Previously on MLRG

Closed loop control

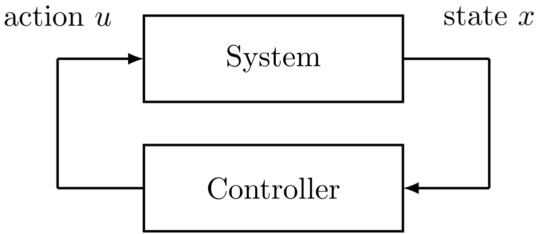

LQR

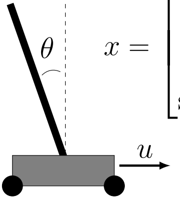

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

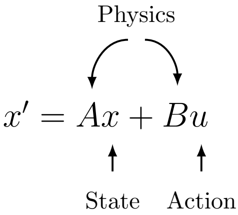

LQR

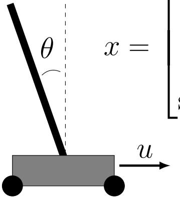

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

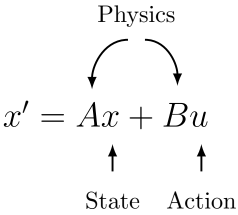

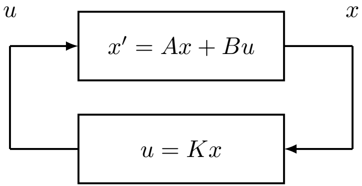

Solving for K

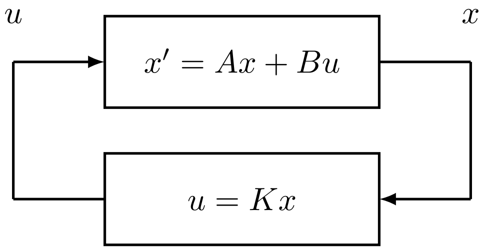

Solving for K

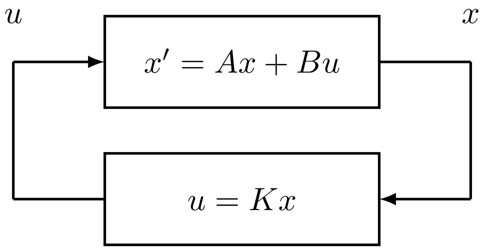

Go to x^* min control

$$J = \int (x - x^*)^\top Q \, (x - x^*) + u^\top R \, u \quad dt$$

Solving for K

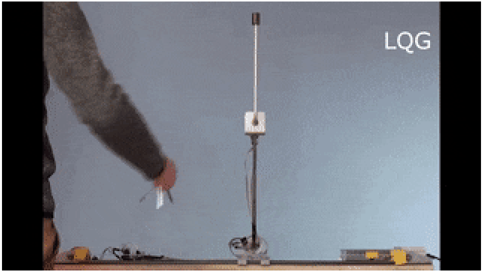

$$J = \int (x - x^*)^\top \underset{\text{Go to } x^*}{Q} (x - x^*) + u^\top \underset{\text{min control}}{R} u \quad dt$$

Learning A, B

$$x' = Ax + Bu$$

Learning A, B

$$x' = Ax + Bu$$

$$\min_{A,B} \|Ax + Bu - x'\|$$

Learning A, B

$$x' = Ax + Bu \qquad \min_{A,B} \|Ax + Bu - x'\|$$

  

Today

What if?

Today

What if?

- Nonlinear?
- Wrong model?

Today

What if?

- Nonlinear?
- Wrong model?

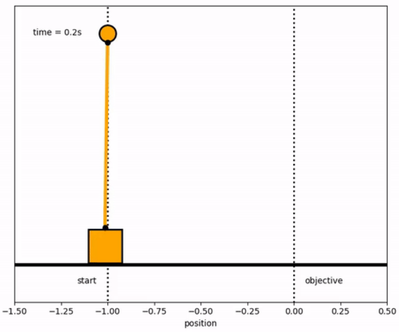

Today

What if?

- Nonlinear?
- Wrong model?

Today

What if?

- Nonlinear?
- Wrong model?
- Learning?
- Safety?

Today

What if?

- Nonlinear?
- Wrong model?
- Learning?
- Safety?

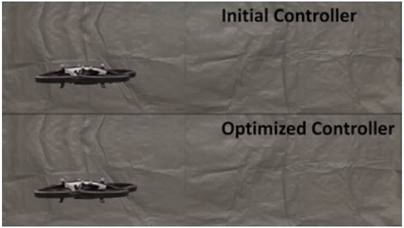

Today

What if?

- Nonlinear?
- Wrong model?
- Learning?
- Safety?

Today

LQR

$\downarrow$

Model Predictive Control

$\downarrow$

$\downarrow$

$\downarrow$

$\downarrow$

safe RL

Today

LQR

$\downarrow$

Model Predictive Control

$\downarrow$ Nonlinear

Learning $\downarrow$

$\downarrow$ Explore/Exploit

Constraints $\downarrow$

safe RL

LQR vs MPC

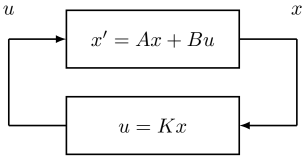

Find K a priori

LQR vs MPC

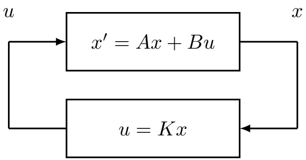

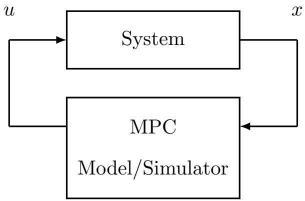

Find K a priori

Find u at every step

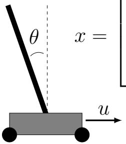

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

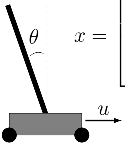

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

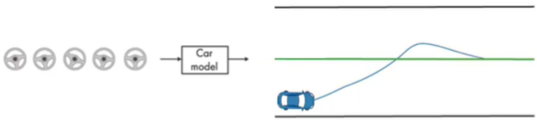

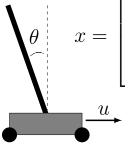

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

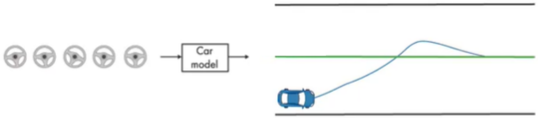

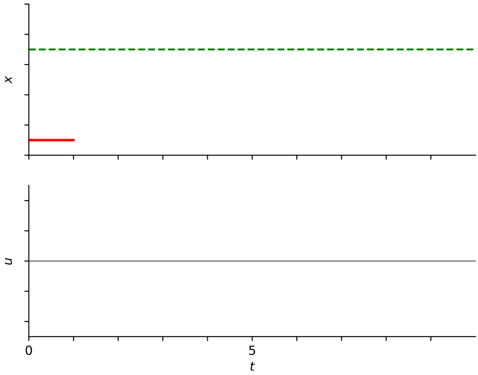

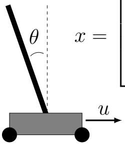

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

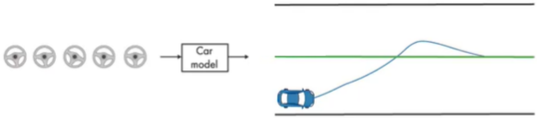

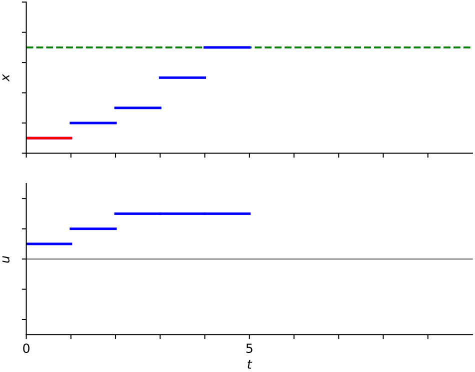

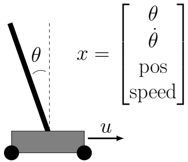
θ
u
x = [θ, θ̇, pos, speed]

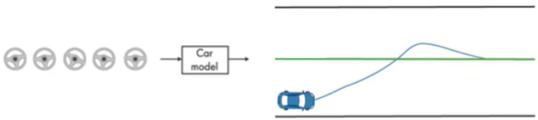
Car model

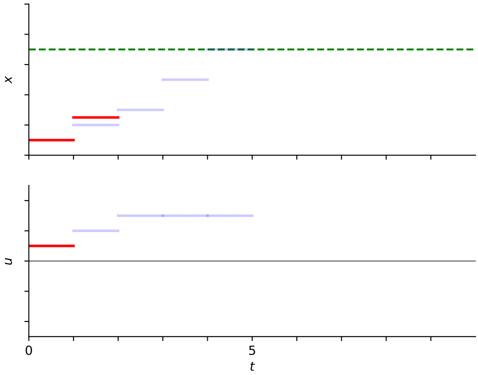
x
u
0
5
t

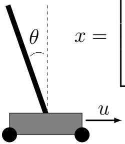

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

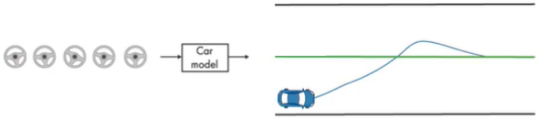

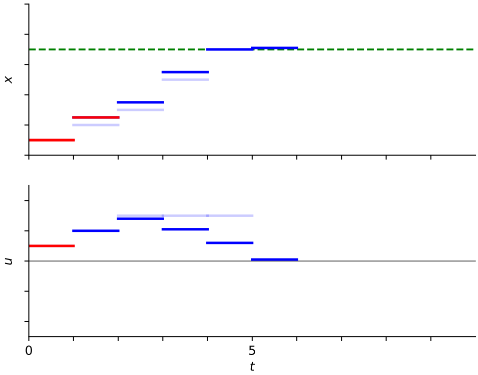

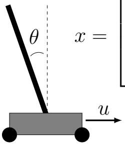

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

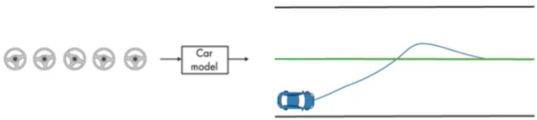

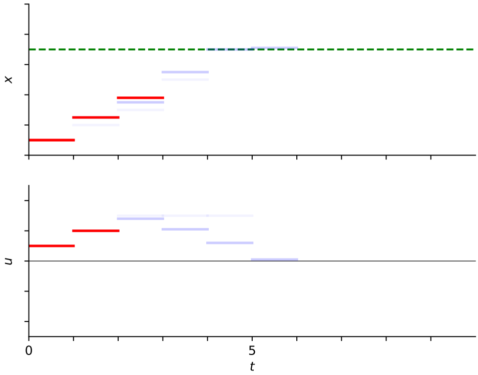

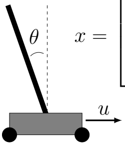

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

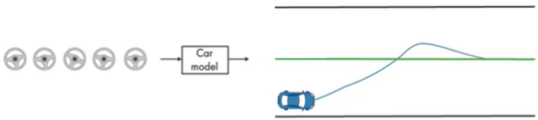

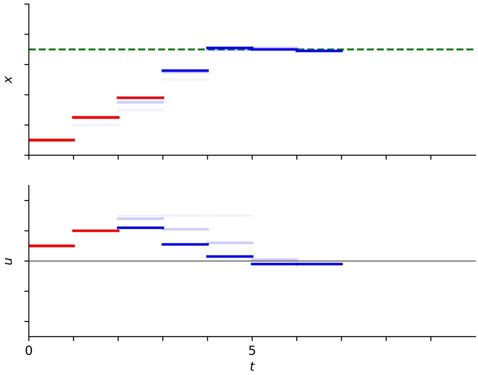

LQR vs MPC

Off/Online　　Full/Local model　　Non-linear

LQR vs MPC

Off/Online Full/Local model Non-linear

LQR vs MPC

Off/Online Full/Local model Non-linear

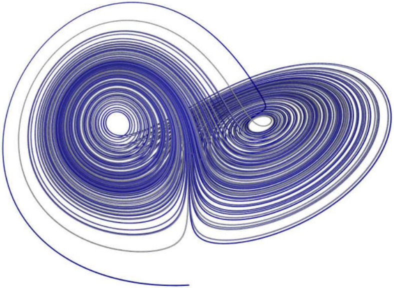

Example

$$x_{i+1} = A(x_t)x_i + B(x_t)u_i$$

Example

$$x_{i+1} = A(x_t)x_i + B(x_t)u_i$$

$$\begin{aligned} \min_{u_{t+1},\dots,u_{t+H}} \quad & \sum_{i=1}^{H} \dots \\ s.t. \quad & x_{i+1} = A(x_t)x_i + B(x_t)u_i \end{aligned}$$

Example

$$x_{i+1} = A(x_t)x_i + B(x_t)u_i$$

$$\begin{aligned} \min_{u_{t+1},\dots,u_{t+H}} \quad & \sum_{i=1}^{H} x_{t+i}^\top Q x_{t+i} + u_{t+i}^\top R u_{t+i} \\ s.t. \quad & x_{i+1} = A(x_t)x_i + B(x_t)u_i \end{aligned}$$

Example

$$x_{i+1} = A(x_t)x_i + B(x_t)u_i$$

$$\begin{aligned} \min_{u_{t+1},\dots,u_{t+H}} \quad & \sum_{i=1}^{H} x_{t+i}^\top Q x_{t+i} + u_{t+i}^\top R u_{t+i} \\ s.t. \quad & x_{i+1} = A(x_t)x_i + B(x_t)u_i \end{aligned}$$

or $\quad u_t = K x_t$

or $\quad u_t, \dots, u_{t+H} = NN(x_t, \theta) \quad$ more flexible and hungry

MPC

Replace full planning by local optimization

Learning

"Online LQR"

Learning

"Online LQR"

True dynamics: $x' = Ax + Bu$

Model of the system: $x' = A_t x + B_t u$

Learning

"Online LQR"

True dynamics: $x' = Ax + Bu$

Model of the system: $x' = A_t x + B_t u$

Start with (A_0, B_0), update as you go

Learning safely?

Safety   Correct models   Explore/Exploit   Uncertainty

Safety

Safety

Safety

Ensure we don't crash

avoid known bad outcomes / stay within known safe regime

A safe starting point

Uncertainty about the dynamics

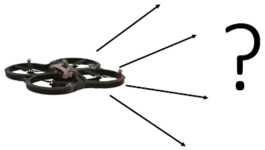

A safe starting point

Uncertainty about the dynamics

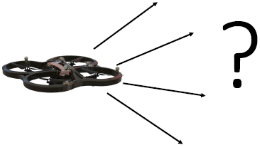

There is a safe controller that can take over if $x \in \mathcal{X}_{\text{safe}}$

A safe starting point

Uncertainty about the dynamics

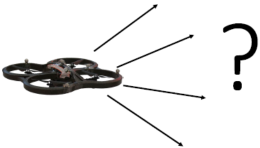

There is a safe controller that can take over if $x \in \mathcal{X}_{\text{safe}}$

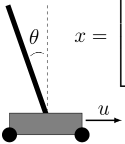

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \\ \text{pos} \\ \text{speed} \end{bmatrix}$$

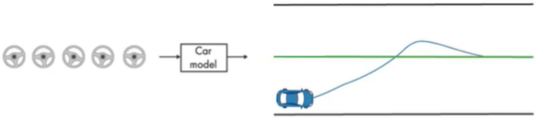

A safe starting point

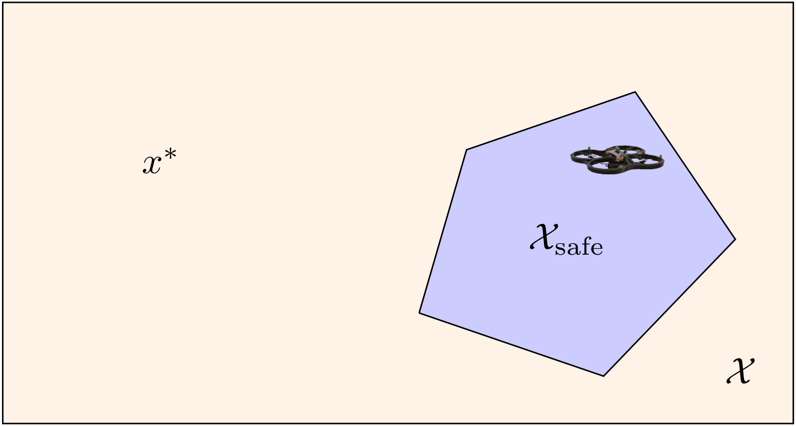

How do we get to x^*?

Safety objectives

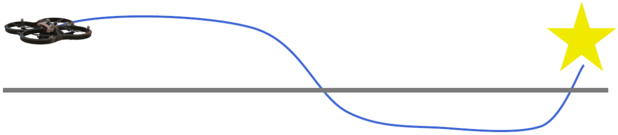

Safety objectives

$$\text{Loss} = J$$

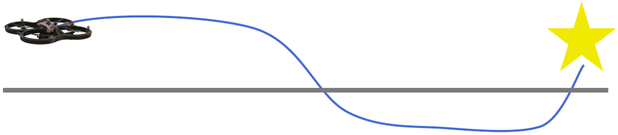

Safety objectives

$$\text{Loss} = J$$

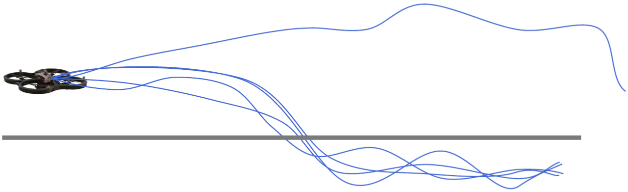

Safety objectives

$$\text{Loss} = \mathbb{E}[J]$$

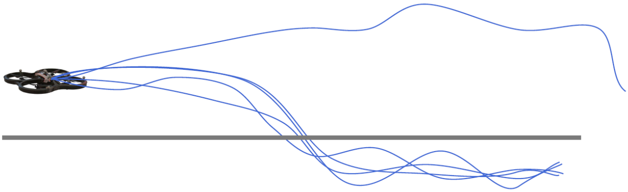

Safety objectives

$$\text{Loss} = \mathbb{E}[J]$$

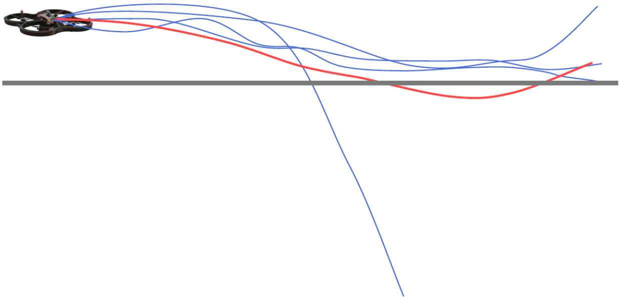

Safety objectives

$$\text{Loss} = \mathbb{E}[J] + \mathbb{E}[(J - \mathbb{E}[J])^2]$$

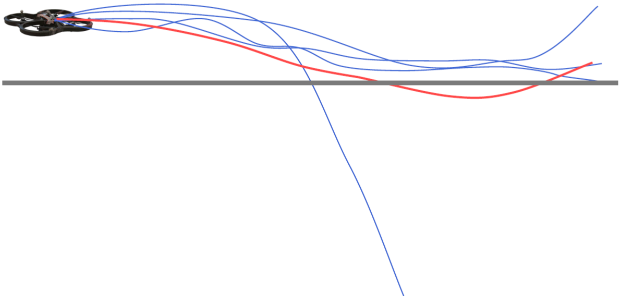

Safety objectives

$\text{Loss} = \mathbb{E}[J]$, but with safety constraints

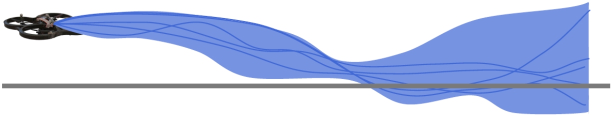

Safety constraints

Each step is optimization $\rightarrow$ constrained opt.

Safety constraints

Each step is optimization $\rightarrow$ constrained opt.

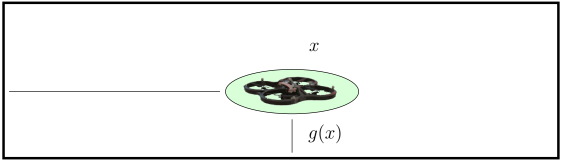

Safety constraints

Each step is optimization $\rightarrow$ constrained opt.

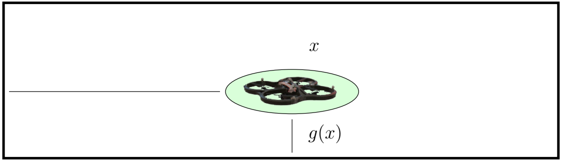

$g(x) \geq 0$

Safety constraints

Each step is optimization $\rightarrow$ constrained opt.

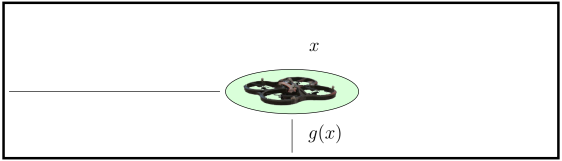

$g(x) \geq 0$ $\theta \rightarrow u \rightarrow x$ $g(\theta) \geq 0$

Safety constraints

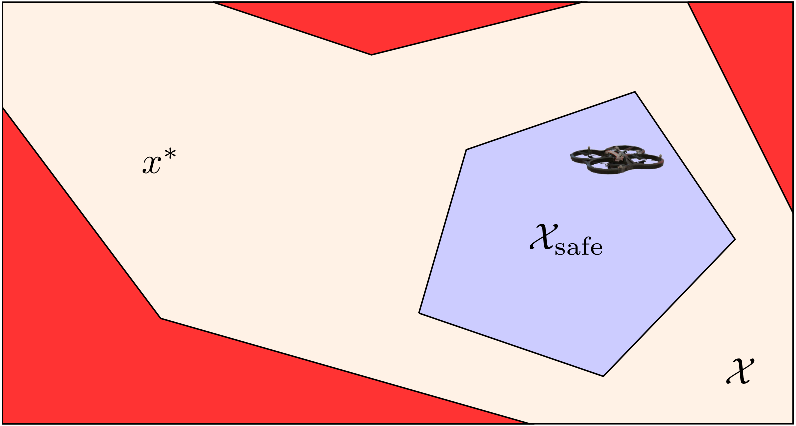

How do we get to x^*?

Uncertain transition dynamics

$$x' = f(x, u)$$
$$f(x, u) = h(x, u) + e(x, u)$$

h is known

e is not, but we have a probabilistic model

Uncertain transition dynamics

$$x' = f(x, u)$$
$$f(x, u) = h(x, u) + e(x, u)$$

h is known

e is not, but we have a probabilistic model

For (μ, σ) and β, with high probability

$$|\mu(x, u) - e(x, u)| \leq \beta\sigma(x, u)$$

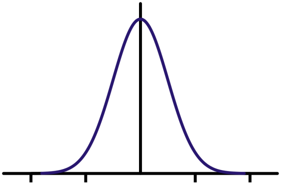

Optimization with Uncertainty

Gaussian Processes Crash Course

Kernel $k(x, y)$

Optimization with Uncertainty

Gaussian Processes Crash Course

Kernel $k(x, y)$

Squared exponential

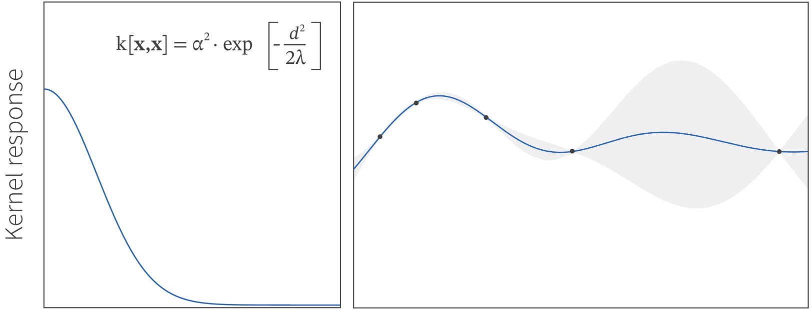

Optimization with Uncertainty

Gaussian Processes Crash Course

Kernel $k(x, y)$

Periodic

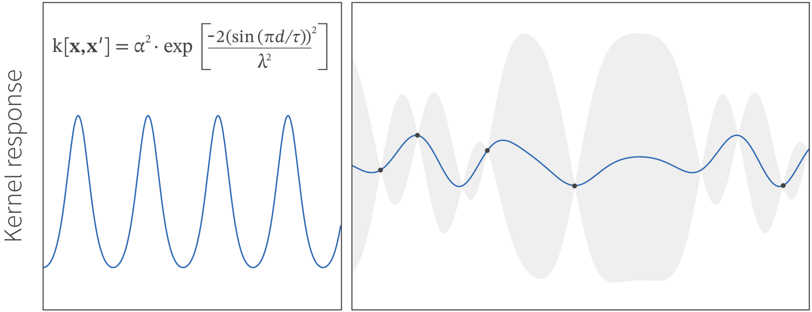

Optimization with Uncertainty

Gaussian Processes Crash Course

Kernel $k(x, y)$

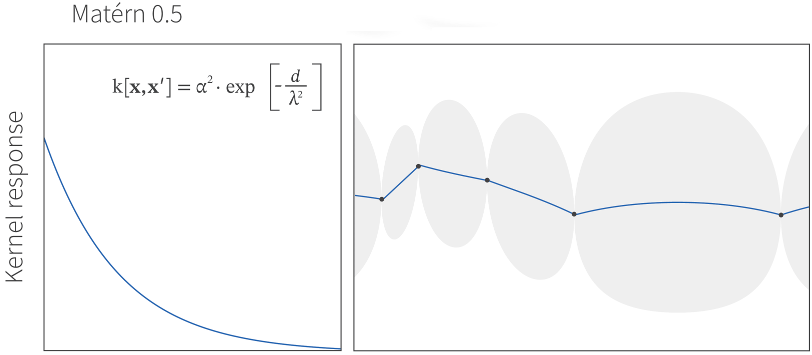

Optimization with Uncertainty

Gaussian Processes Crash Course

Kernel $k(x, y)$

Linear regression: $\|Xw - y\|^2$

Kernel regression: $\|Kw - y\|^2$ where $K_{ij} = k(x_i, x_j)$

$+$ Covariances

Exploration/Exploitation

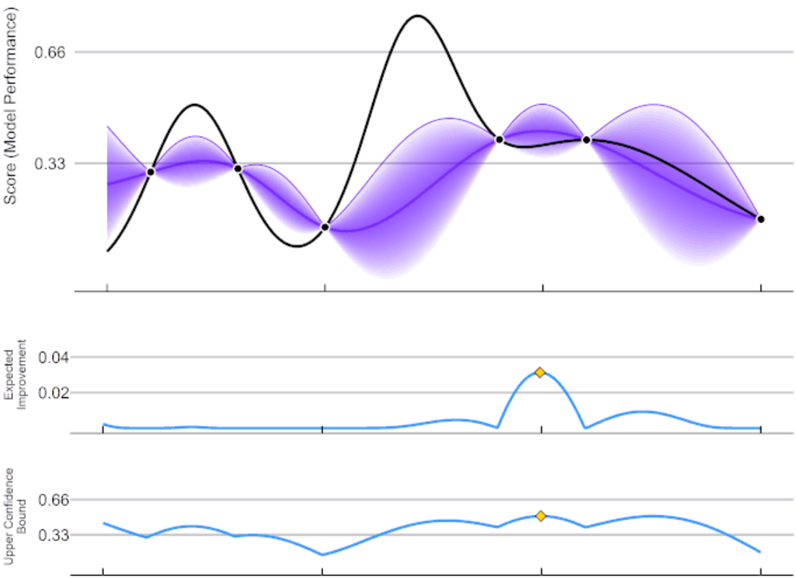

Exploration/Exploitation

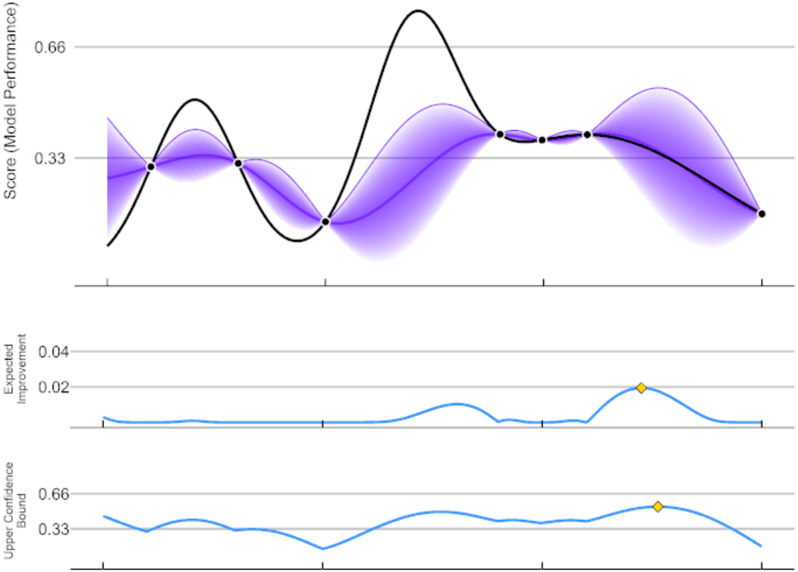

Exploration/Exploitation

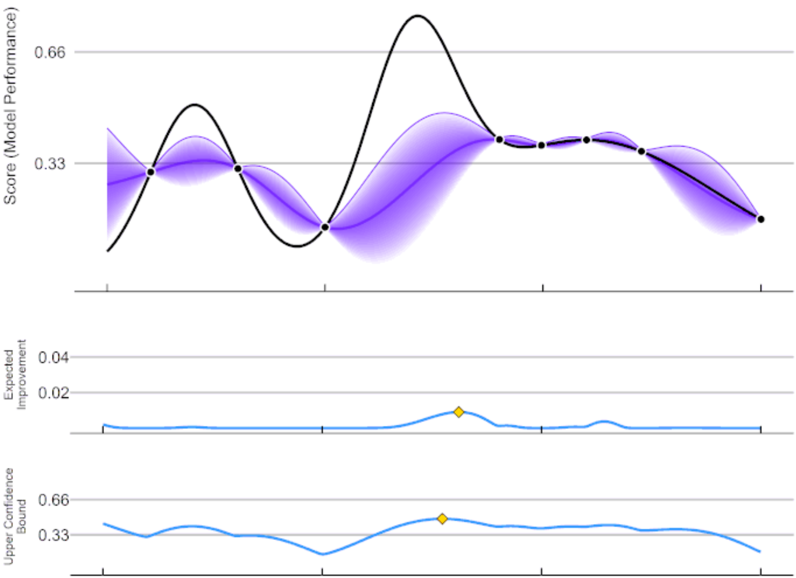

Exploration/Exploitation

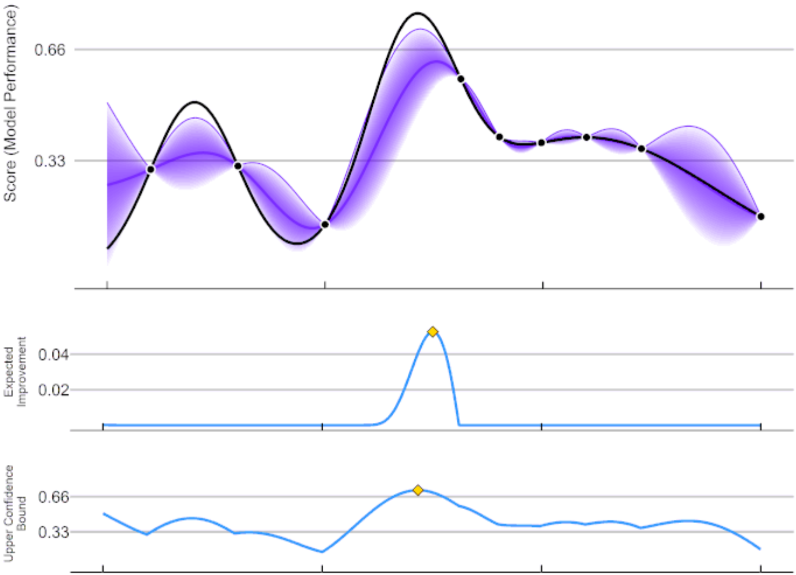

Exploration/Exploitation

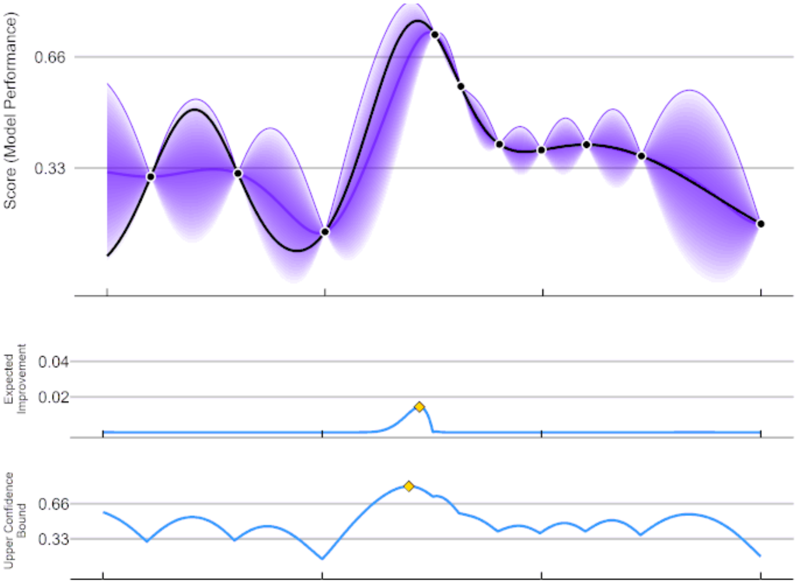

Exploration/Exploitation

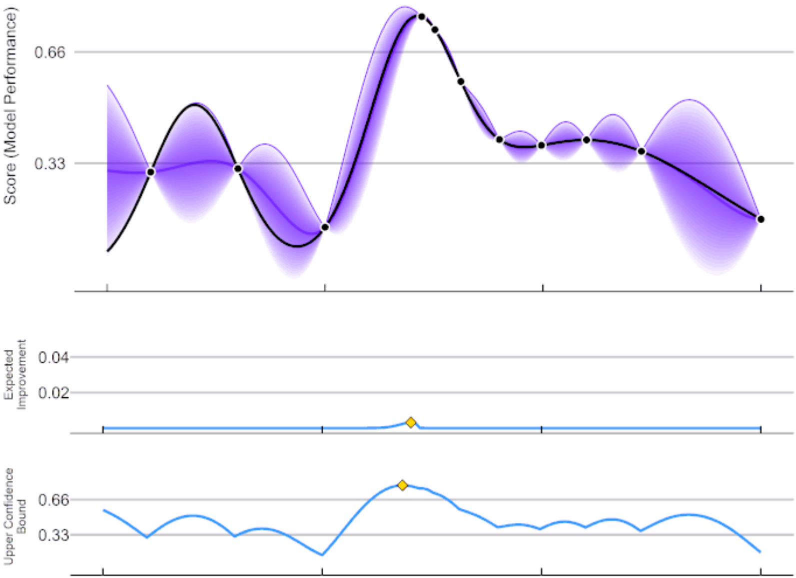

Exploration/Exploitation

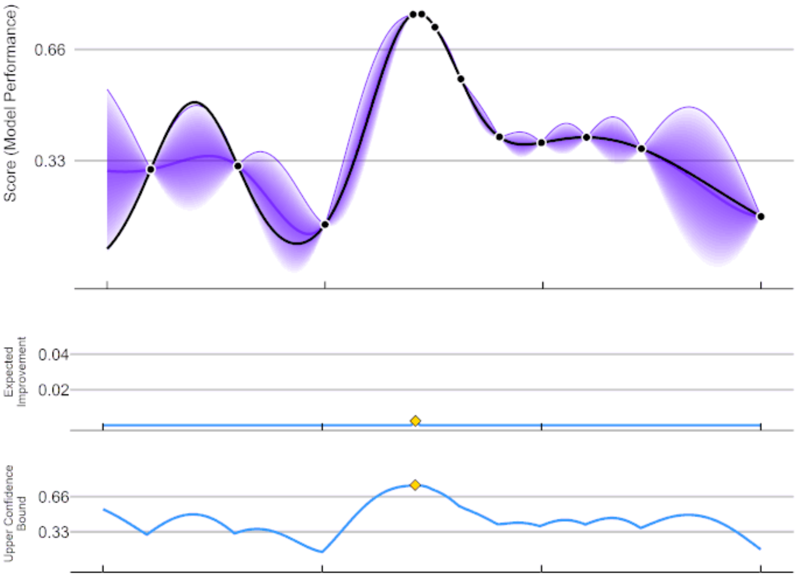

Safe Explore/Exploit

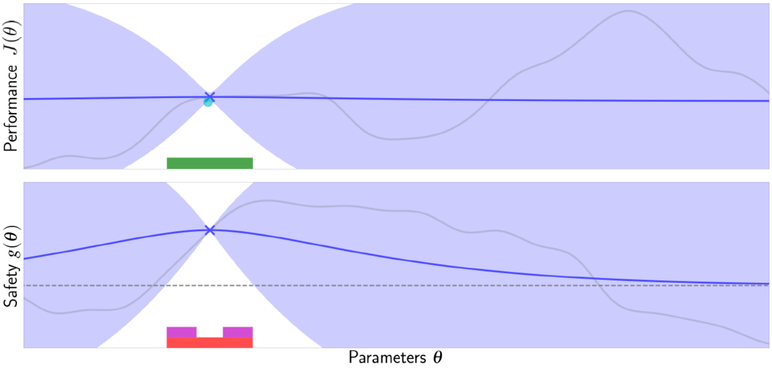

Safe Explore/Exploit

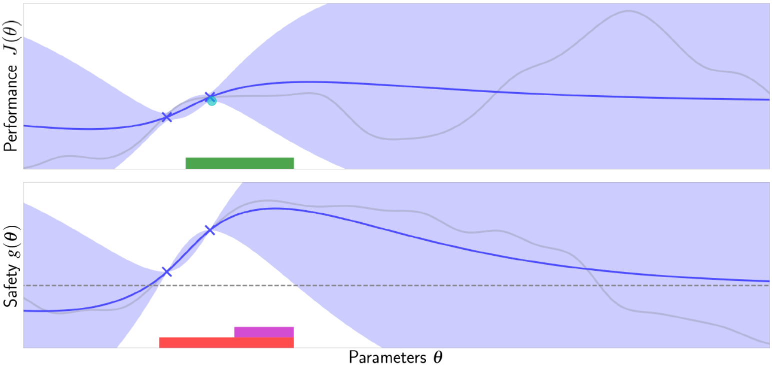

Safe Explore/Exploit

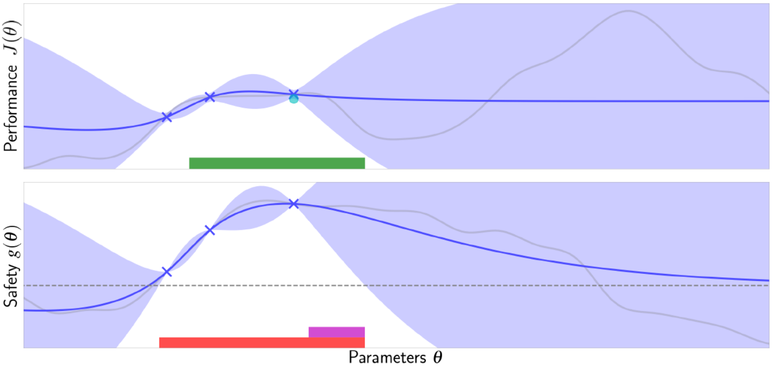

Safe Explore/Exploit

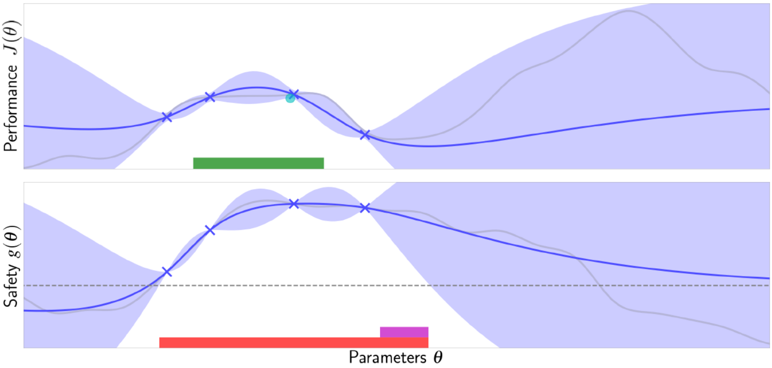

Safe Explore/Exploit

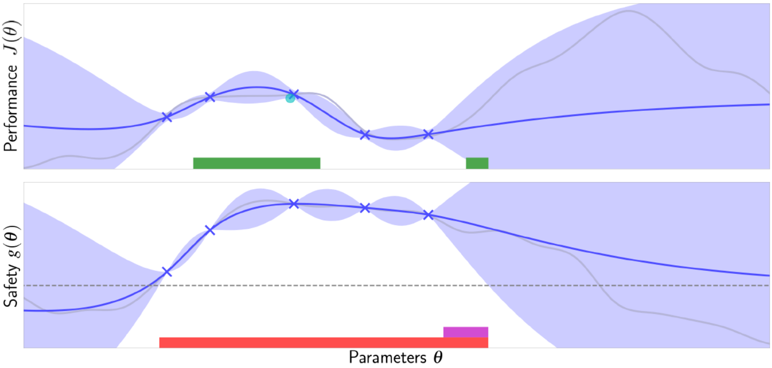

Safe Explore/Exploit

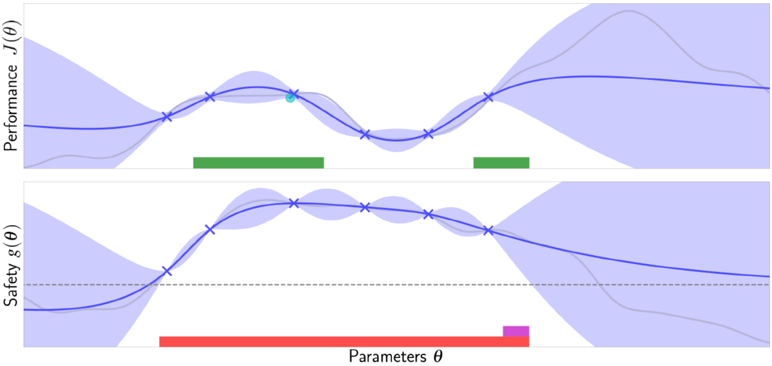

Safe Explore/Exploit

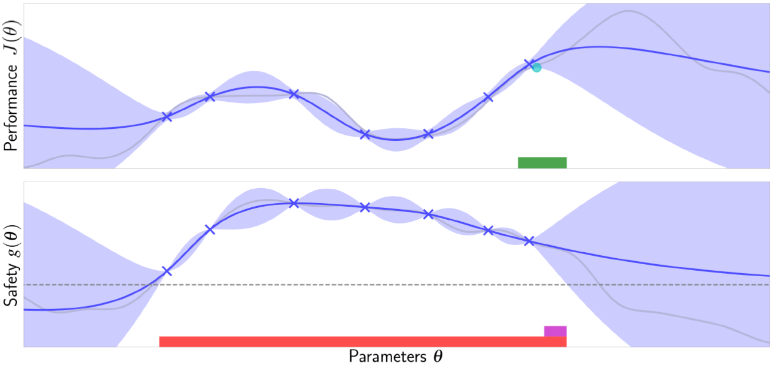

Safe Explore/Exploit

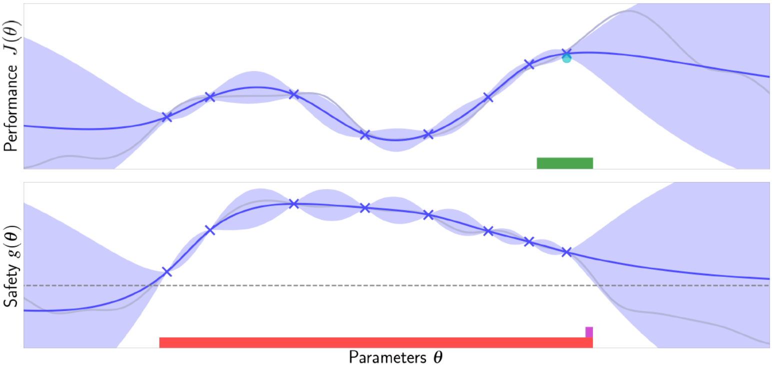

Safe Explore/Exploit

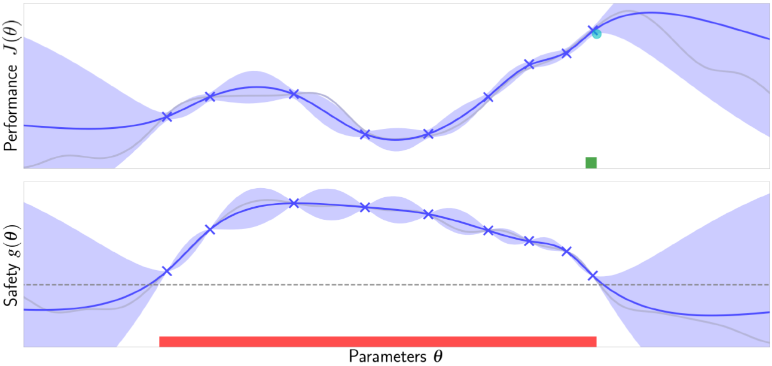

Safe Explore/Exploit

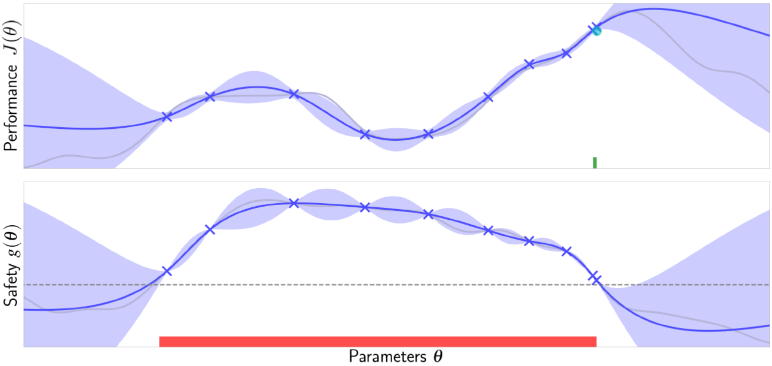

Safe Explore/Exploit

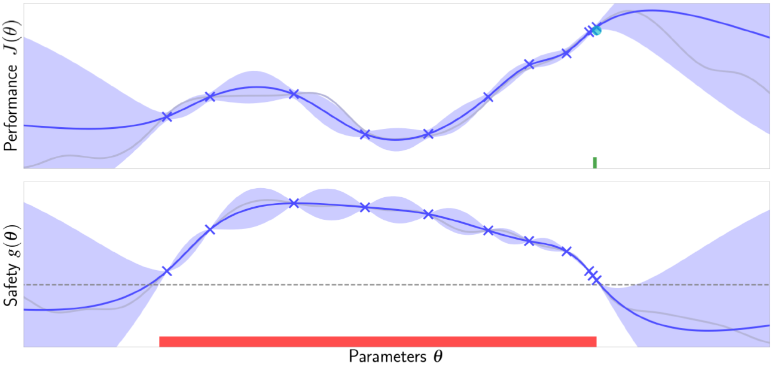

Ideal result

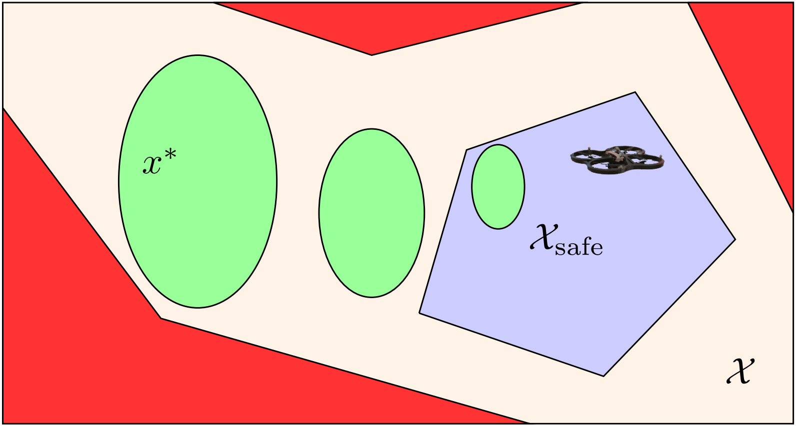

Slight issue: propagating uncertainty

$g(x)$ might be linear/convex but $g(\theta)$ is not

Slight issue: propagating uncertainty

$g(x)$ might be linear/convex but $g(\theta)$ is not

$$g(f(x, u = Model(x, \theta)))$$

Slight issue: propagating uncertainty

$g(x)$ might be linear/convex but $g(\theta)$ is not

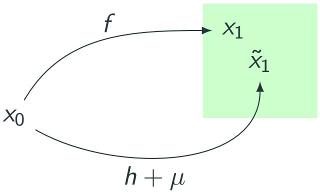

$$|\mu(x, u) - e(x, u)| \leq \beta\sigma(x, u)$$

Slight issue: propagating uncertainty

$g(x)$ might be linear/convex but $g(\theta)$ is not

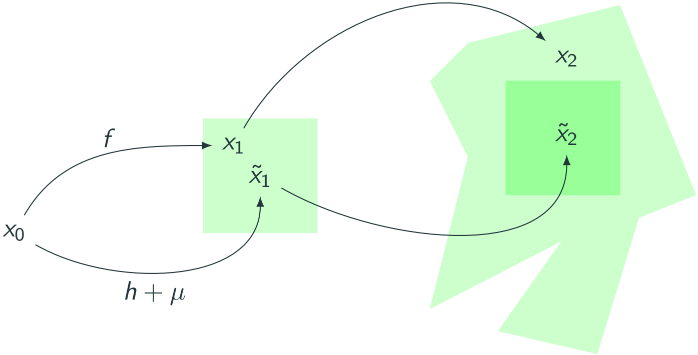

$$|\mu(x, u) - e(x, u)| \leq \beta\sigma(x, u)$$

Linearized uncertainty propagation

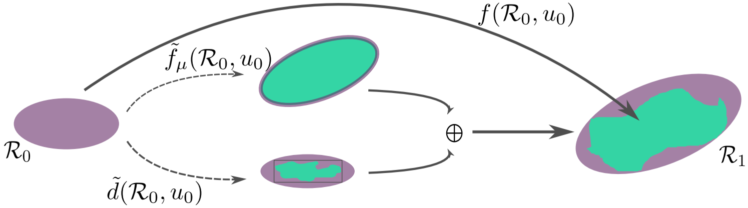

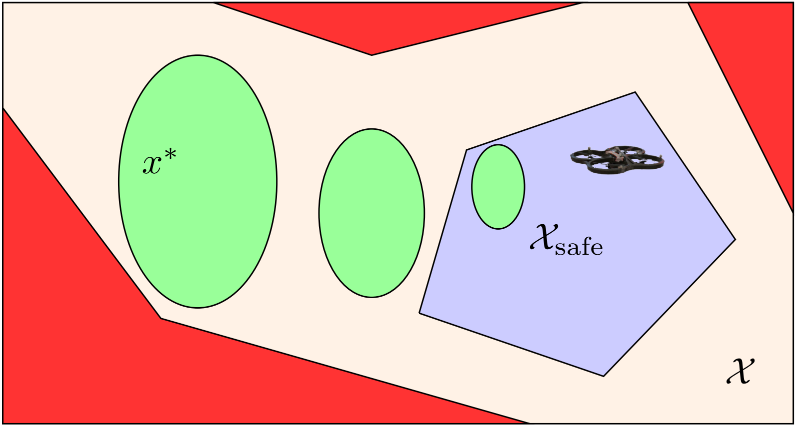
x^*
$\mathcal{X}_{\text{safe}}$
$\mathcal{X}$

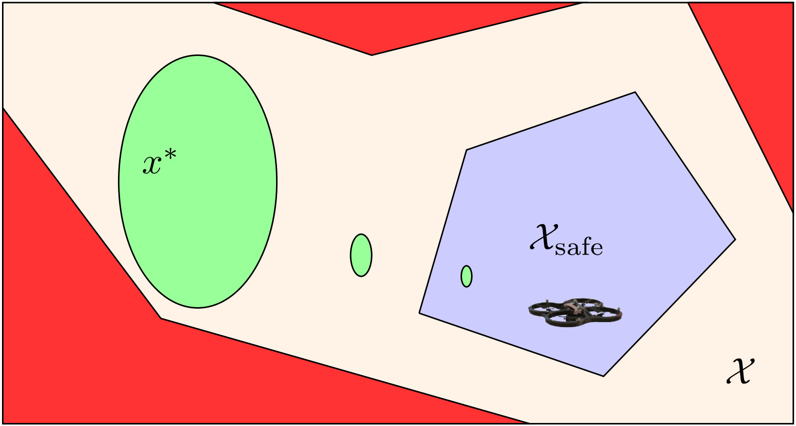
x^*
$\mathcal{X}_{\text{safe}}$
$\mathcal{X}$

One last thing...

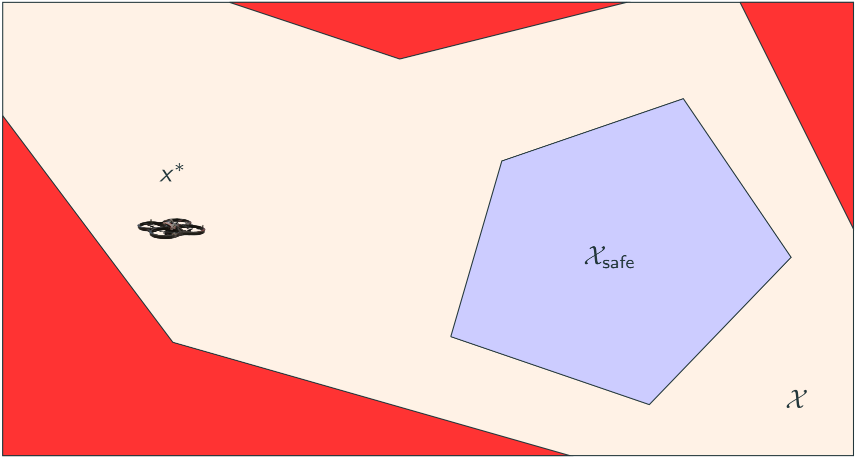

One last thing...

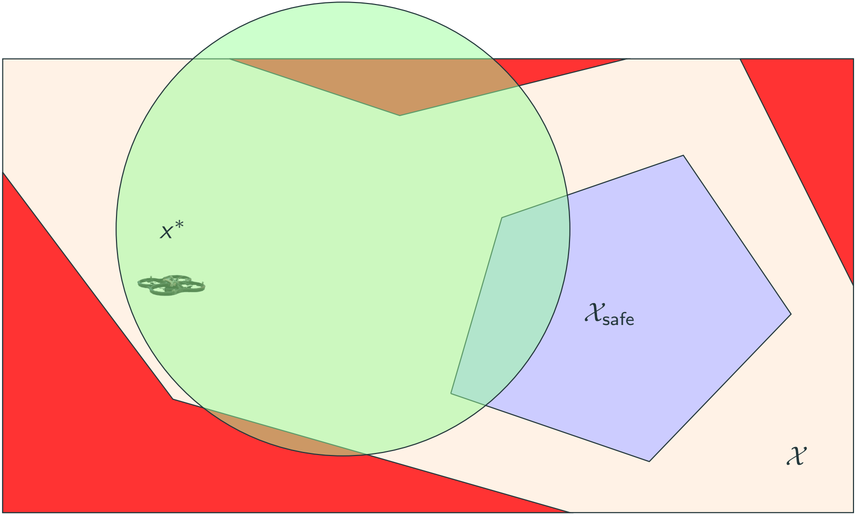

Simultaneous planning

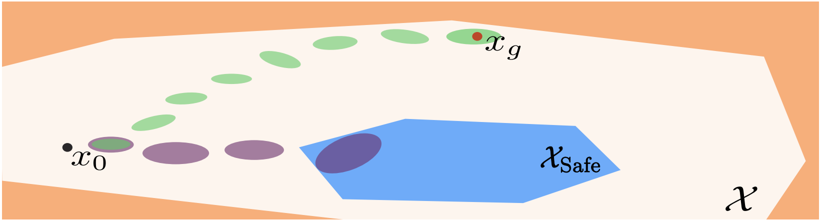

Simultaneous planning

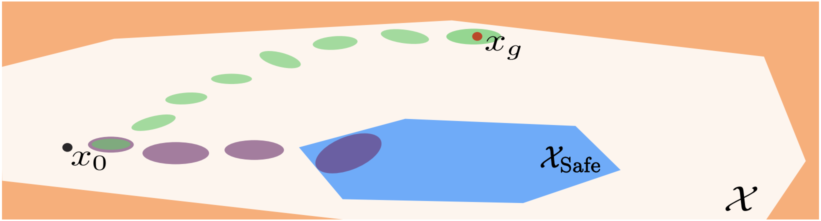

- A safe default controller
- A definition of the boundaries
- A well specified Gaussian Process
- The Lipschitz constant of the model error
- Bayesian Optimization Model Predictive Control