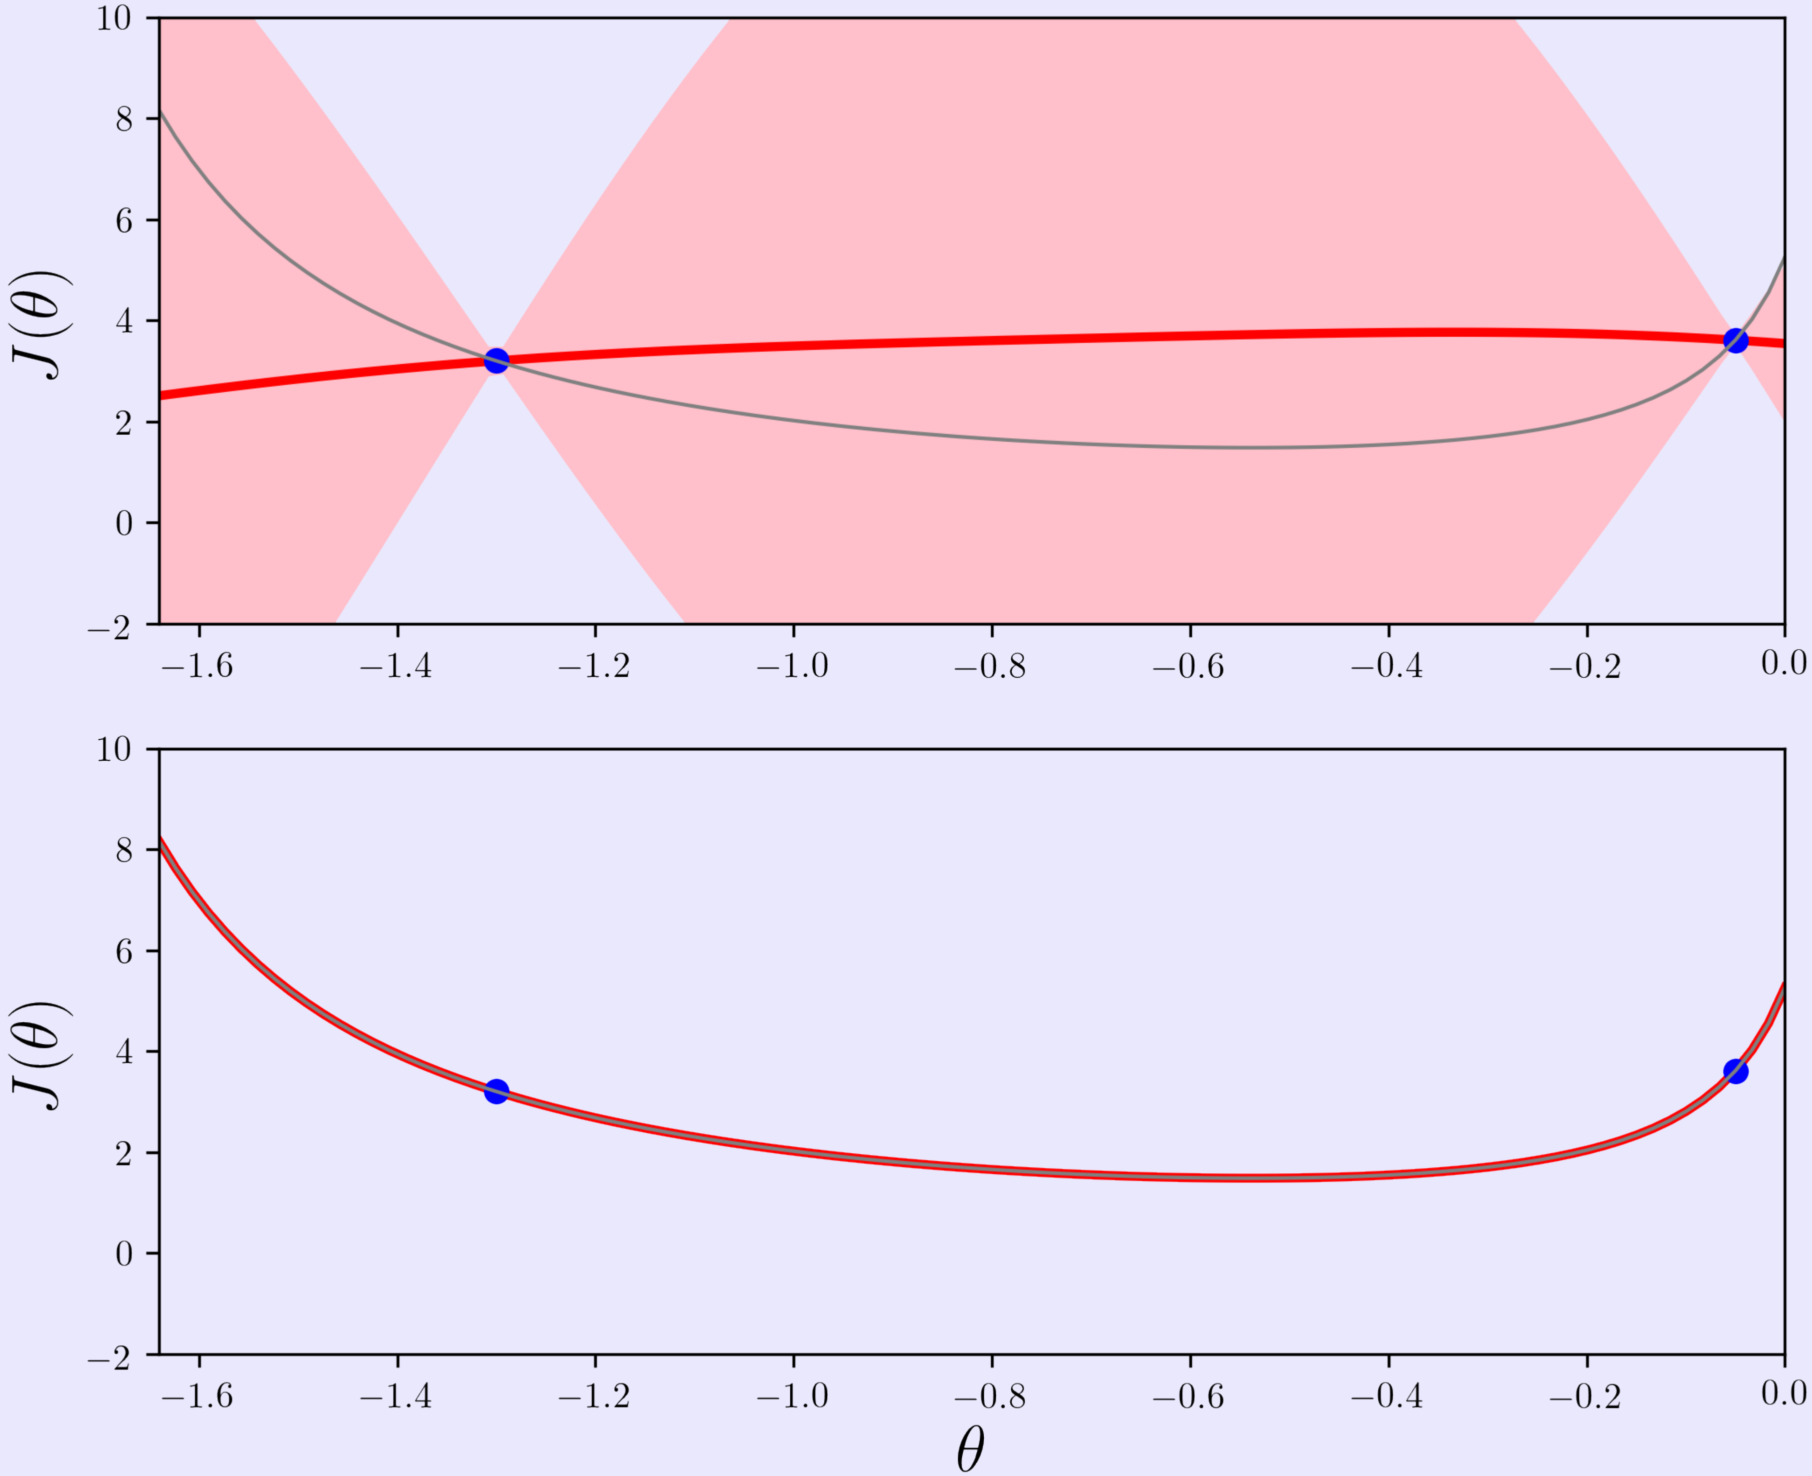

Standard kernel

Parametric LQR kernel

$a = 0.9,\ b = 1$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

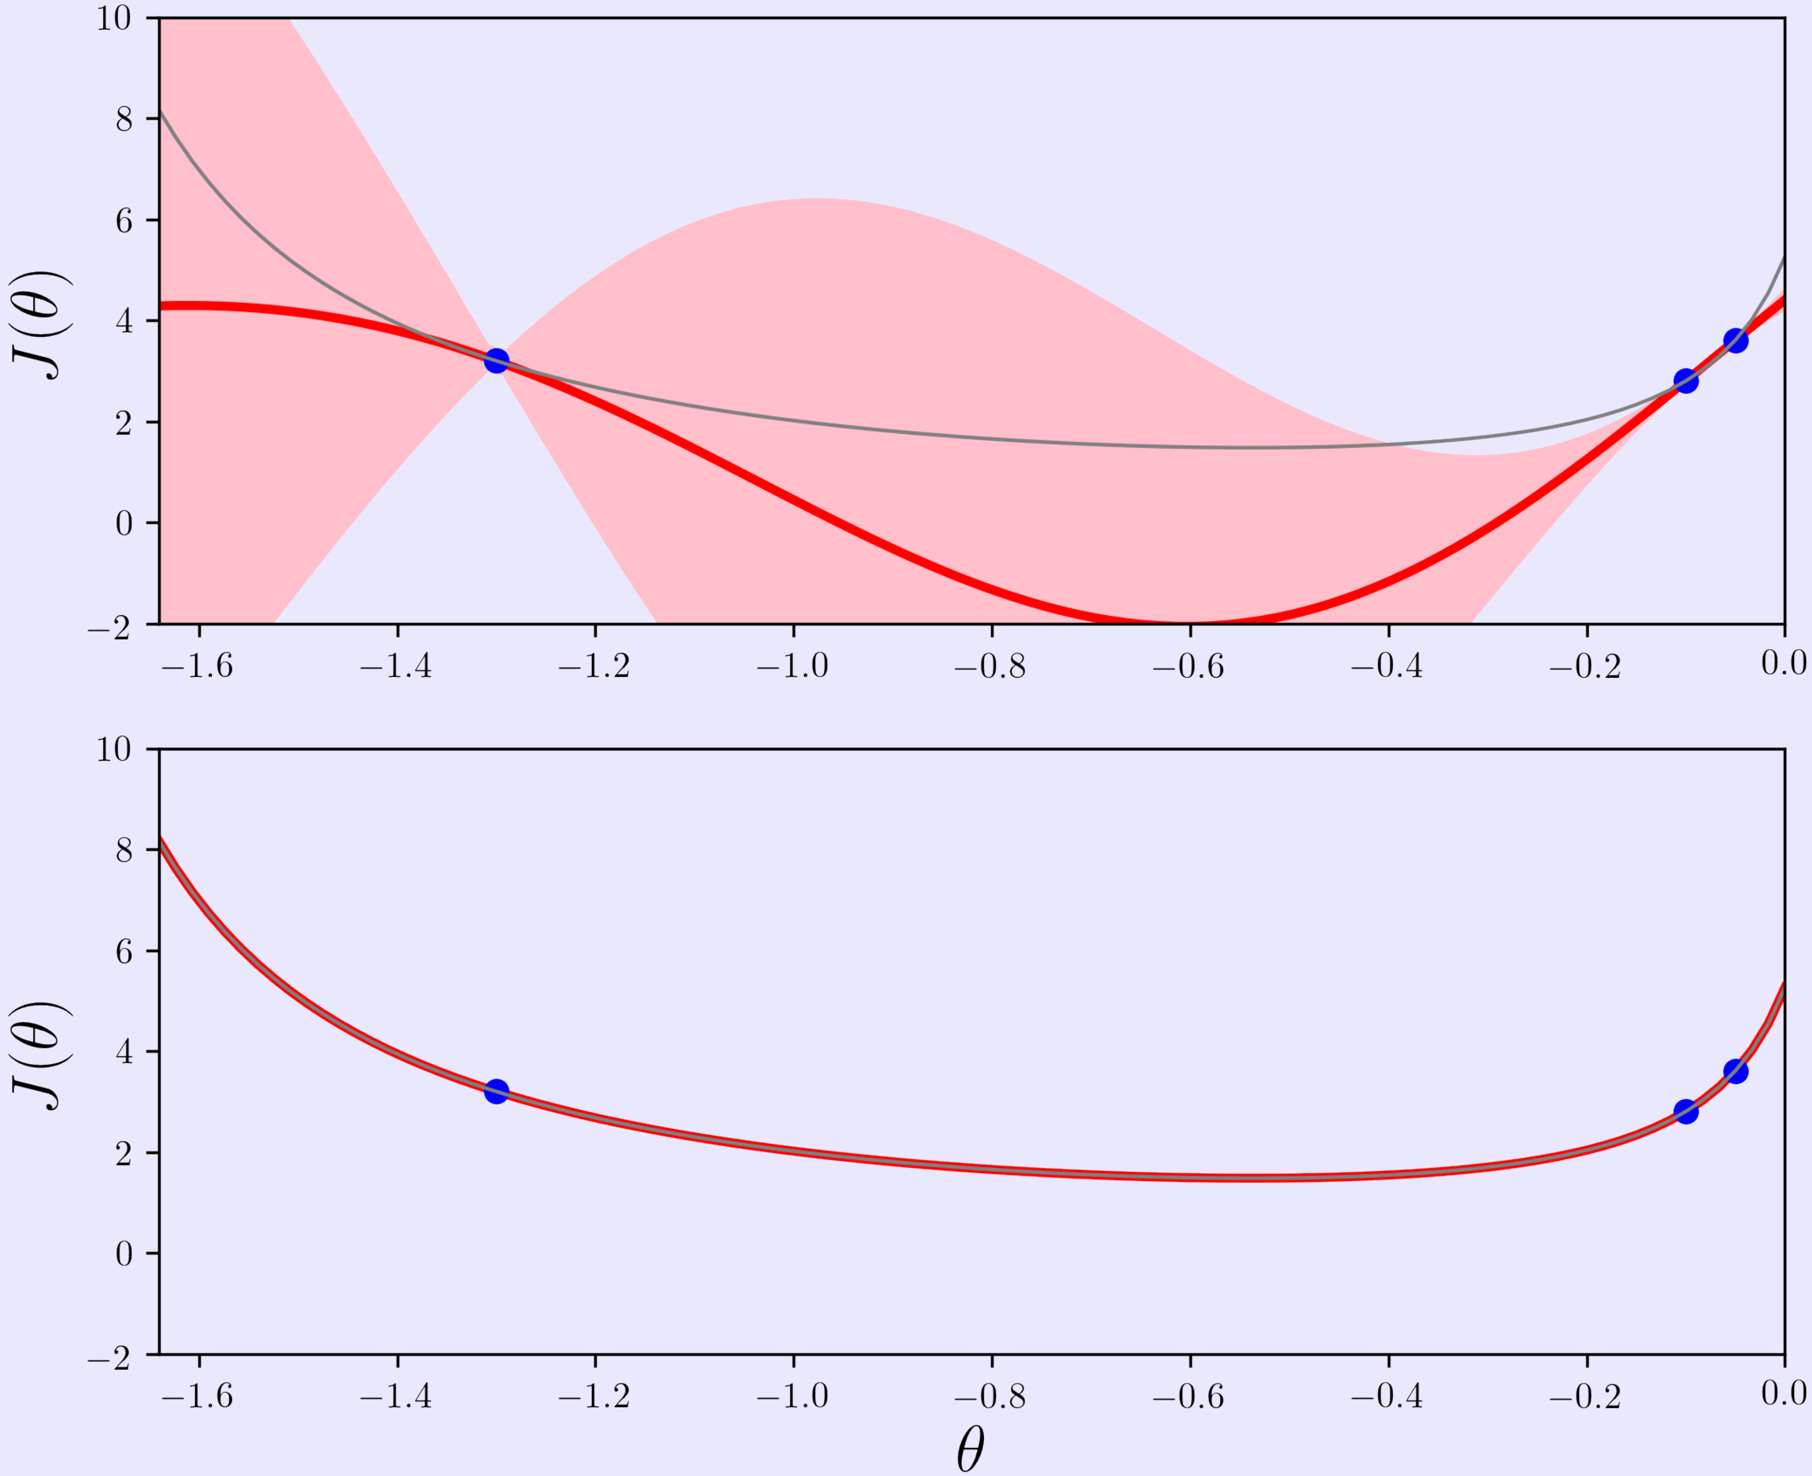

Standard kernel

Parametric LQR kernel

$a = 0.9,\ b = 1$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

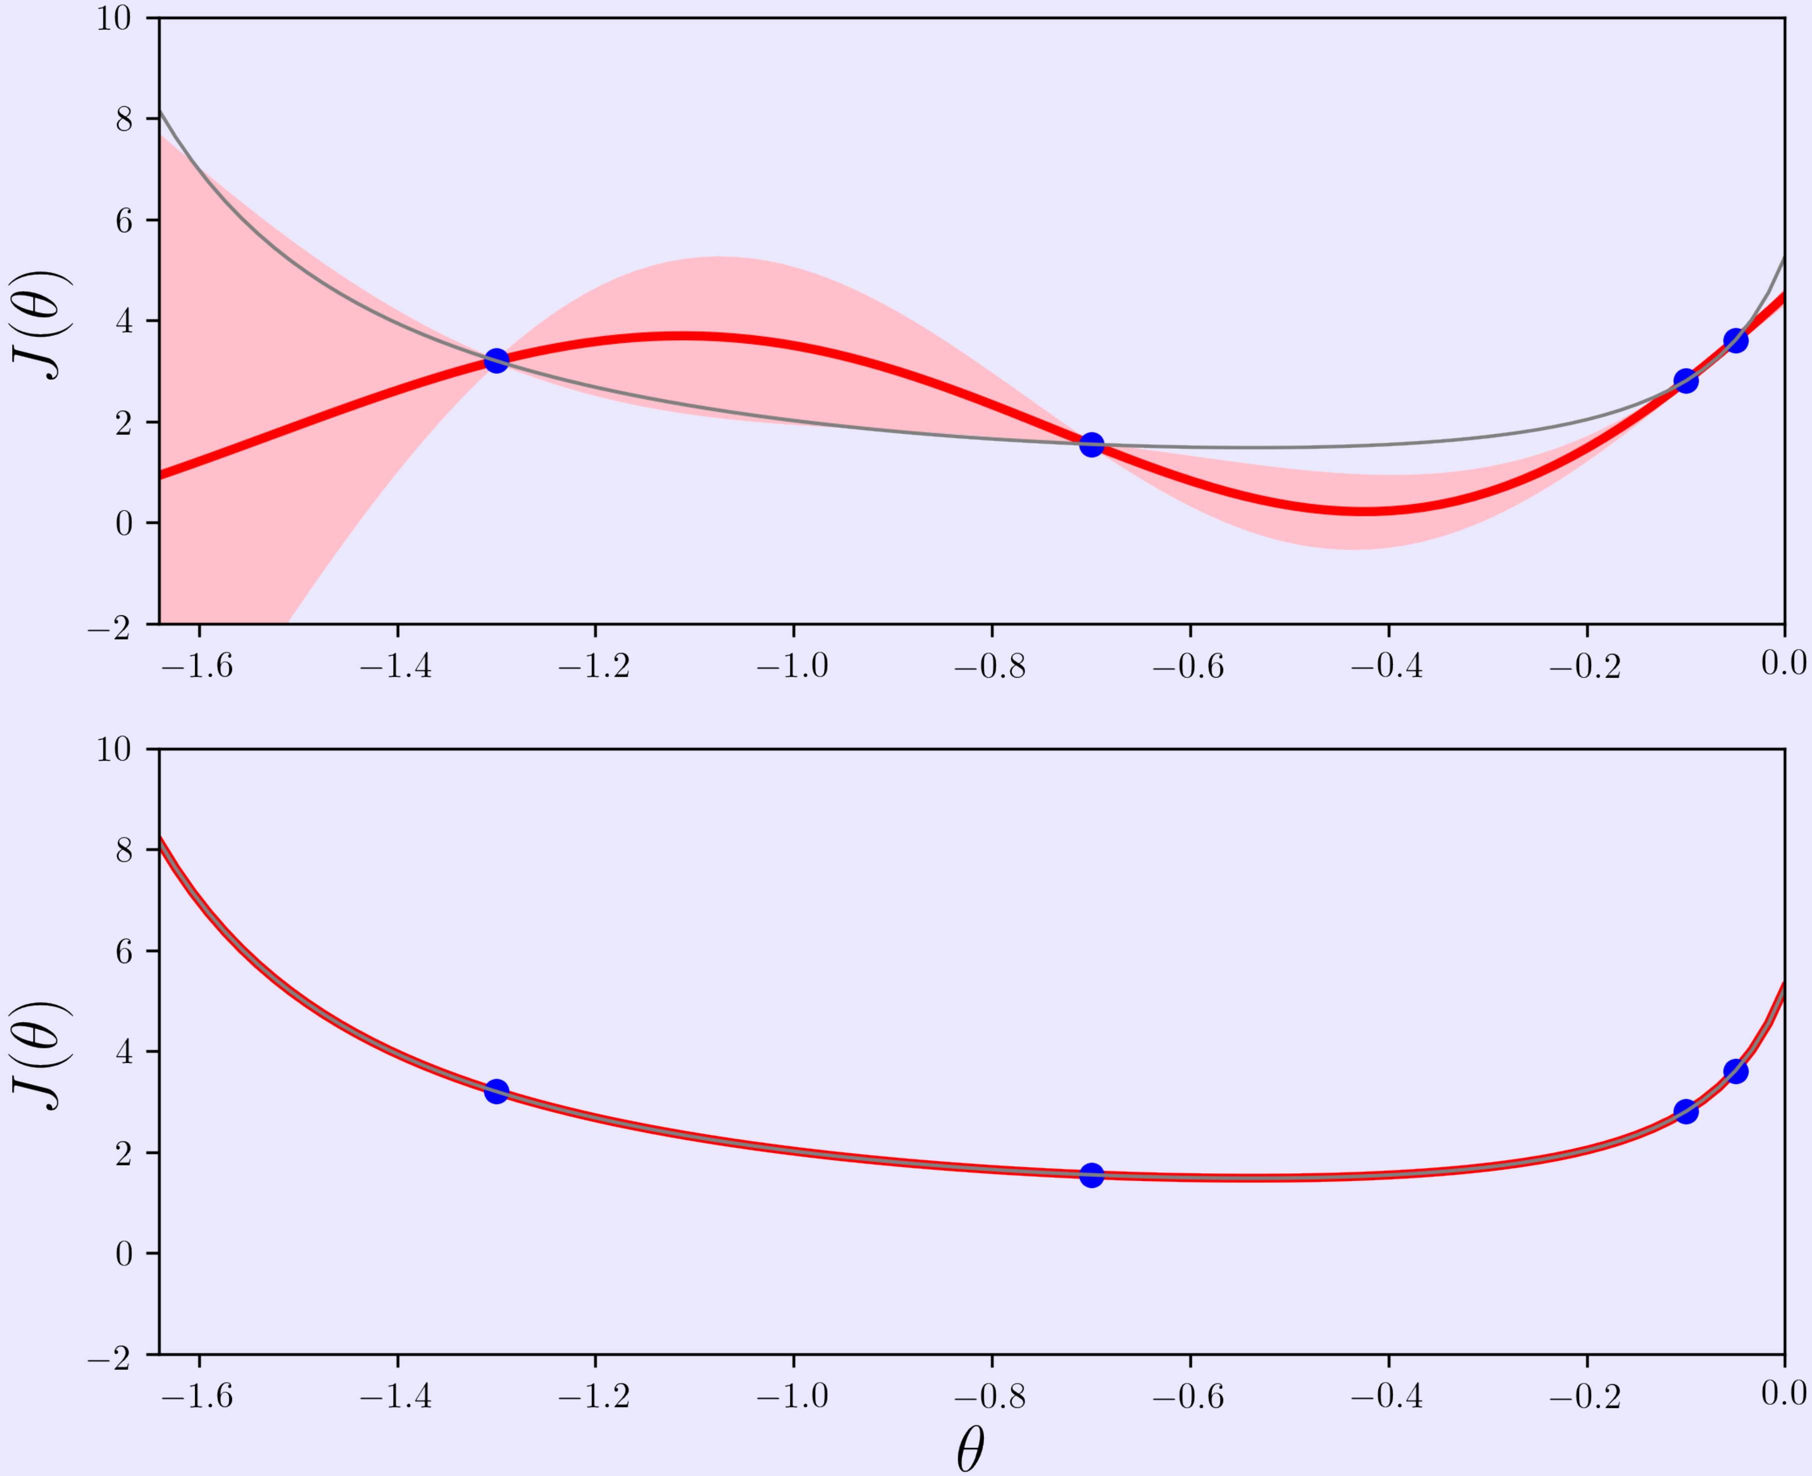

Standard kernel

Parametric LQR kernel

$a = 0.9,\ b = 1$

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

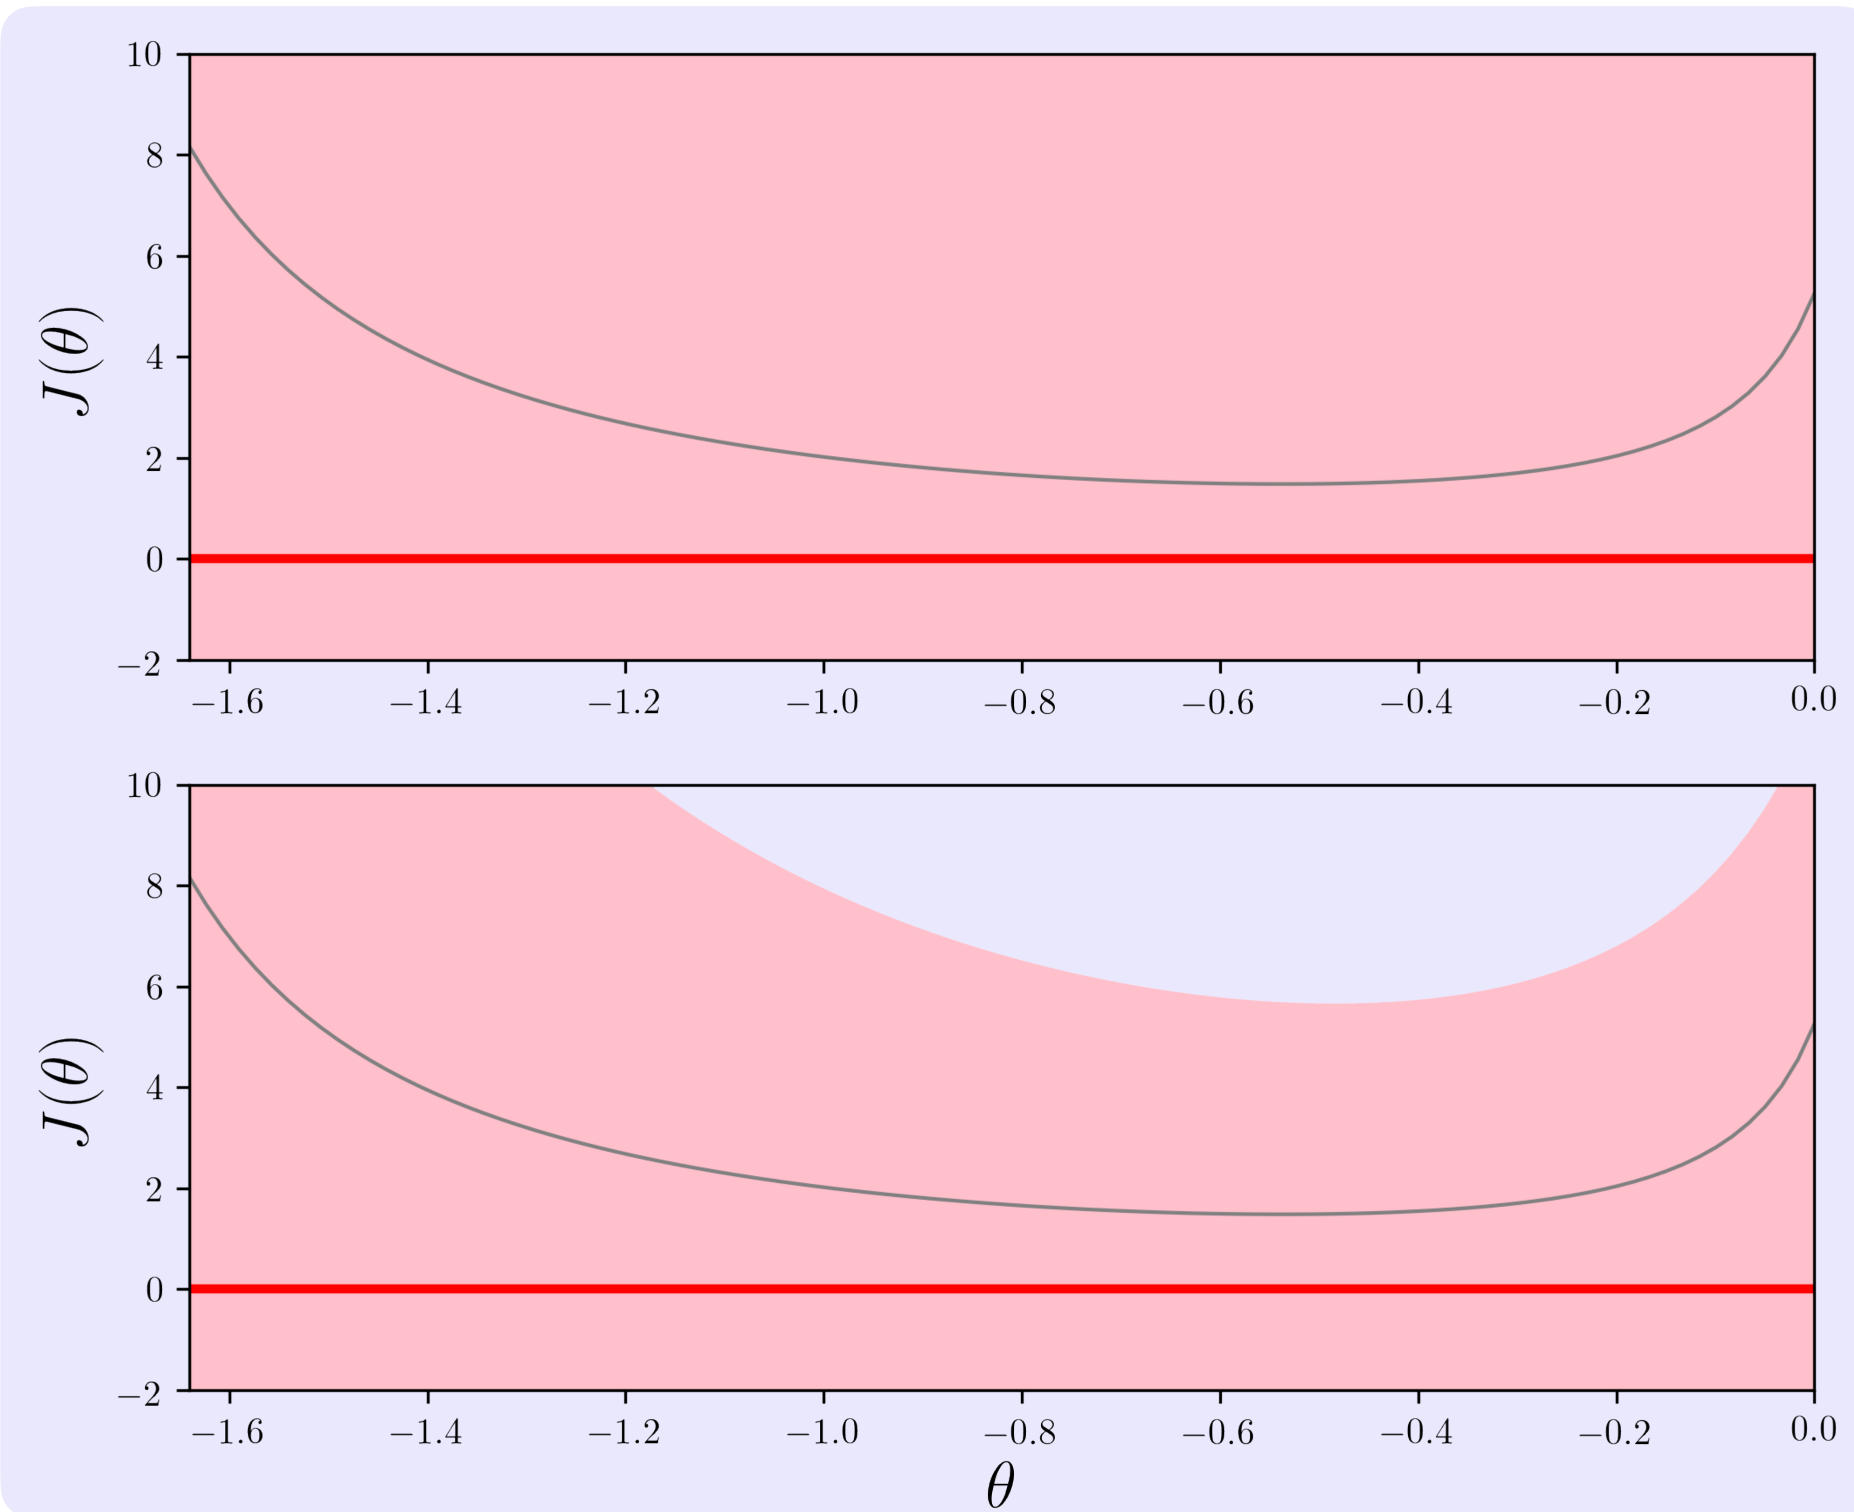

Standard kernel

Parametric LQR kernel

$a = 0.8,\ b = 0.9$

~10% OFF!

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

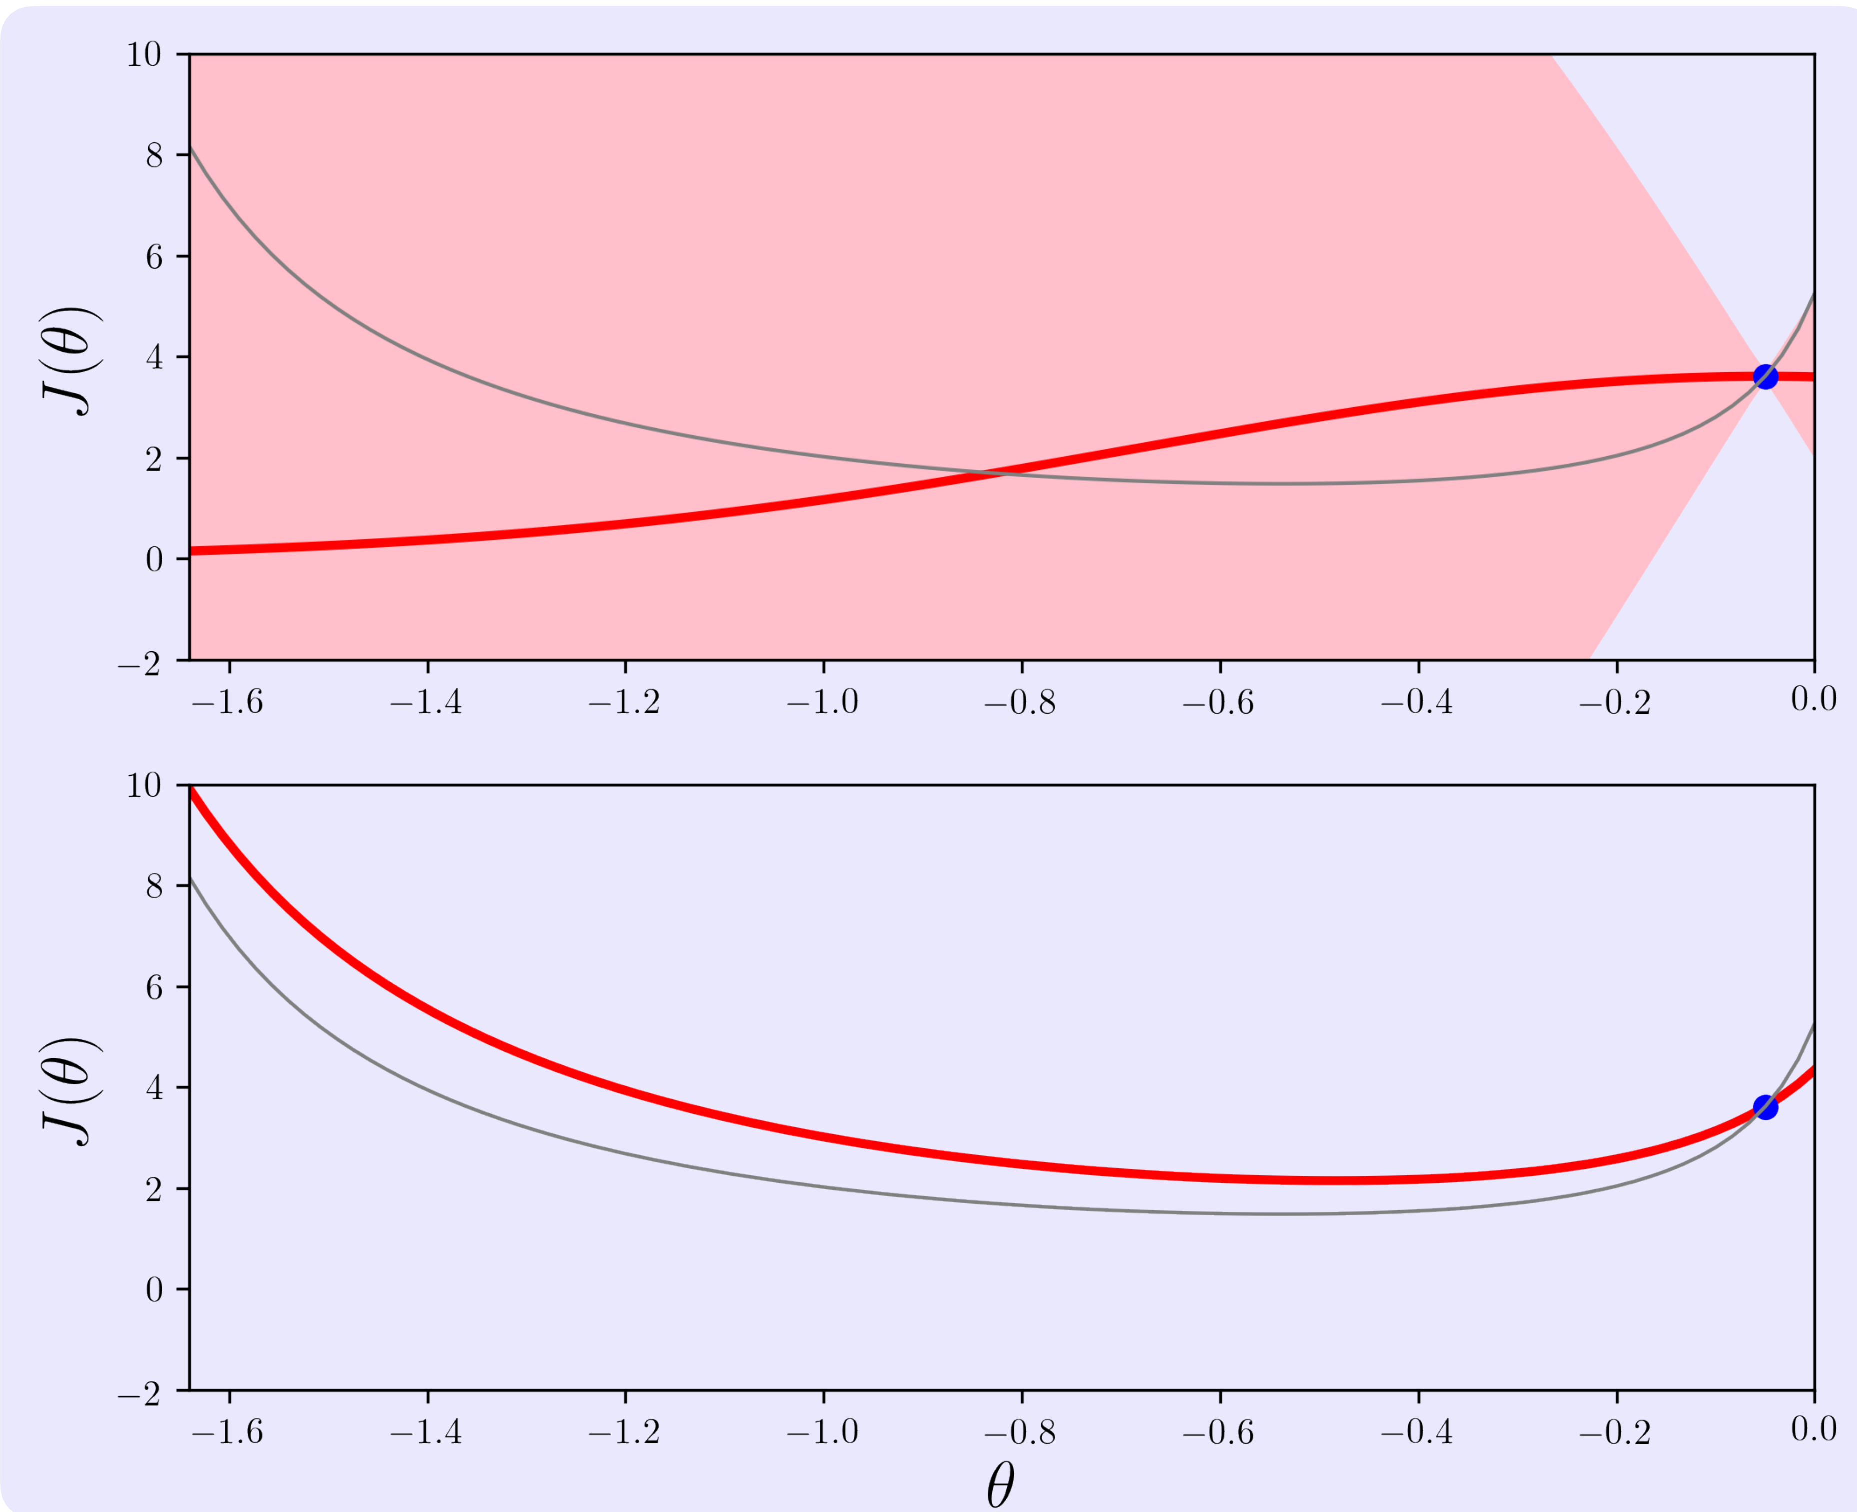

Standard kernel

Parametric LQR kernel

$a = 0.8,\ b = 0.9$

~10% OFF!

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

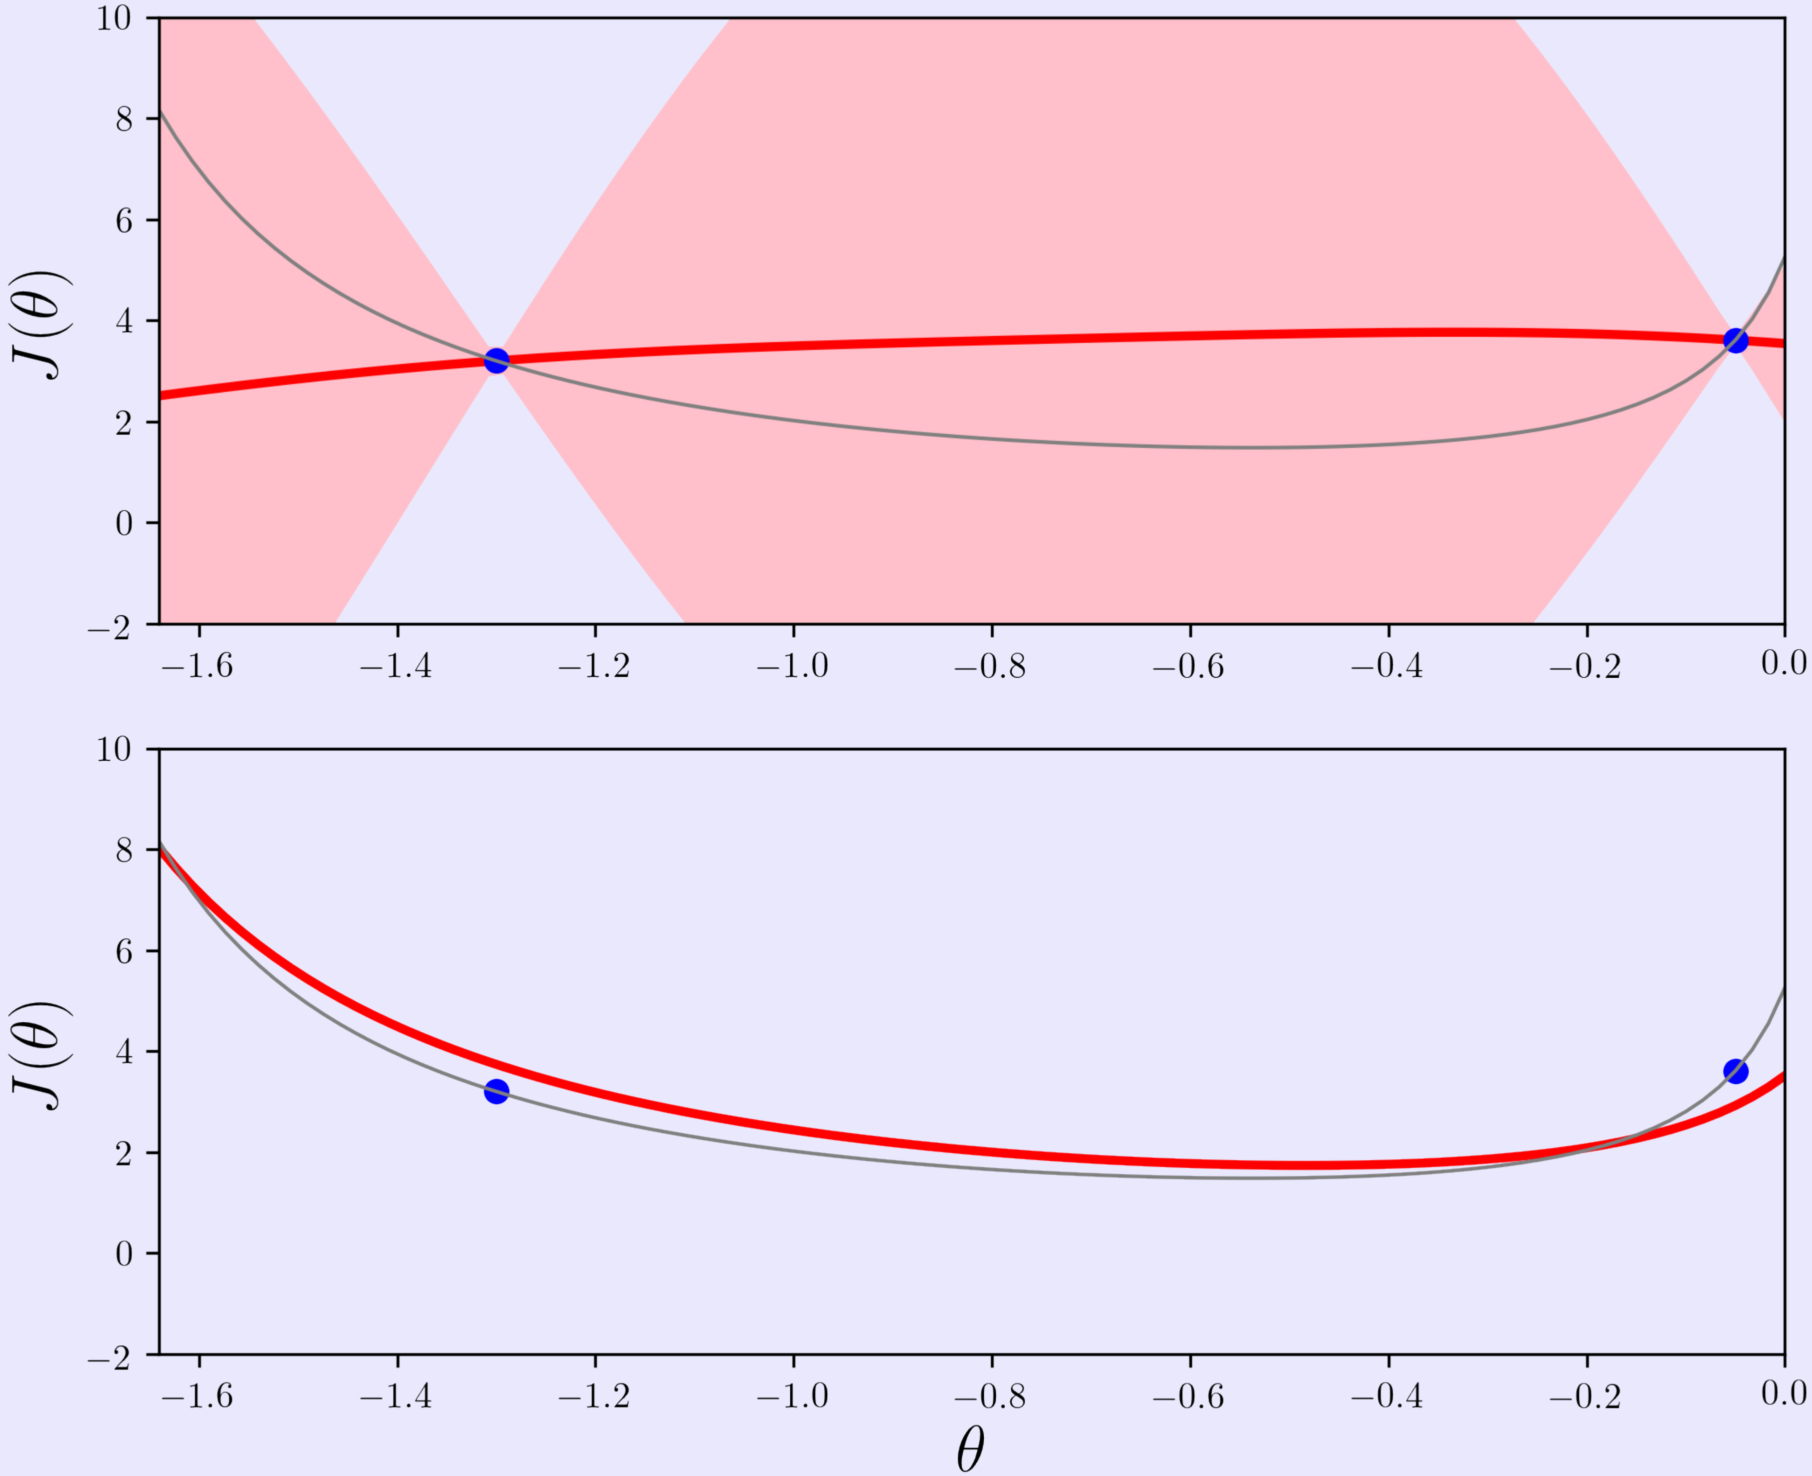

Standard kernel

Parametric LQR kernel

$a = 0.8,\ b = 0.9$

~10% OFF!

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

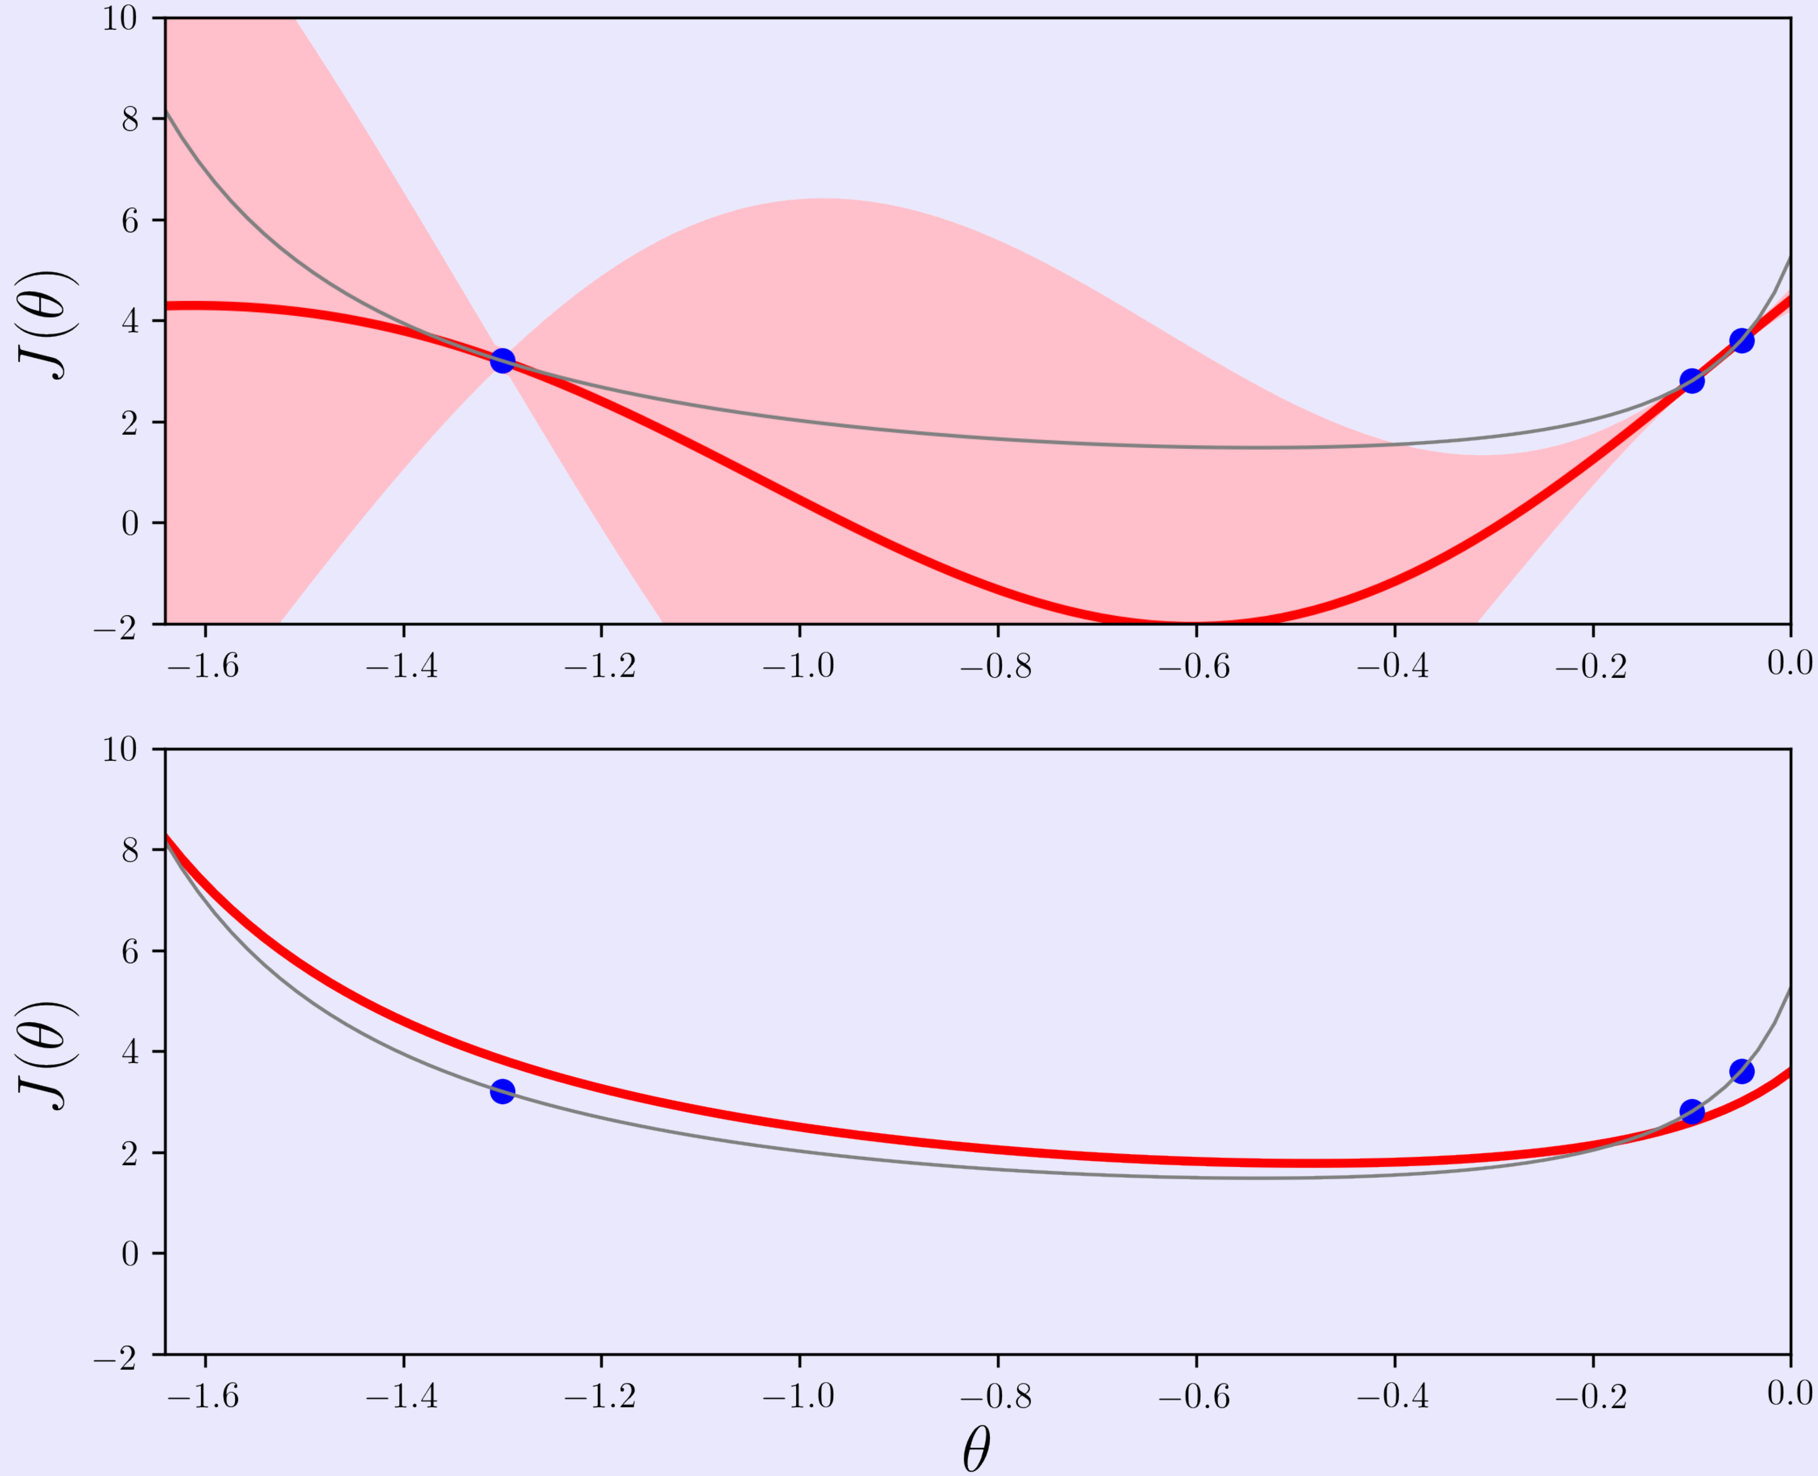

Standard kernel

Parametric LQR kernel

$a = 0.8,\ b = 0.9$

~10% OFF!

Parametric LQR kernel

True system $x_{t+1} = 0.9x_t + u_t + v_t,\ v_t \sim \mathcal{N}(0,1)$

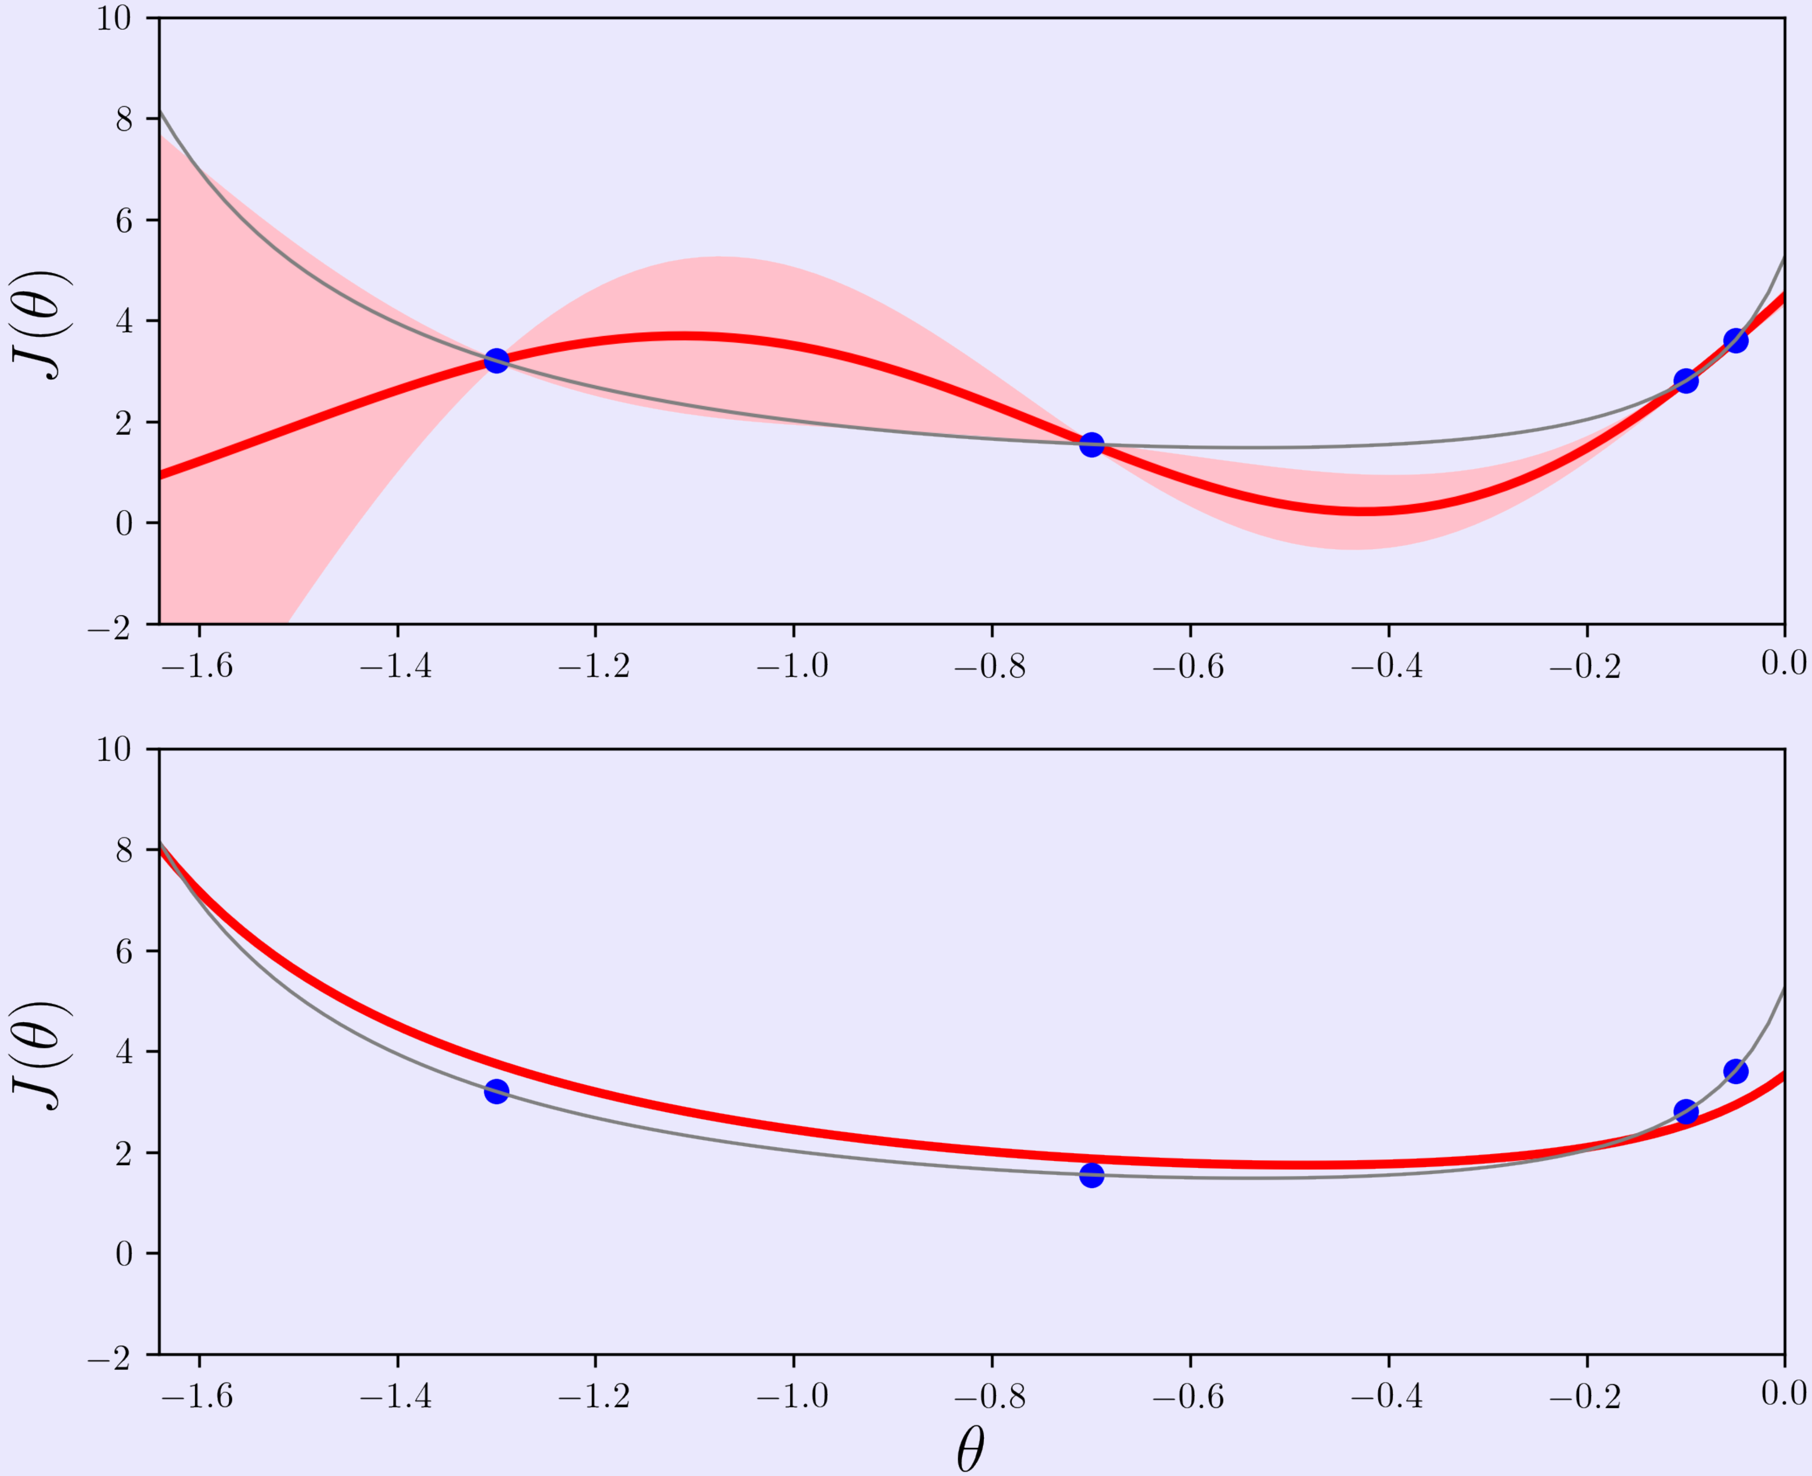

Standard kernel

Parametric LQR kernel

$a = 0.8,\ b = 0.9$

~10% OFF!

Goal

✴Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Steps

✓Parametric LQR kernel

✴**Non-parametric LQR kernel**

✴Simulated results

Non-parametric LQR kernel

Consider a scalar LTI system $(\hat{a}, \hat{b})$

$$x_{t+1} = \hat{a} x_t + \hat{b} u_t + v_t,\ v_t \sim \mathcal{N}(0, v)$$

Uncertain model (a, b)

$$x_{t+1} = a x_t + b u_t + v_t, \quad v_t \sim \mathcal{N}(0, v)$$
$$a \in [a_{\min}, a_{\max}], b \in [b_{\min}, b_{\max}]$$

Stochastic cost: one feature

$$J_{\mathrm{LQR}}(f) = w\, \phi_{(a,b)}(f), \quad w \sim \mathcal{N}(0, \sigma_w^2)$$

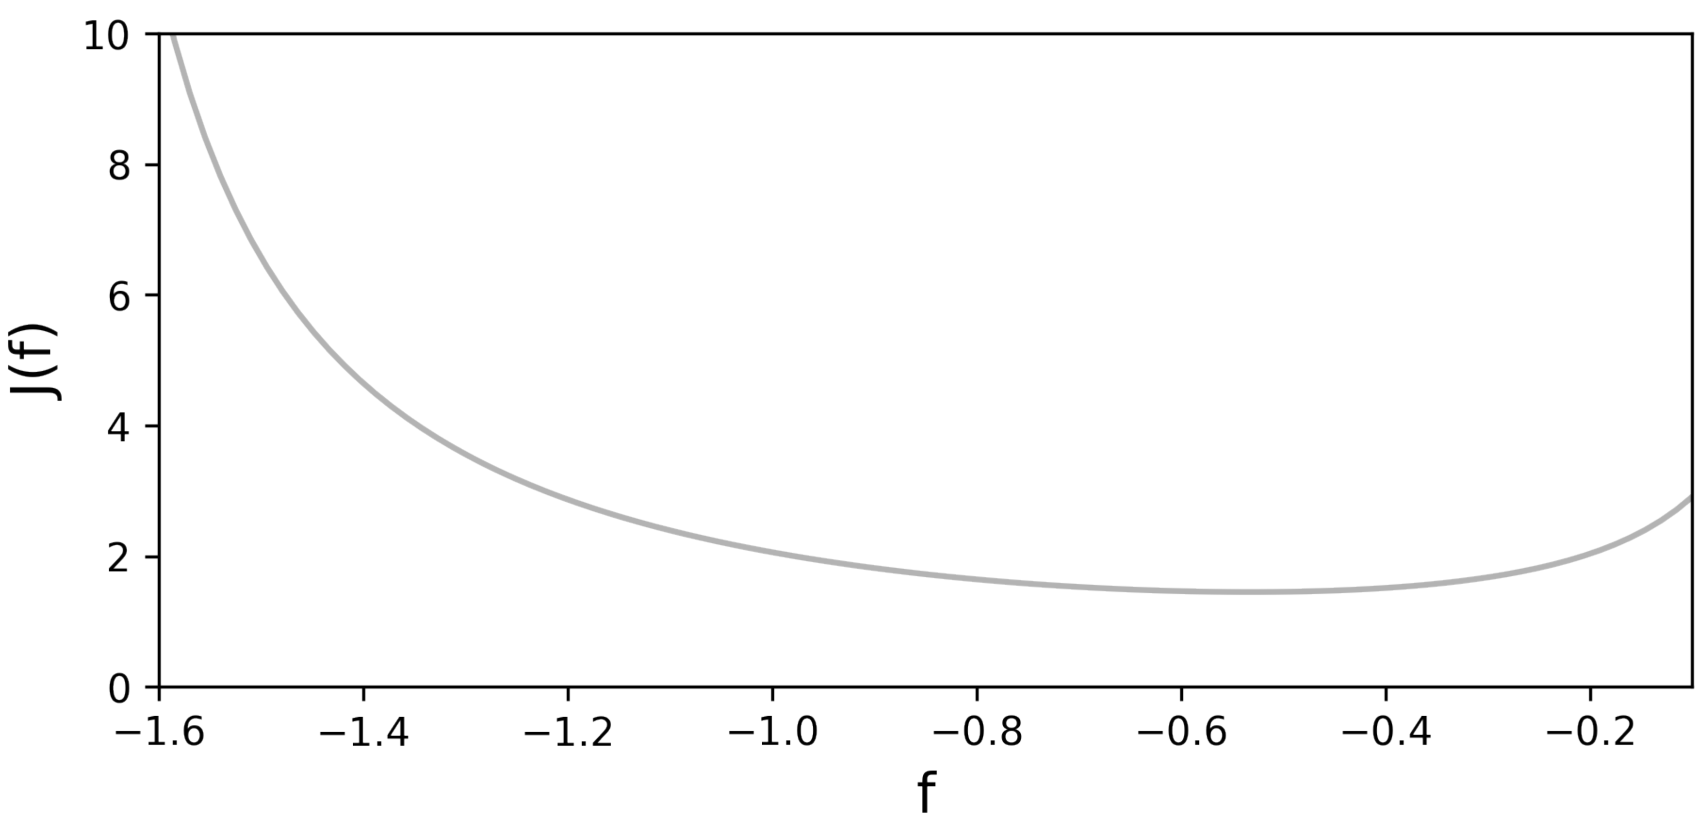

Non-parametric LQR kernel

Consider a scalar LTI system $(\hat{a}, \hat{b})$

$$x_{t+1} = \hat{a}x_t + \hat{b}u_t + v_t,\ v_t \sim \mathcal{N}(0, v)$$

Uncertain model (a, b)

$$x_{t+1} = ax_t + bu_t + v_t, \quad v_t \sim \mathcal{N}(0, v)$$
$$a \in [a_{\min}, a_{\max}], b \in [b_{\min}, b_{\max}]$$

Stochastic cost: m features

$$J_{\mathrm{LQR}}(f) = \underbrace{\left[\phi_{(a_1,b_1)}(f)\ \ \phi_{(a_2,b_2)}(f)\ \cdots\ \phi_{(a_m,b_m)}(f)\right]}_{=:\Phi^{\mathrm{T}}(f)} w = \Phi^{\mathrm{T}}(f)w,\ w \in \mathbb{R}^m, w \sim \mathcal{N}(0, \Sigma_w)$$

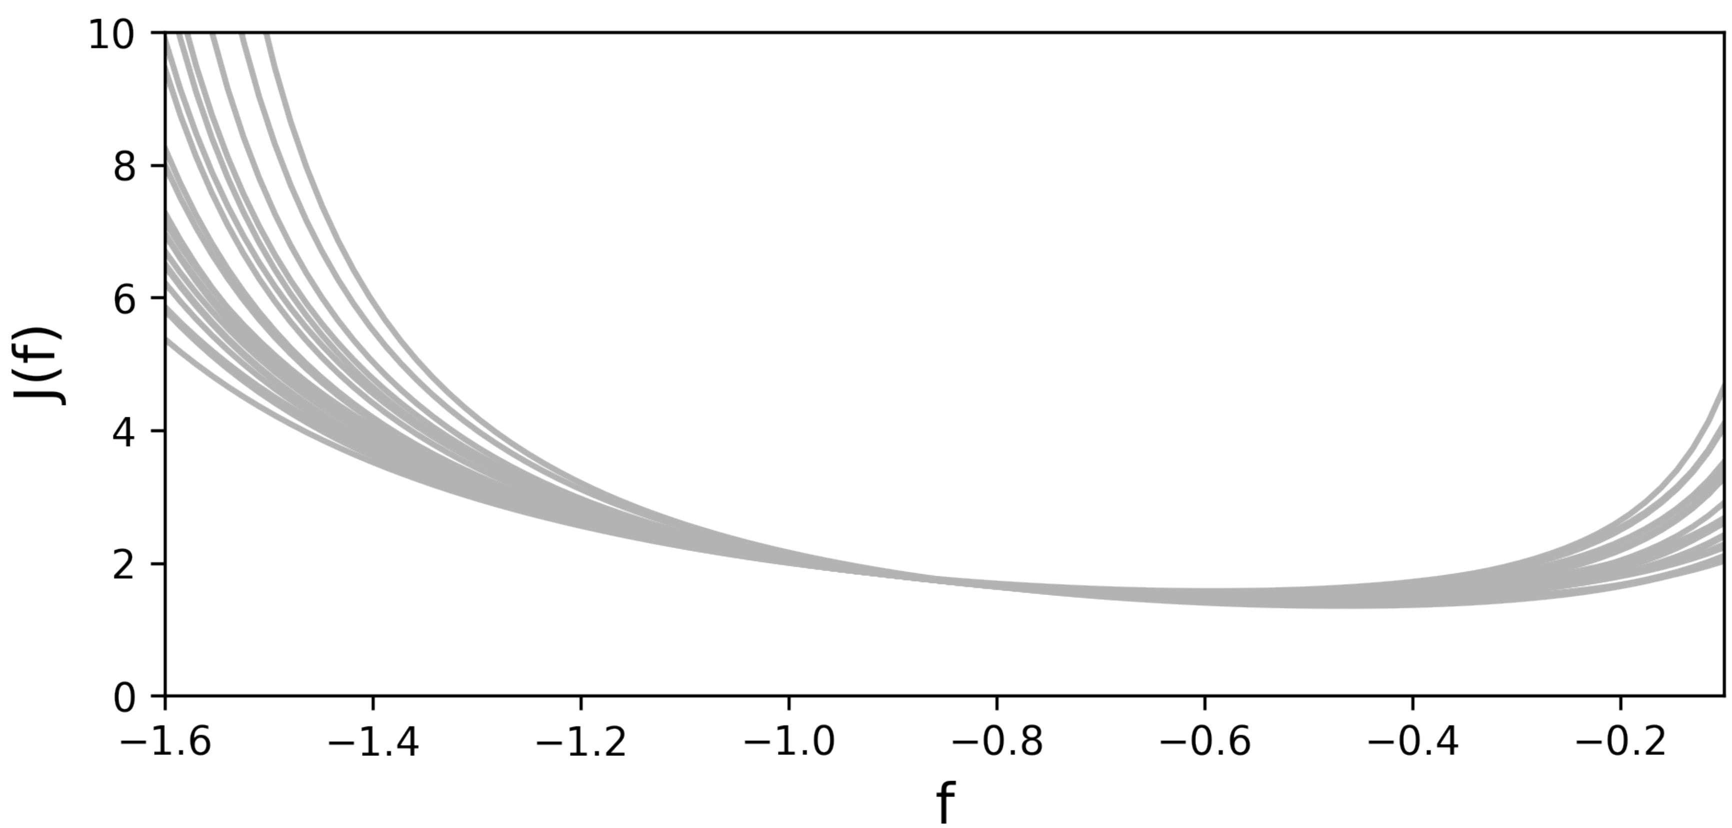

Non-parametric LQR kernel

Consider a scalar LTI system $(\hat{a}, \hat{b})$

$$x_{t+1} = \hat{a} x_t + \hat{b} u_t + v_t,\ v_t \sim \mathcal{N}(0, v)$$

Uncertain model (a, b)

$$x_{t+1} = a x_t + b u_t + v_t, \quad v_t \sim \mathcal{N}(0, v)$$
$$a \in [a_{\min}, a_{\max}], b \in [b_{\min}, b_{\max}]$$

Stochastic cost: m features

$$J_{\mathrm{LQR}}(f) = \underbrace{\left[\phi_{(a_1,b_1)}(f)\ \ \phi_{(a_2,b_2)}(f)\ \cdots\ \phi_{(a_m,b_m)}(f)\right]}_{=:\Phi^{\mathrm{T}}(f)} w = \Phi^{\mathrm{T}}(f) w,\ w \in \mathbb{R}^m, w \sim \mathcal{N}(0, \Sigma_w)$$

Parametric LQR kernel with m features

$$k_{\mathrm{pLQR},m}(f, f') = \Phi^{\mathrm{T}}(f)\, \Sigma_w\, \Phi(f')$$

Non-parametric LQR kernel

Consider a scalar LTI system $(\hat{a}, \hat{b})$

$$x_{t+1} = \hat{a} x_t + \hat{b} u_t + v_t,\ v_t \sim \mathcal{N}(0, v)$$

Uncertain model (a, b)

$$x_{t+1} = a x_t + b u_t + v_t, \quad v_t \sim \mathcal{N}(0, v)$$
$$a \in [a_{\min}, a_{\max}], b \in [b_{\min}, b_{\max}]$$

Stochastic cost: m features

$$J_{\mathrm{LQR}}(f) = \underbrace{\left[\phi_{(a_1,b_1)}(f)\ \ \phi_{(a_2,b_2)}(f)\ \cdots\ \phi_{(a_m,b_m)}(f)\right]}_{=:\Phi^{\mathrm{T}}(f)} w = \Phi^{\mathrm{T}}(f) w,\ w \in \mathbb{R}^m, w \sim \mathcal{N}(0, \Sigma_w)$$

Parametric LQR kernel with m features

$$k_{\mathrm{pLQR},m}(f, f') = \Phi^{\mathrm{T}}(f)\, \Sigma_w\, \Phi(f')$$

Kernel trick

$$\Sigma_w = \sigma_w^2 I,\ \sigma_w \propto 1/m$$

Use as many features as possible: $m \to \infty$

$$k_{\mathrm{nLQR}}(f, f') = \lim_{m \to \infty} k_{\mathrm{pLQR},m}(f, f')$$

Non-parametric LQR kernel

$$= \sigma_{\mathrm{n}}^2 \int_{b_{\min}}^{b_{\max}} \int_{a_{\min}}^{a_{\max}} \phi_{(a,b)}(f) \phi_{(a,b)}(f')\, da\, db$$

Non-parametric LQR kernel

Uncertainty ranges $a \in [0.8, 1.0]\,,\ b \in [0.9, 1.1]$

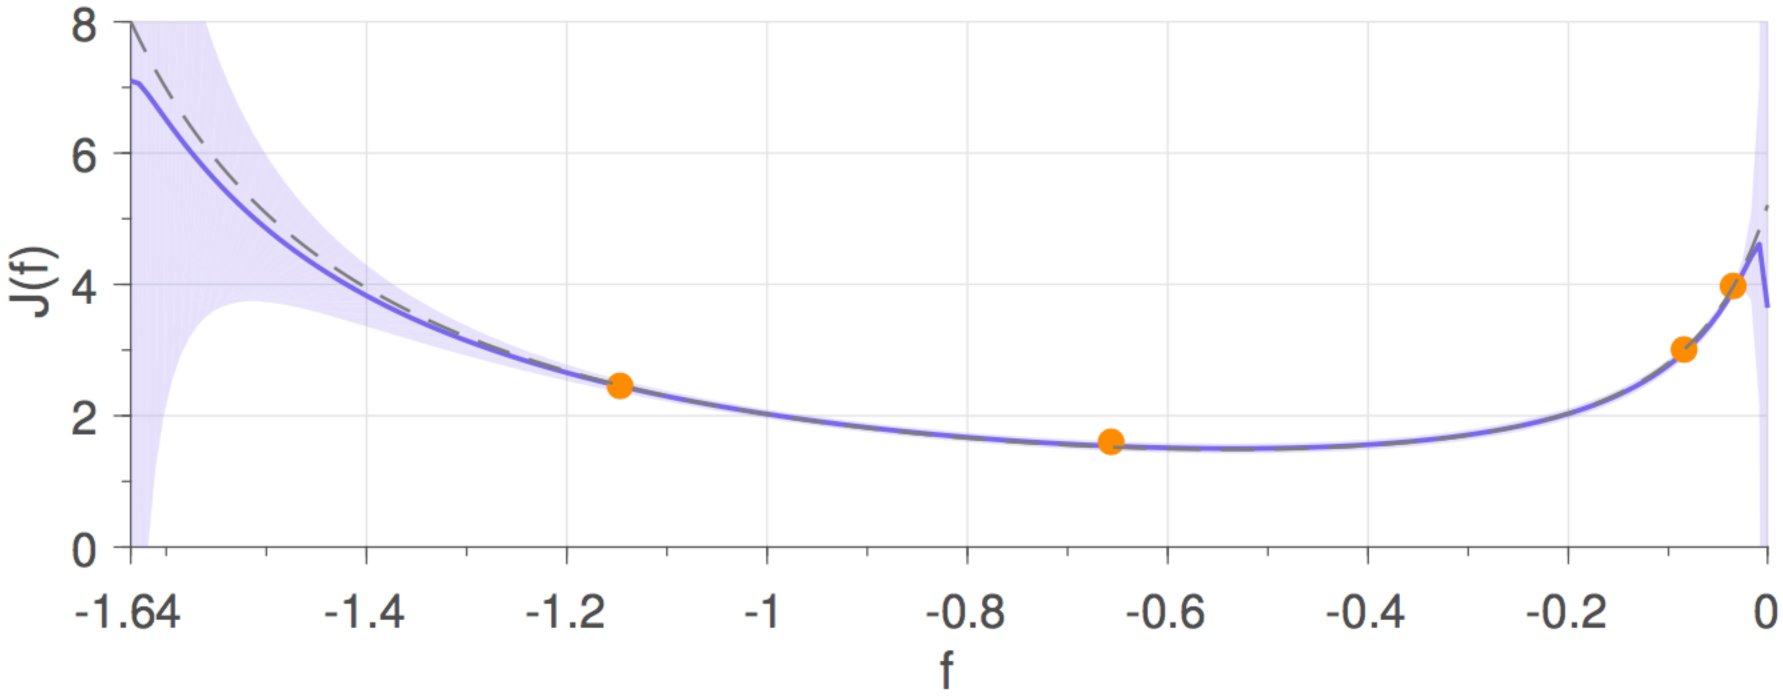

$$(\hat{a}, \hat{b}) = (0.9, 1.0)$$

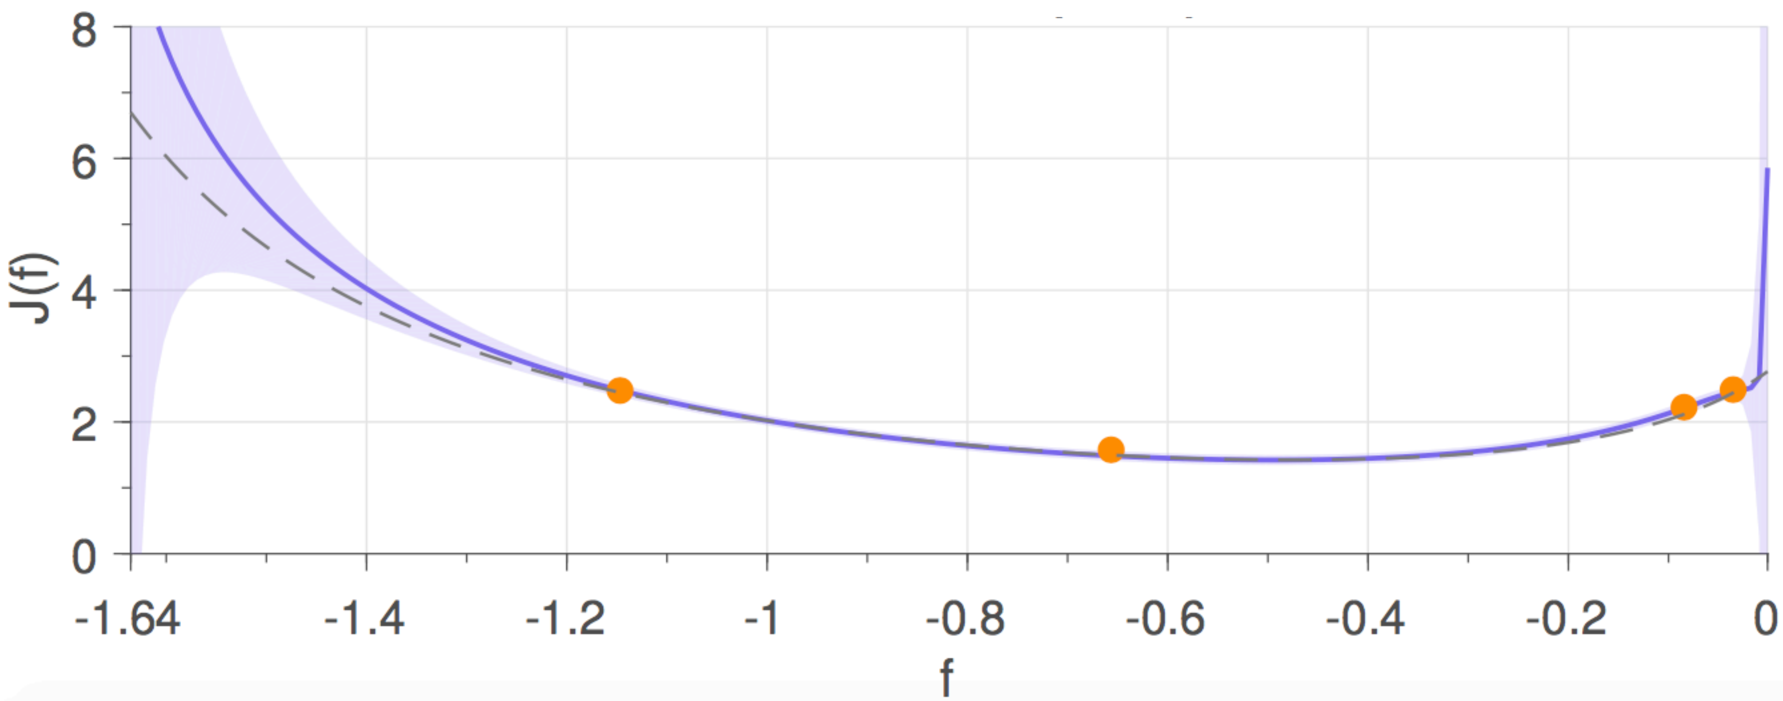

$$(\hat{a}, \hat{b}) = (0.8, 0.9)$$

Goal

✴Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Steps

✓Parametric LQR kernel

✓Non-parametric LQR kernel

✴Simulated results

Simulations

LQR kernel fitting performance: 1000 simulations, 2 evaluations

Kernel	RMSE
Squared exponential	2.49
Parametric LQR kernel	**1.02**
Parametric LQR kernel - infer (a,b)	**1.09**
Non-parametric LQR kernel	**1.22**

Bayesian optimization: 100 runs, 2 evaluations

Kernel	Regret
Squared exponential	1.34
Parametric LQR kernel	**0.30**
Parametric LQR kernel - infer (a,b)	**0.32**
Non-parametric LQR kernel	**0.35**

Combined kernel

Non-linear system

$$x_{t+1} = h(x_t, u_t, v_t)$$

Uncertain model (a, b)

$$x_{t+1} = ax_t + bu_t + v_t, \quad v_t \sim \mathcal{N}(0, v)$$
$$a \in [a_{\min}, a_{\max}], b \in [b_{\min}, b_{\max}]$$

$$J(f) = J_{\mathrm{LQR}}(f) + \underbrace{J_{\Delta}(f)}_{\text{Error term}}$$

$$J(f) \sim \mathcal{GP}(0, k_{\mathrm{LQR}}(f) + k_{\mathrm{SE}}(f))$$

Goal

✓Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Steps

✓Parametric LQR kernel

✓Non-parametric LQR kernel

✓Simulated results

Conclusions

✓Parametric LQR kernel

✓Non-parametric LQR kernel

✓Sample efficiency for LQR problems

✓Non-linear systems

Goal

✓Incorporate LQR controller structure into kernel

Consider

✓Scalar linear system

Steps

✓Parametric LQR kernel

✓Non-parametric LQR kernel

✓Simulated results

Conclusions

✓Parametric LQR kernel

✓Non-parametric LQR kernel

✓Sample efficiency for LQR problems

✓Non-linear systems

Future work

✴Extensions to vector systems

✴Real system experiments

Thank you for your attention

Backup slides

✳Gaussian process: mean and kernel

$$J(\theta) \sim \mathcal{GP}(\mu(\theta), k(\theta, \theta'))$$

✳kernel defines correlation

$$k(\theta, \theta') = \mathbb{C}\mathrm{ov}\left(J(\theta), J(\theta')\right)$$

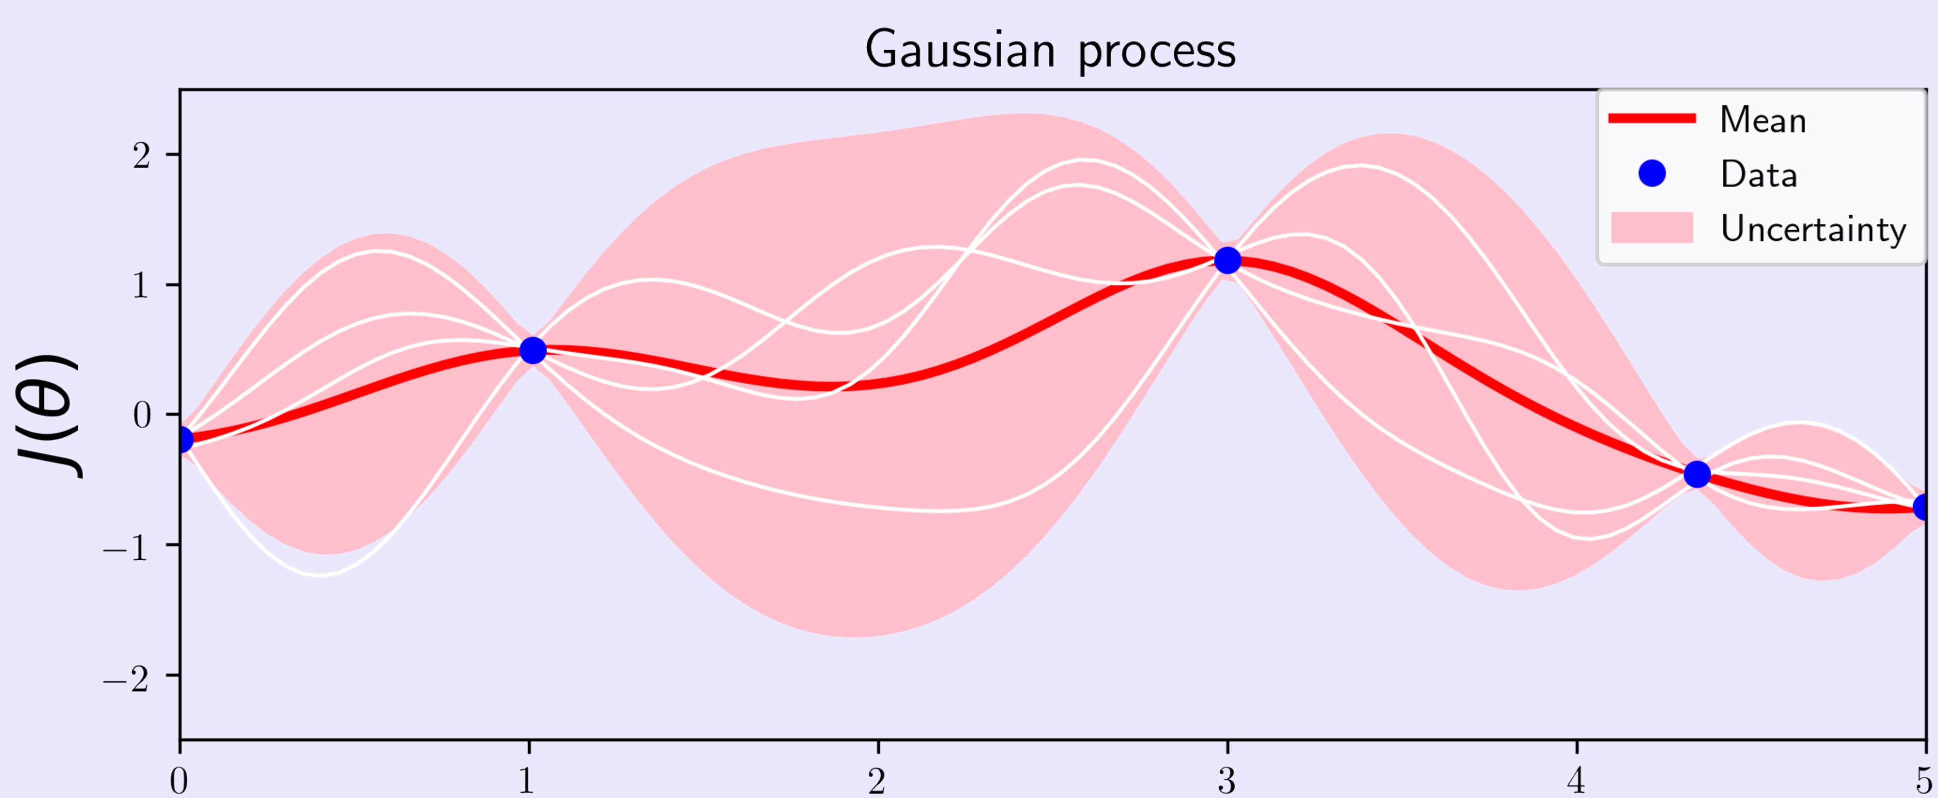

Iterat

$\theta_i,$

Cost $J(\theta)$

$J : \mathbb{R}^M \to \mathbb{R}$

θ_{i+1}

✳Gaussian process: mean and kernel

$$J(\theta) \sim \mathcal{GP}(\mu(\theta), k(\theta, \theta'))$$

✳kernel defines correlation

$$k(\theta, \theta') = \mathbb{C}\mathrm{ov}\left(J(\theta), J(\theta')\right)$$

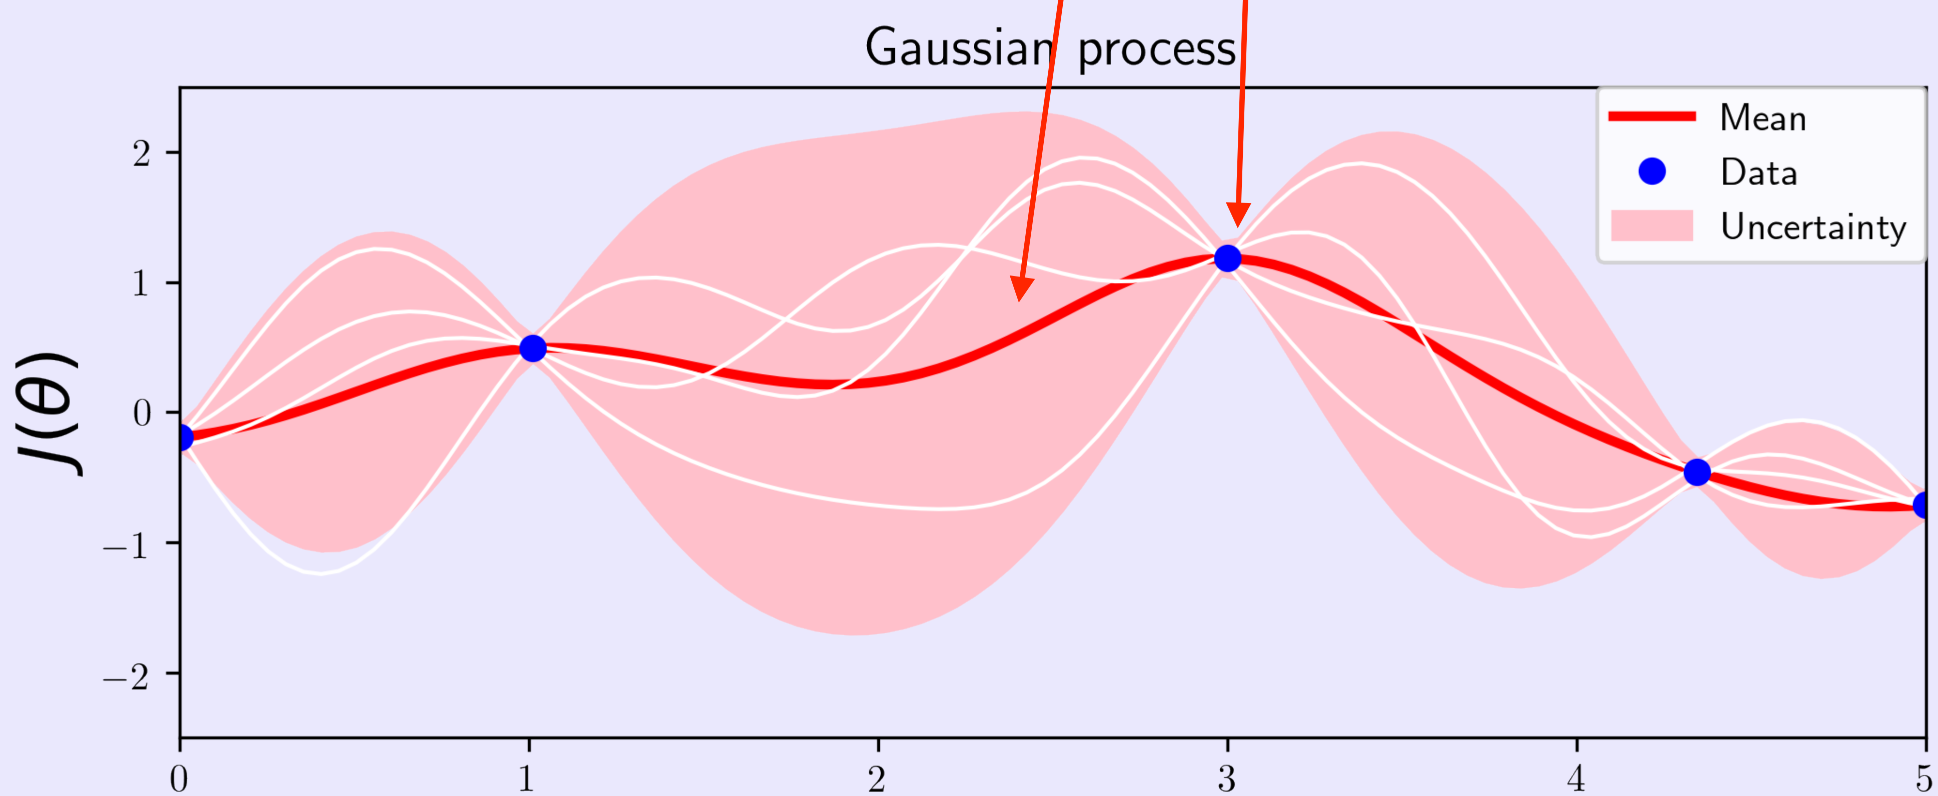

Alonso Marco | Max Planck Institute for Intelligent Systems

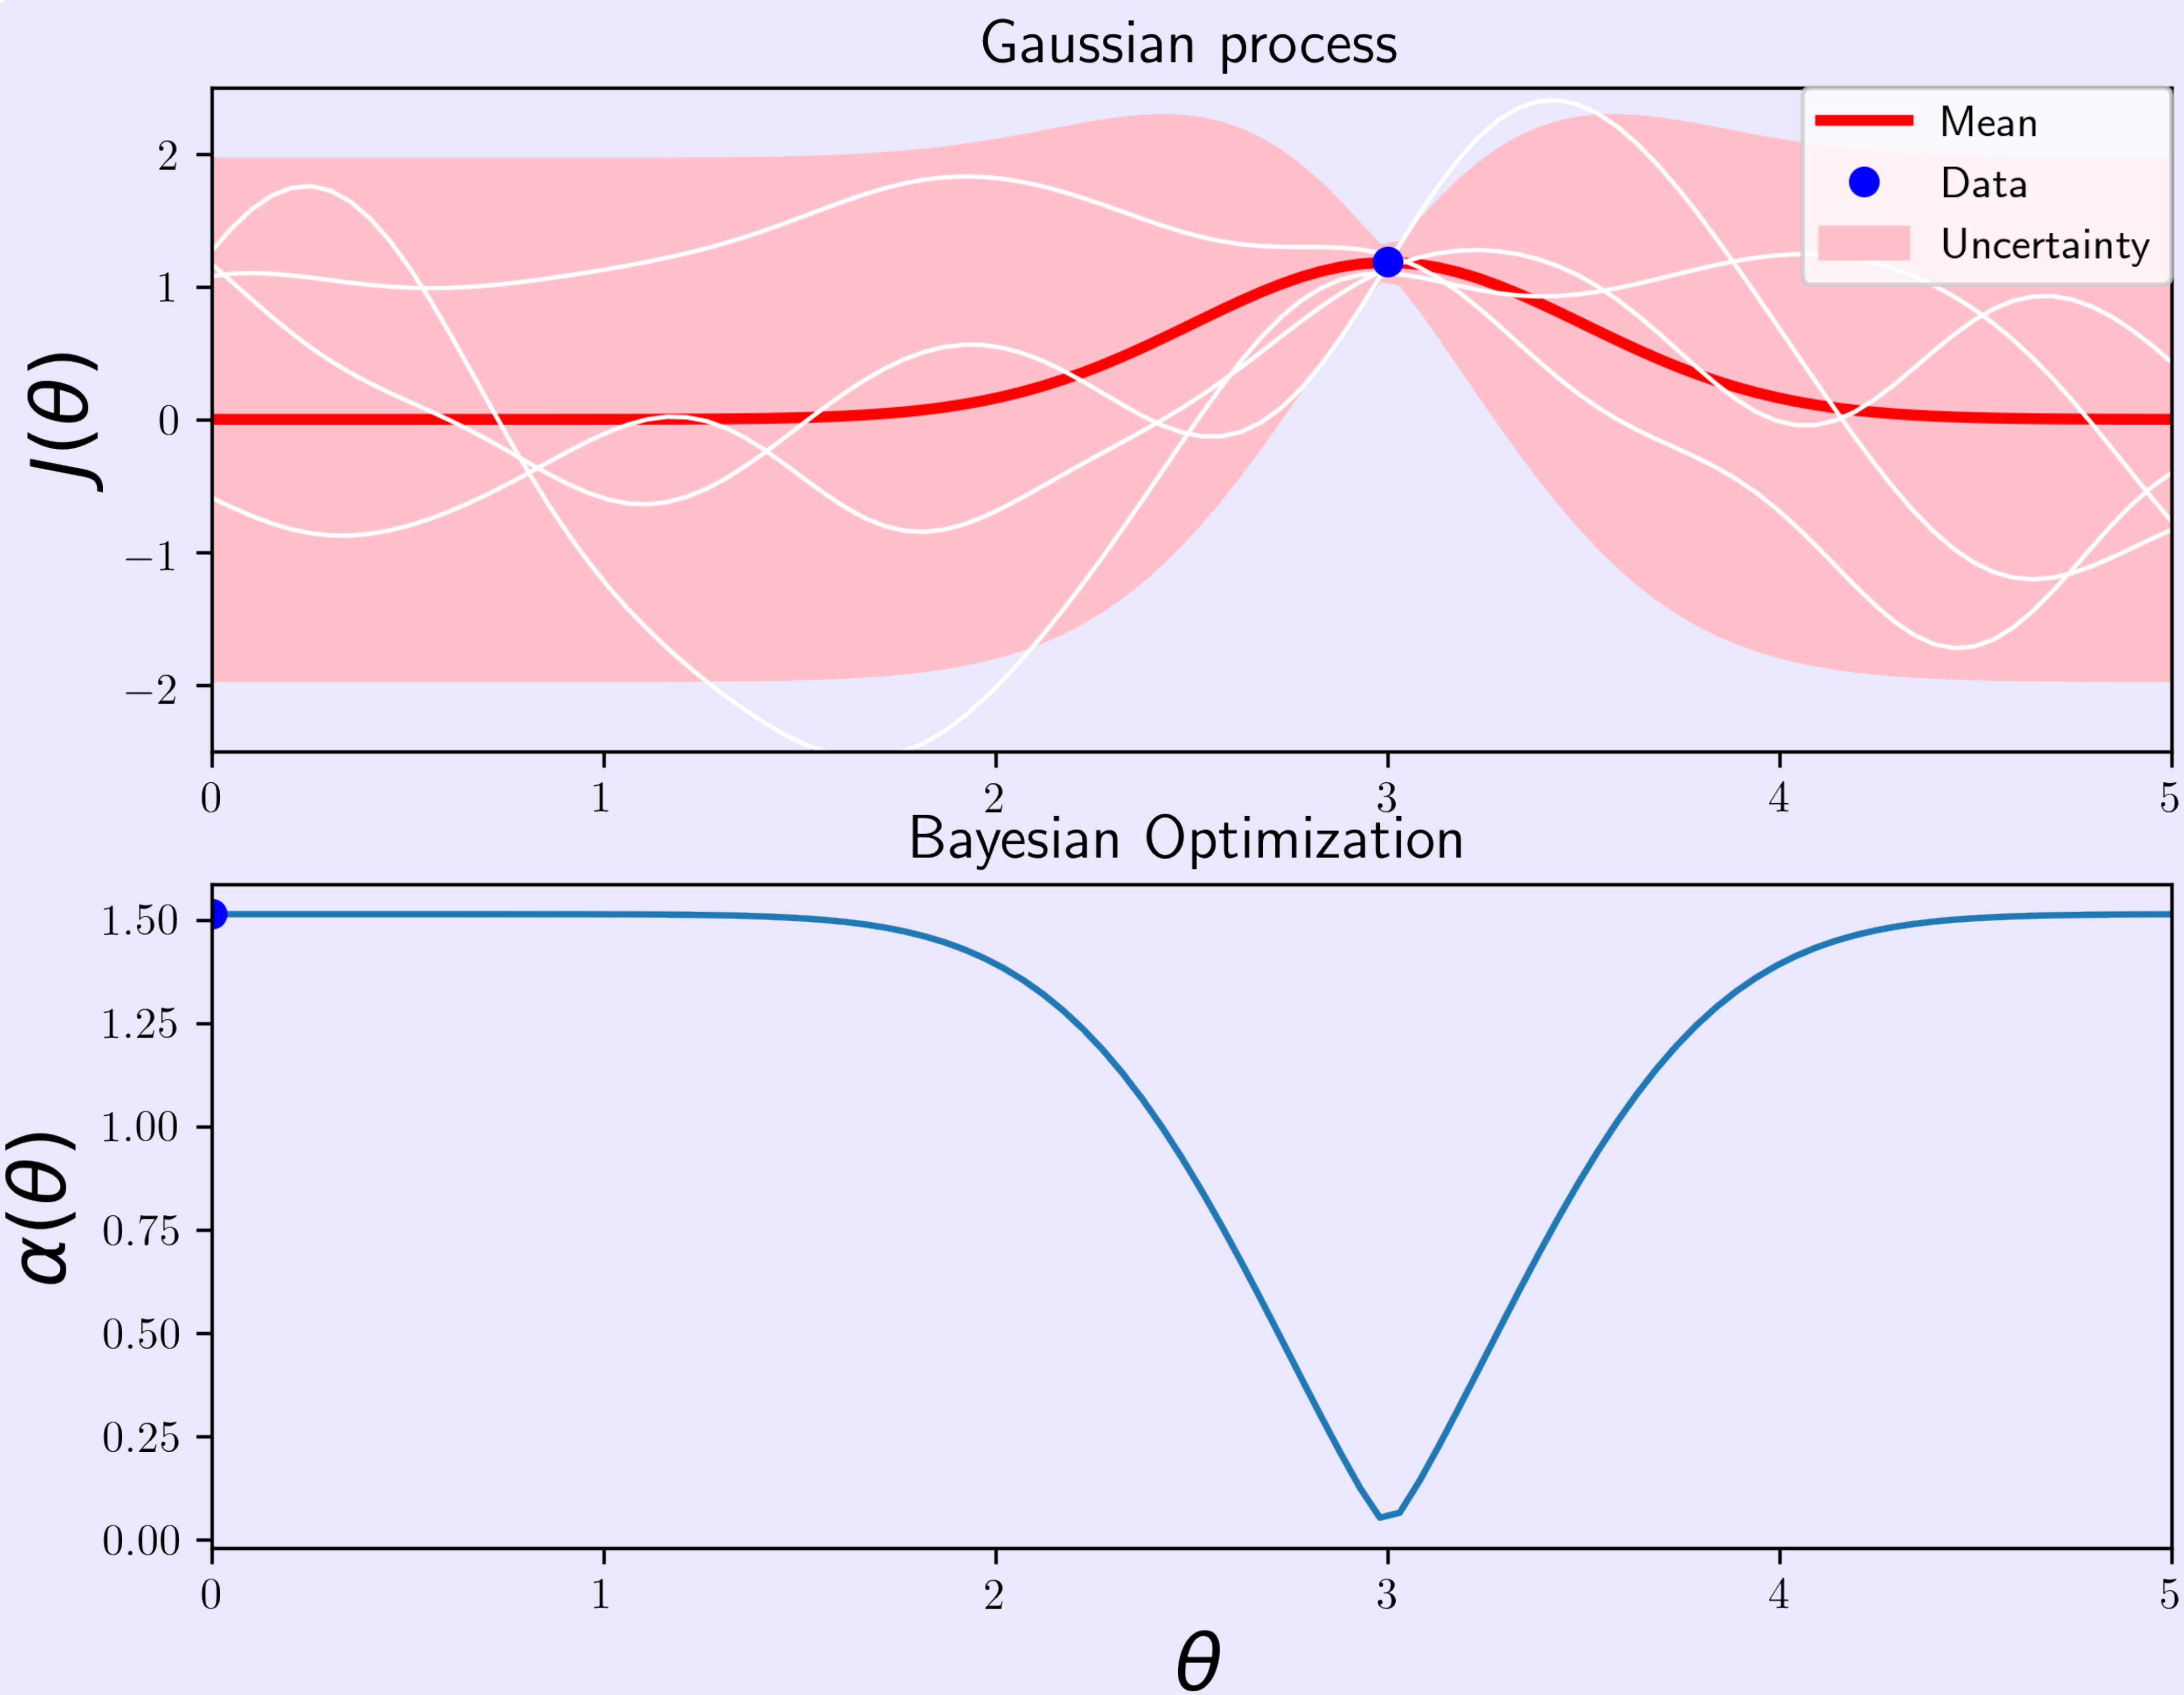

Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5
Bayesian Optimization
$\alpha(\theta)$
1.50
1.25
1.00
0.75
0.50
0.25
0.00
0
1
2
3
4
5
θ

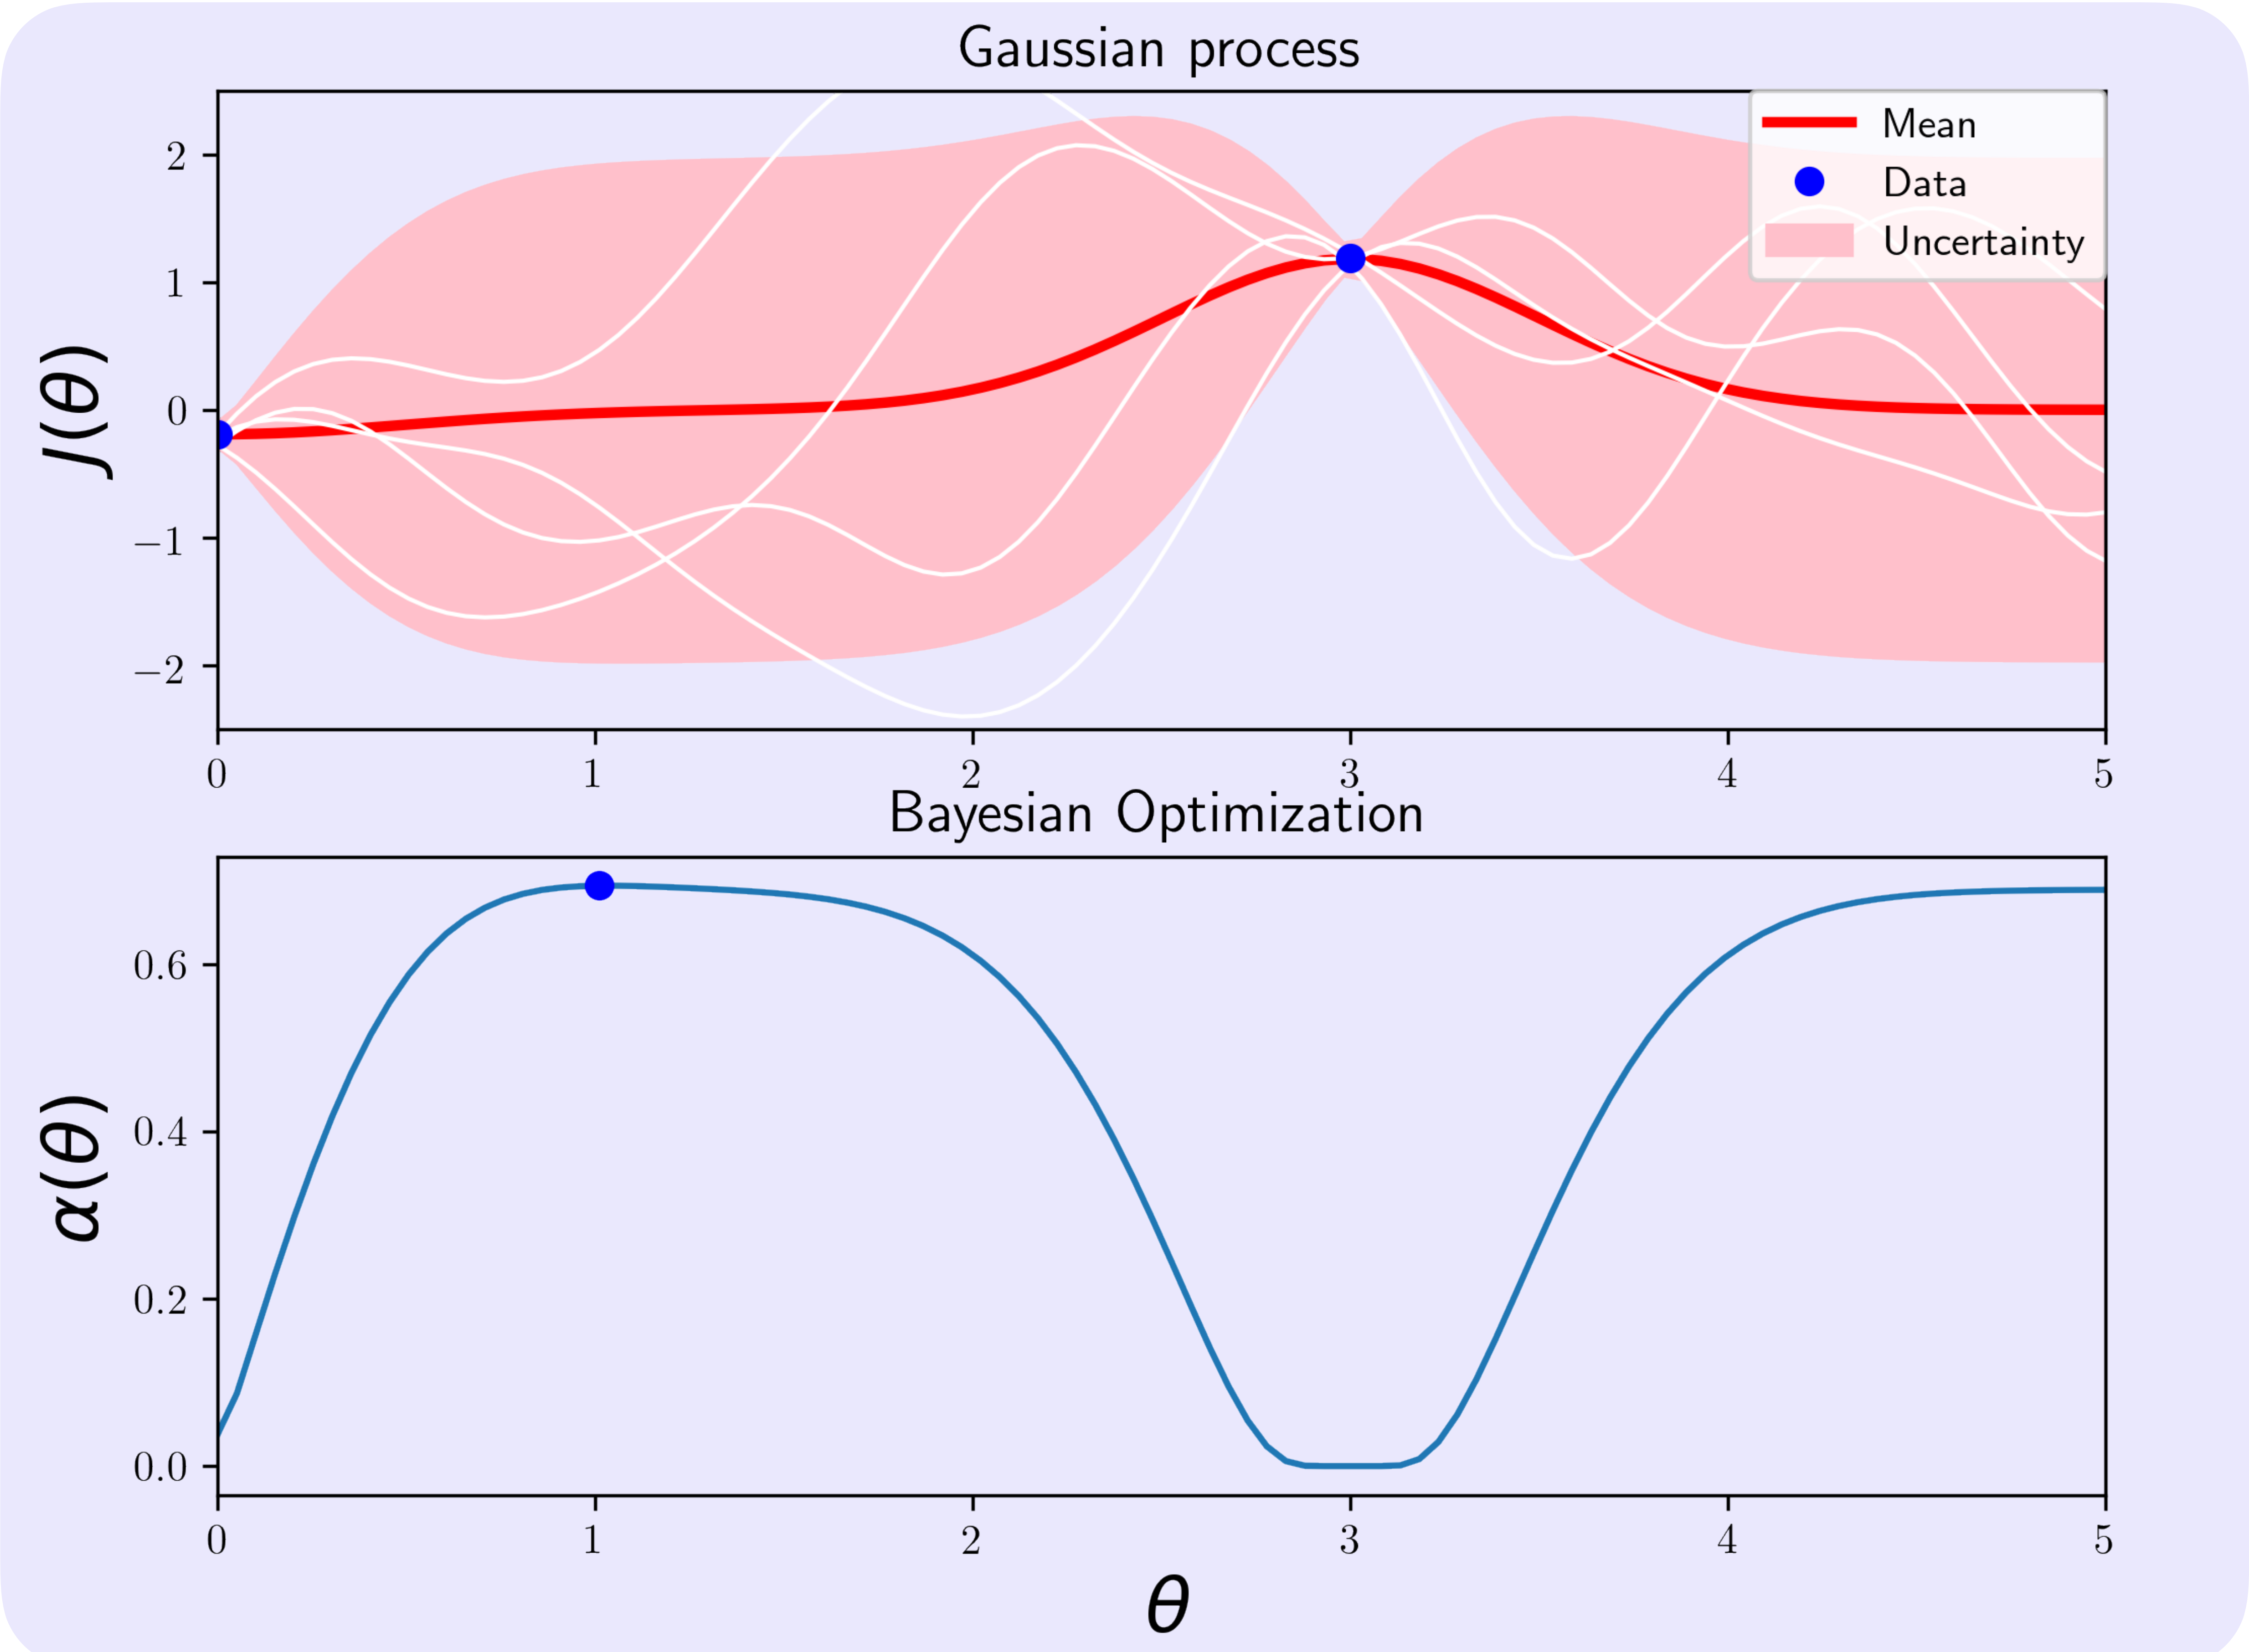

Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5
Bayesian Optimization
$\alpha(\theta)$
0.6
0.4
0.2
0.0
0
1
2
3
4
5
θ

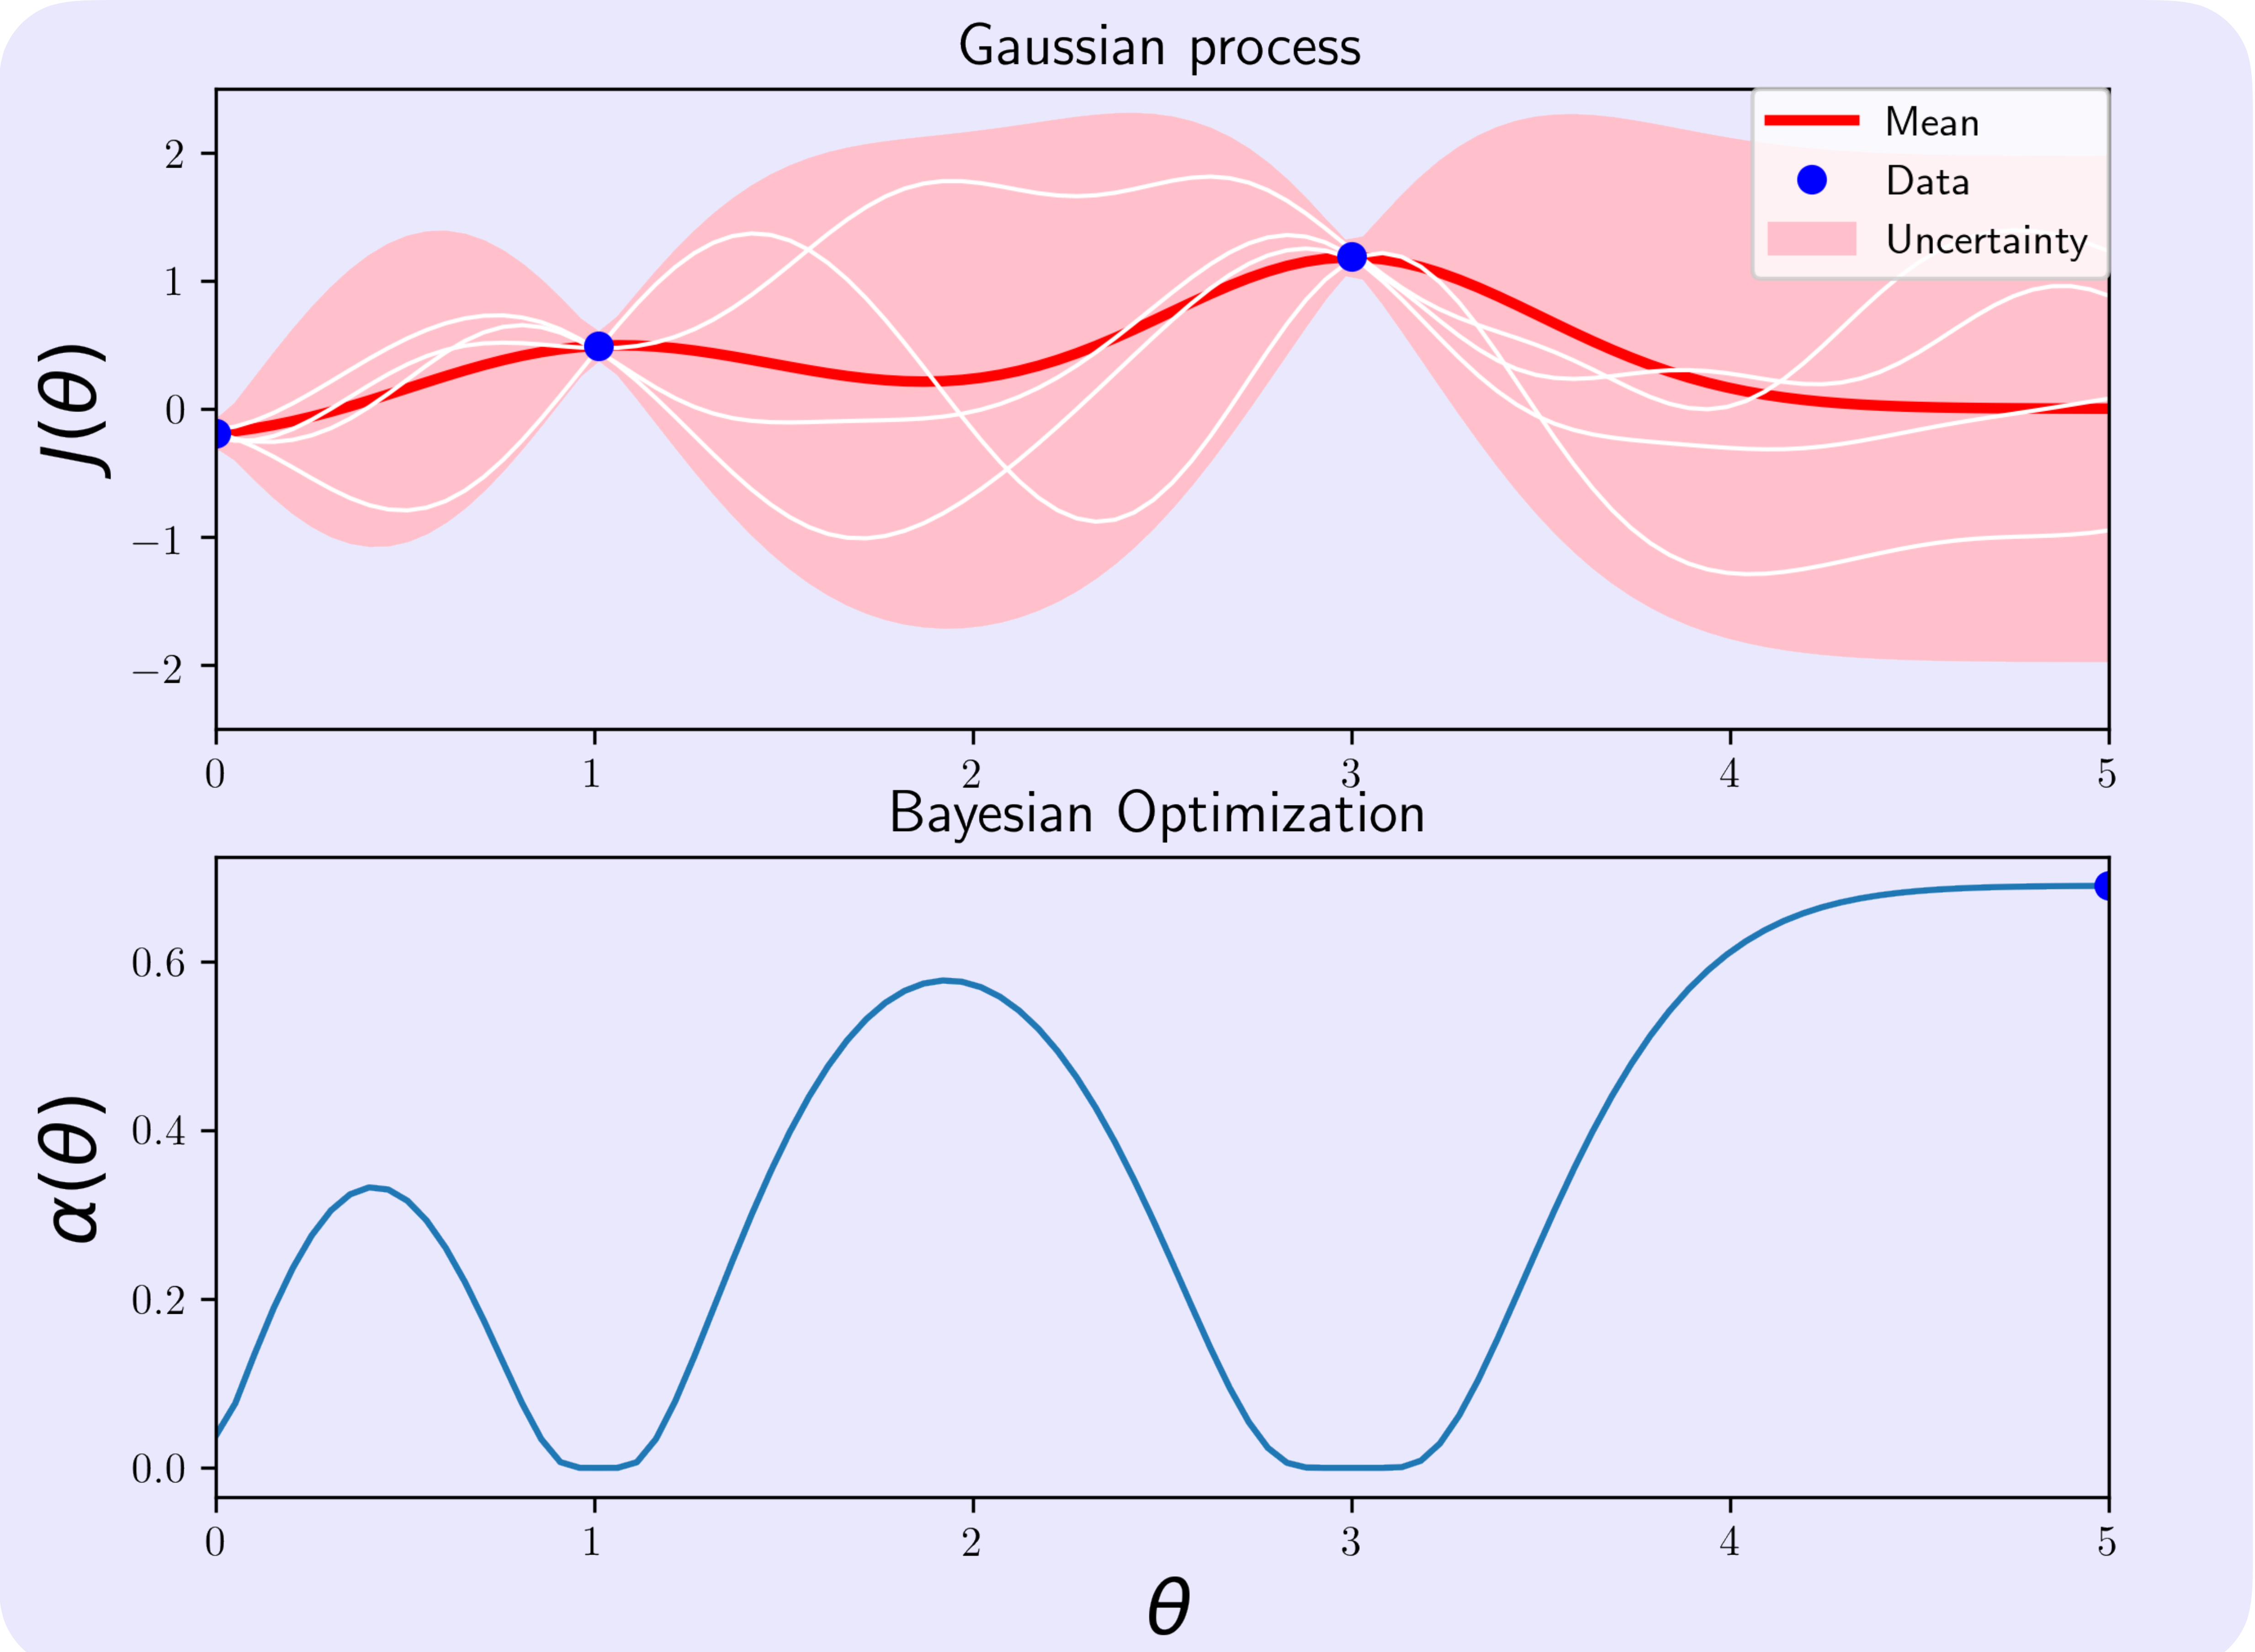
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5
Bayesian Optimization
$\alpha(\theta)$
0.6
0.4
0.2
0.0
0
1
2
3
4
5
θ

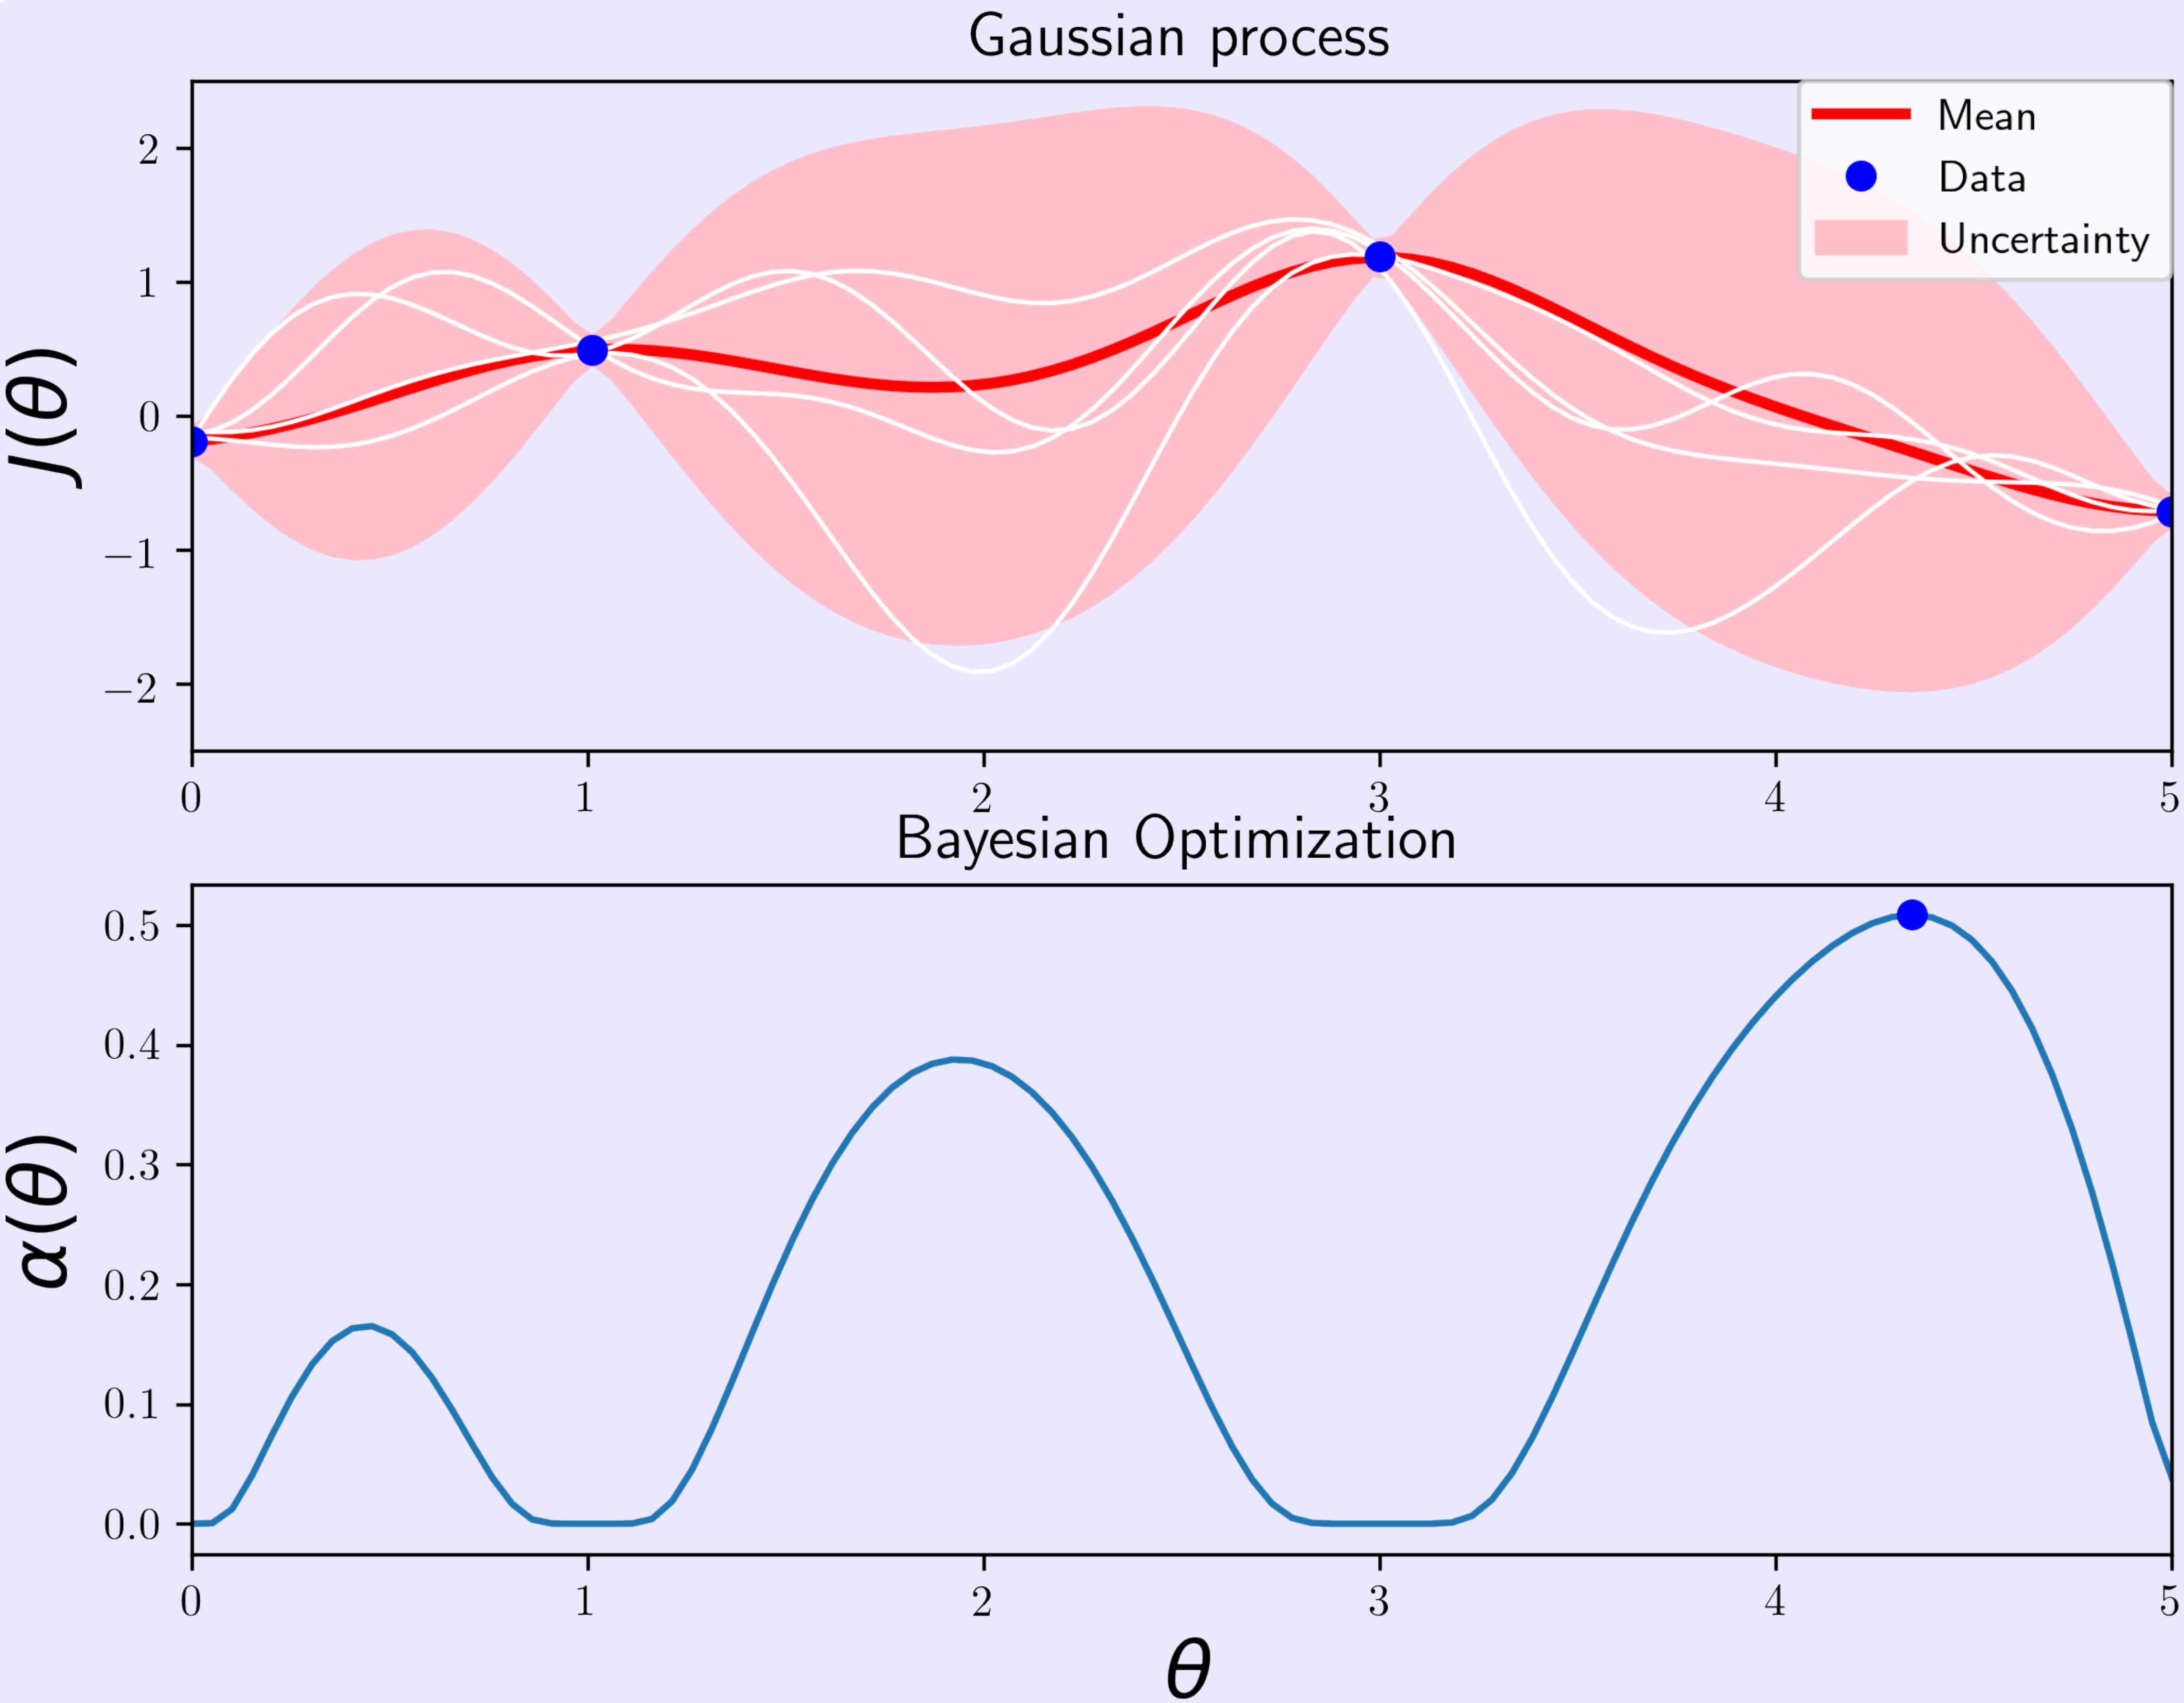
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5
Bayesian Optimization
$\alpha(\theta)$
0.5
0.4
0.3
0.2
0.1
0.0
0
1
2
3
4
5
θ

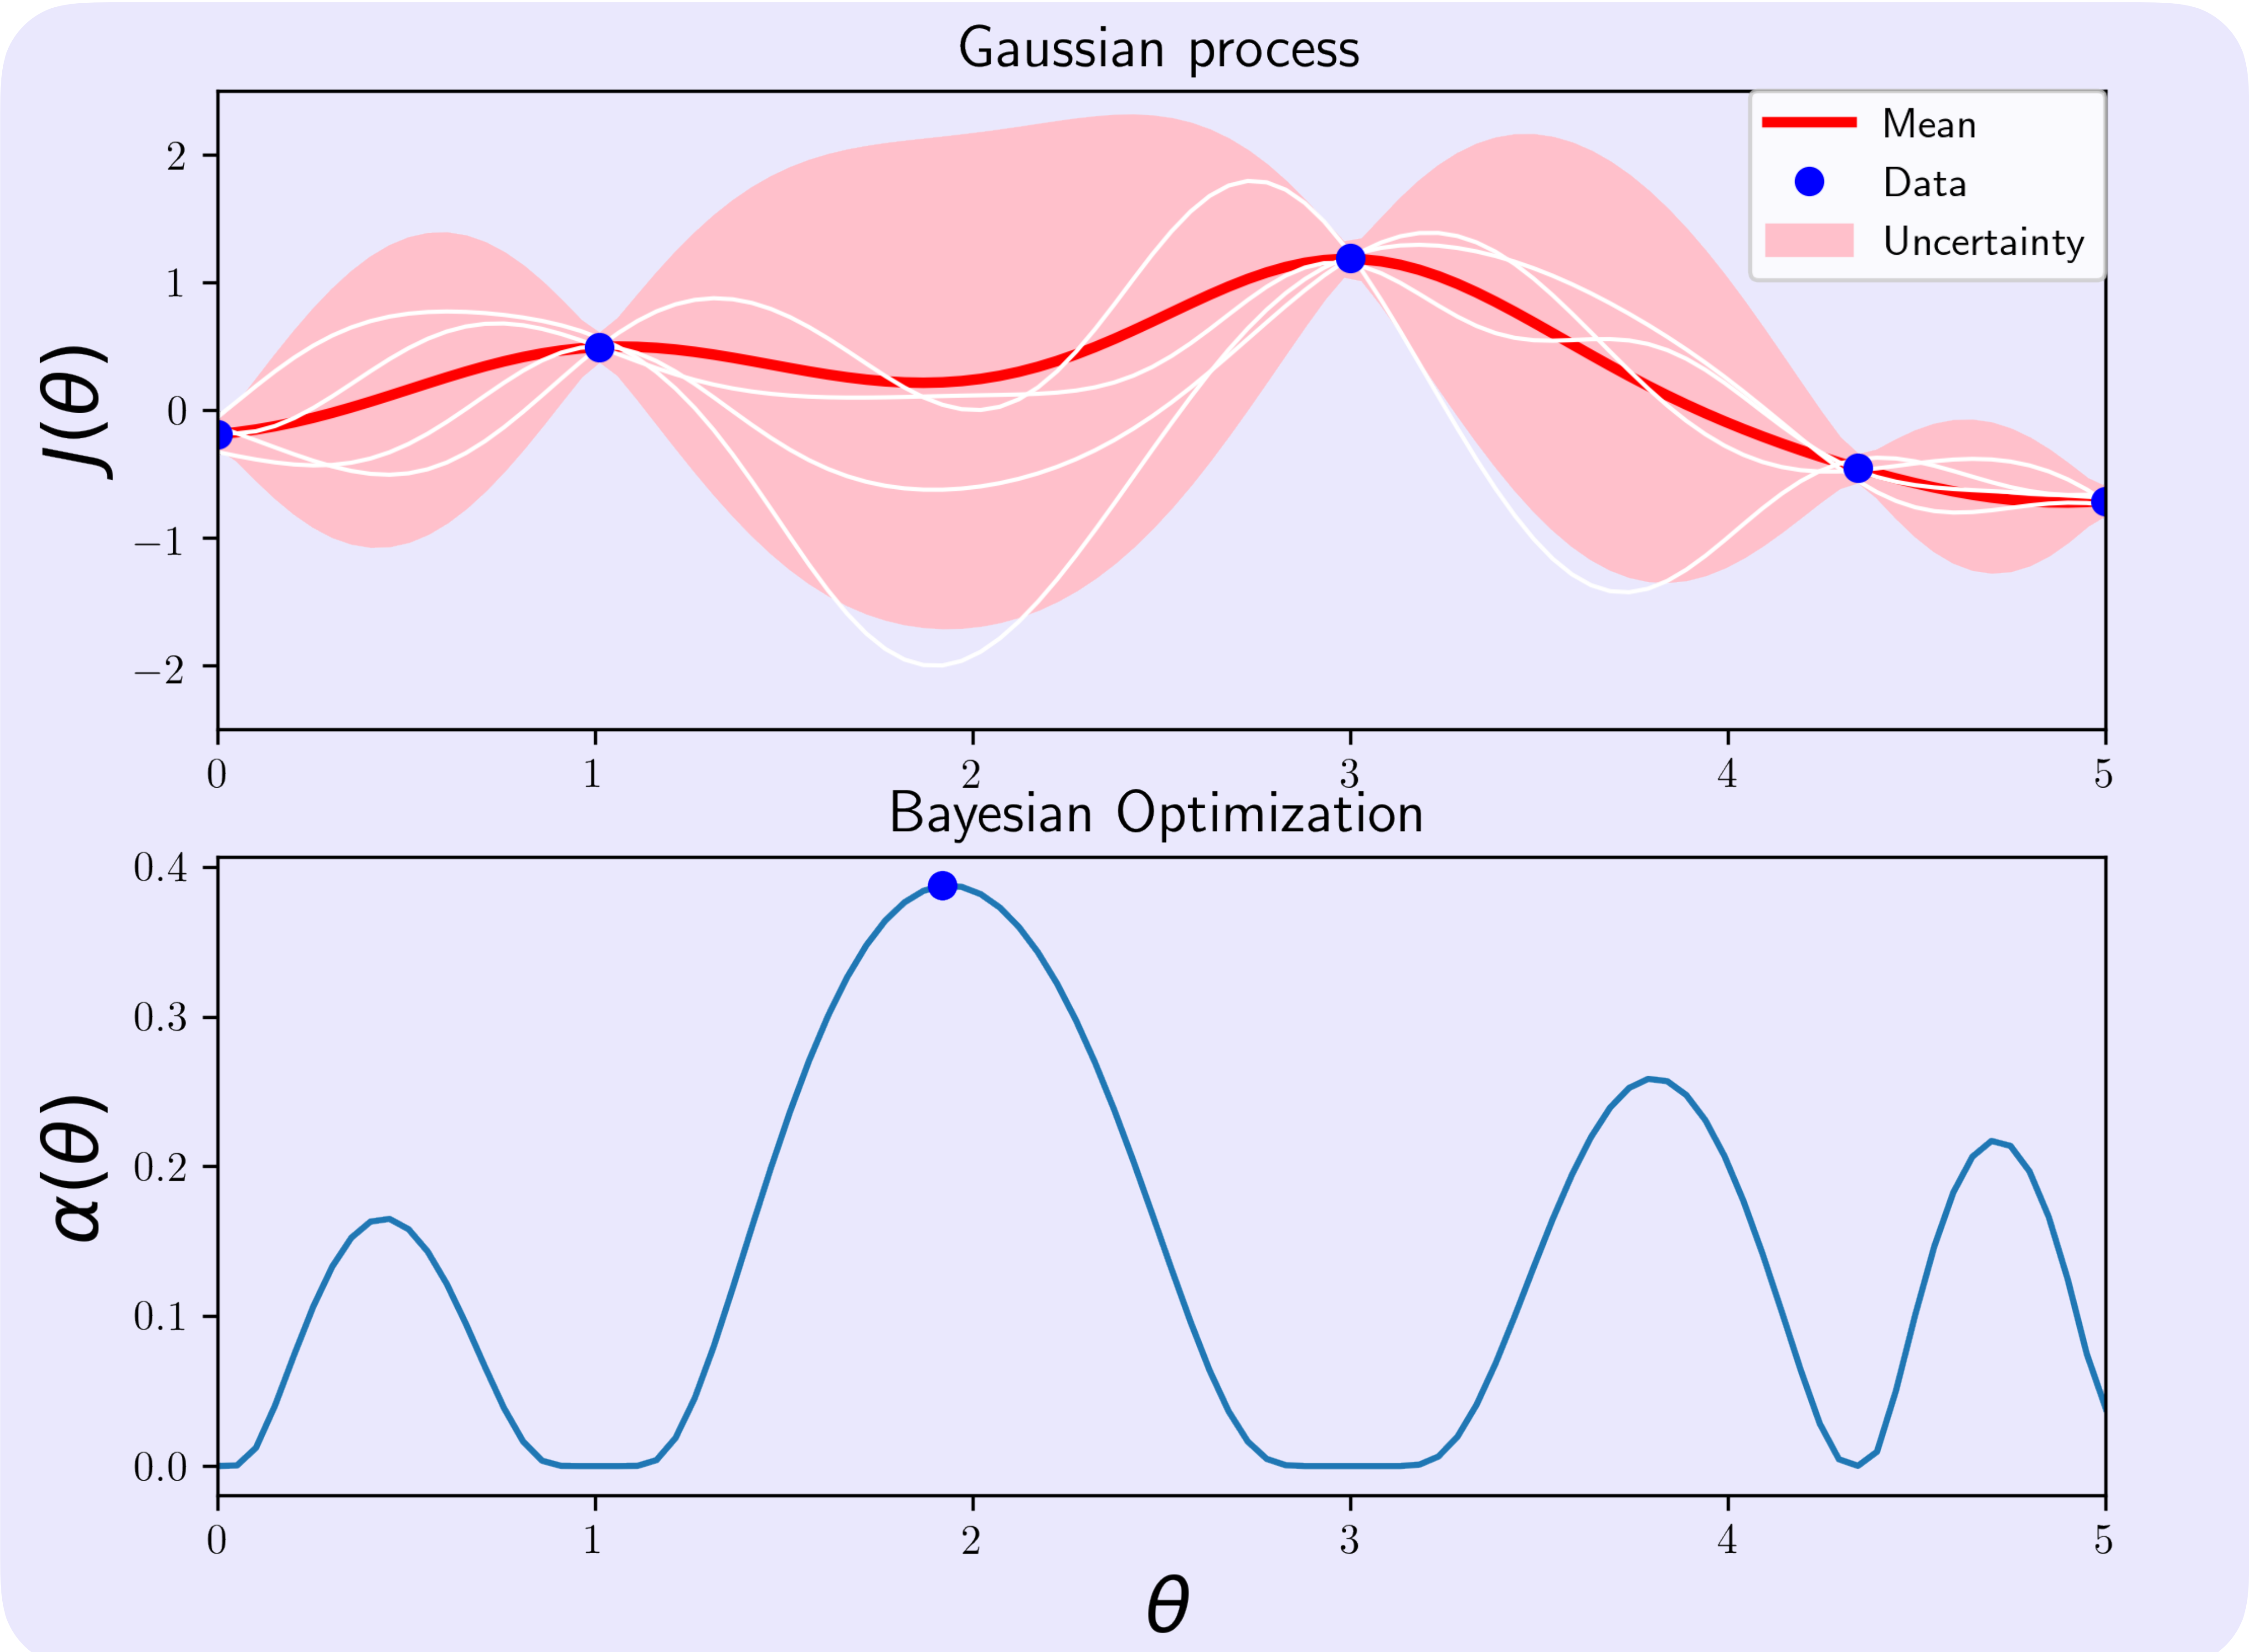
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
2
1
0
−1
−2
0
1
2
3
4
5
Bayesian Optimization
$\alpha(\theta)$
0.4
0.3
0.2
0.1
0.0
0
1
2
3
4
5
θ

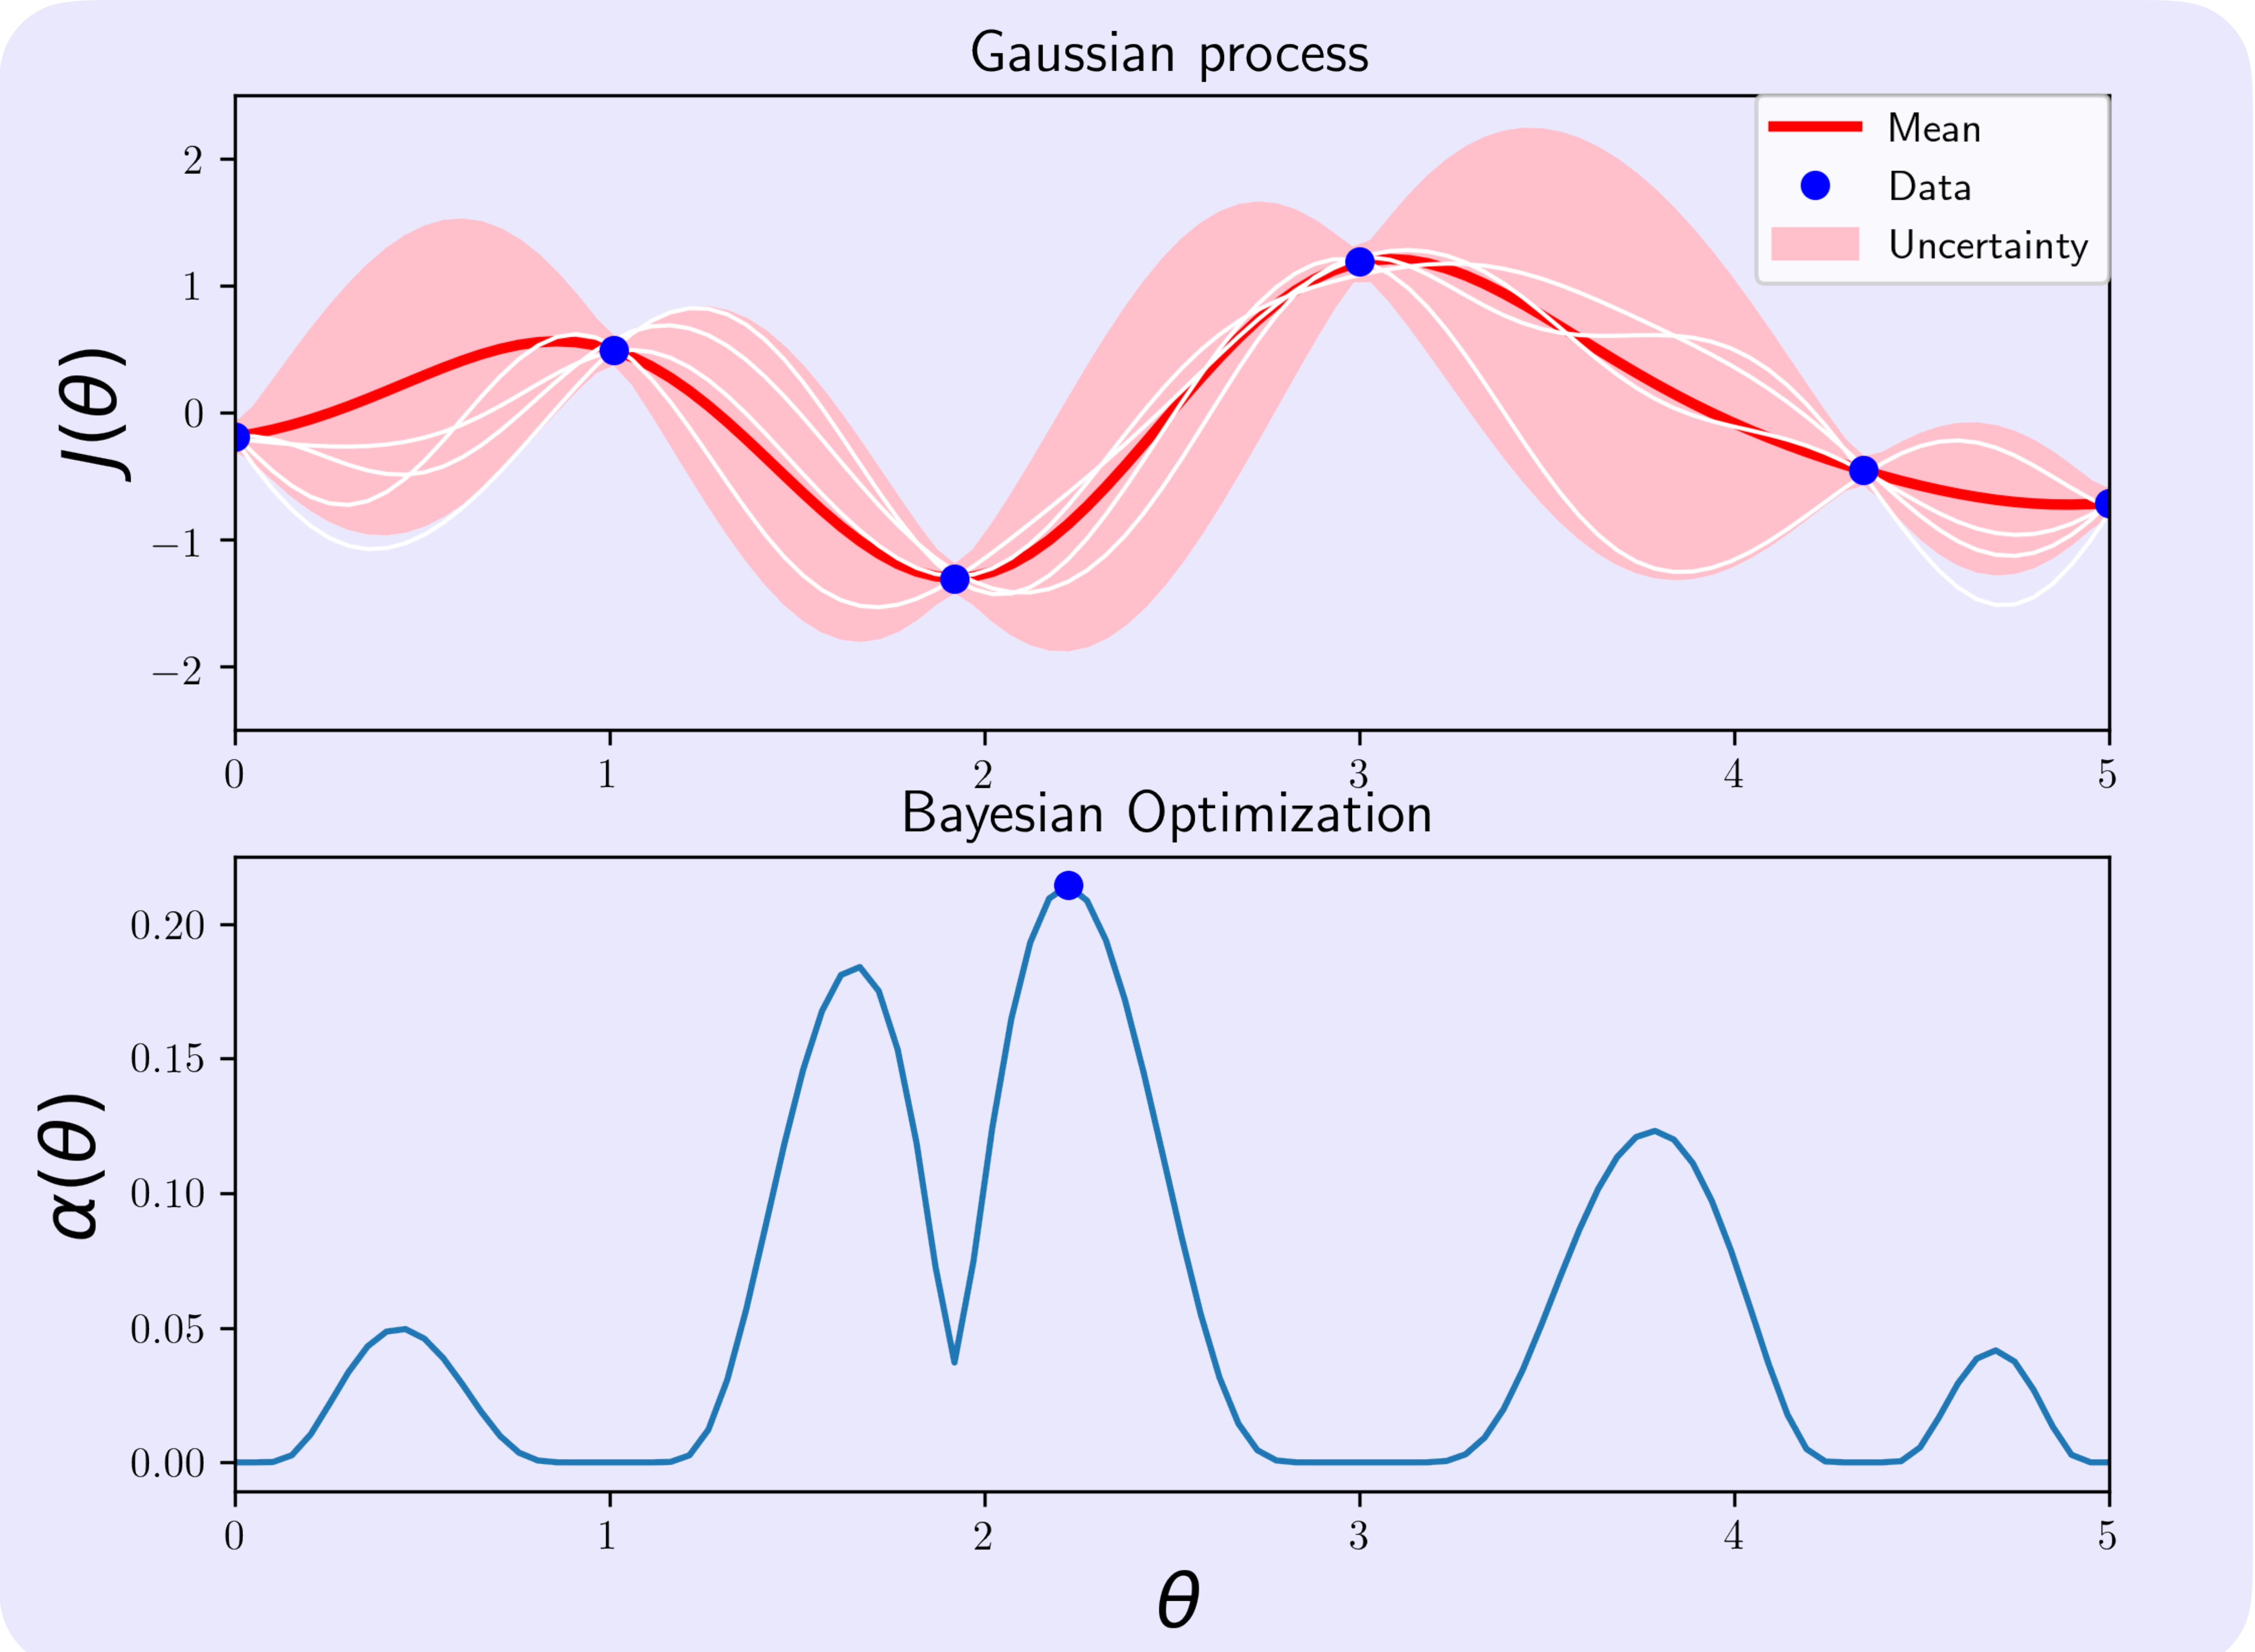
Gaussian process
Mean
Data
Uncertainty
$J(\theta)$
Bayesian Optimization
$\alpha(\theta)$
θ