Automatic Gain Tuning

based on

Gaussian Process Global Optimization

(= Bayesian Optimization)

https://is.tuebingen.mpg.de/publications/marco_icra_2016

A simple PD control example

Global optimal gains, θ to get a minimum cost J ?

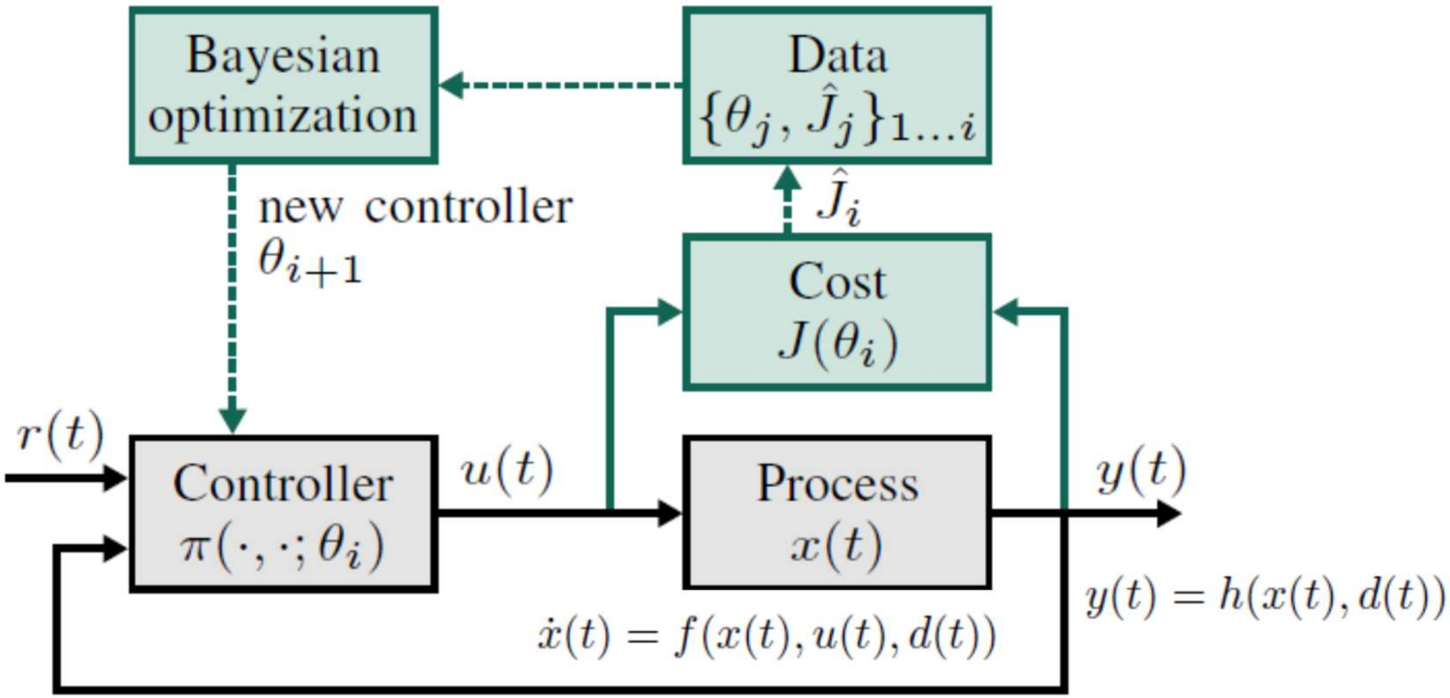

$$u(t) = \pi(y(t), r(t); \theta)$$

$$= \theta_1\big(r(t) - y(t)\big) + \theta_2 \dot{y}(t)$$

$$J = \int_0^T \|y(t) - r(t)\|^2 + \|u(t)\|^2 \, dt$$

A simple PD control example

Procedure of Bayesian Optimization

1. GP prior before observing any data

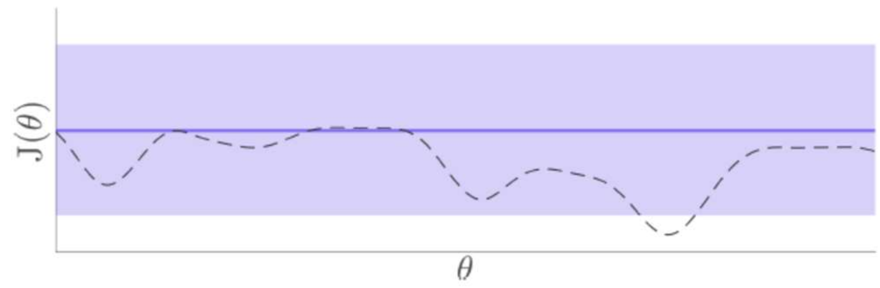

2. GP posterior, after five noisy evaluations

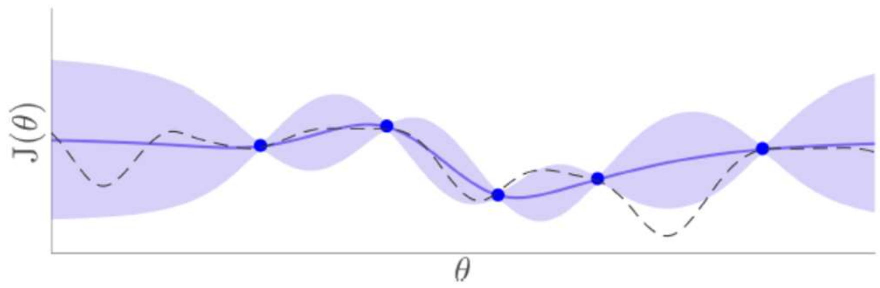

3. The next parameters θ_{next} are chosen at the maximum of Acquisition function

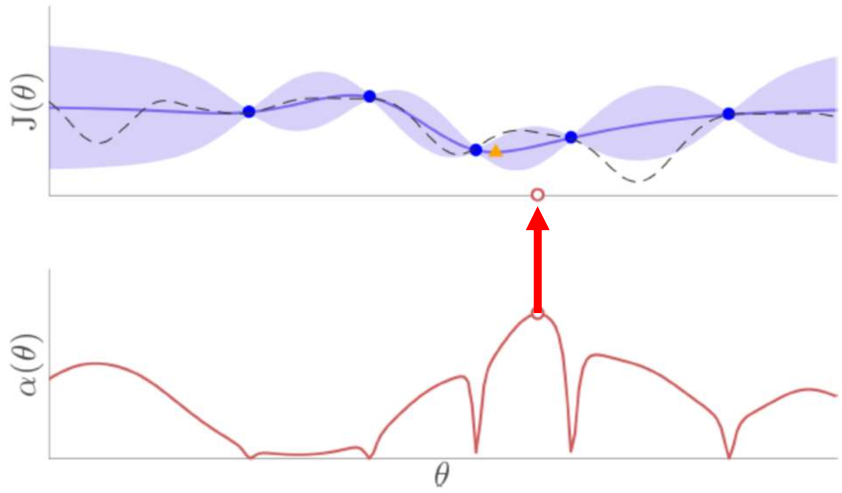

Repeat until you can find a globally optimal θ

Acquisition function

Idea: build a probabilistic model of the function f

LOOP

- choose new query point(s) to evaluate
  decision criterion: acquisition function $\alpha_t(\cdot)$
- update model

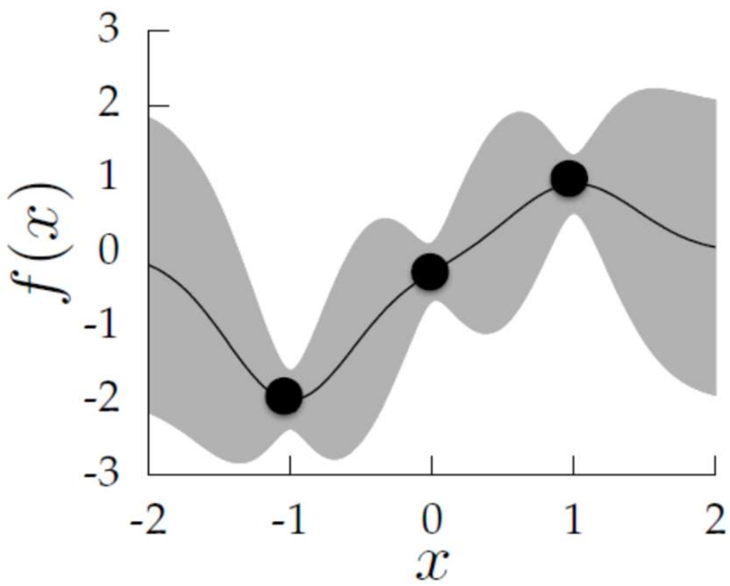

$$x_* = \underset{\mathcal{X} \subset \mathbb{R}^d}{\operatorname{argmin}} \; f(x)$$

$$x_t = \underset{\mathcal{X} \subset \mathbb{R}^d}{\operatorname{argmax}} \; \alpha_t(x)$$

$$t = 1, \cdots, T$$

Acquisition function and Entropy Search

$$\underset{x_t \in \mathfrak{X}}{\text{maximize}}\ \alpha_t(x_t) \qquad t = 1, \cdots, T$$

$$I(a;b) = H(a) - H(a|b)$$
$$= H(b) - H(b|a)$$

Point to query

Location of global optimum

Observed Data

$$\alpha_t(x) = I(\{x, y\}; x_* \mid D_t)$$

$$= H(x_*|D_t) - H(x_*|D_t U(x,y))$$

Information gain, I: Mutual information for an observed data

➔ Reduction of uncertainty in the location x_* by selecting points (x, y) that are expected to cause the largest reduction in entropy of distribution $H(x_*|D_t)$

Control design problem

propose the use of **Entropy Search**, a recent algorithm for global Bayesian optimization, as the minimizer for the LQR tuning problem. ES employs a **Gaussian process** (GP) as a non-parametric model capturing the knowledge about the unknown cost function.

consider a system that follows a discrete-time nonlinear dynamic model

$$\boldsymbol{x}_{k+1} = \boldsymbol{f}(\boldsymbol{x}_k, \boldsymbol{u}_k, \boldsymbol{w}_k)$$

system states x_k, control input u_k, and zero-mean process noise w_k at time instant k

Control design problem

common way to measure the performance of a control system is through a quadratic cost function such as

$$J = \lim_{K \to \infty} \frac{1}{K} \mathbb{E}\left[\sum_{k=0}^{K} \boldsymbol{x}_k^T \boldsymbol{Q} \boldsymbol{x}_k + \boldsymbol{u}_k^T \boldsymbol{R} \boldsymbol{u}_k \right]$$

cost captures a trade-off between control performance (keeping x_k small) and control effort(keeping u_k small)

Nonlinear control design problem is intractable in general.

linear model is often sufficient for control design

$$\tilde{\boldsymbol{x}}_{k+1} = \boldsymbol{A}_{\mathrm{n}} \tilde{\boldsymbol{x}}_k + \boldsymbol{B}_{\mathrm{n}} \boldsymbol{u}_k + \boldsymbol{w}_k$$

LQR tuning problem

Linear Quadratic Regulator (LQR)

$$\boldsymbol{u}_k = \boldsymbol{F}\boldsymbol{x}_k$$

static gain matrix F can readily be computed by solving the discrete-time infinite-horizon LQR problem for the nominal model (A_n, B_n) and the weights (Q, R).

$$\boldsymbol{F} = \mathrm{lqr}(\boldsymbol{A}_\mathrm{n}, \boldsymbol{B}_\mathrm{n}, \boldsymbol{Q}, \boldsymbol{R}).$$

goal of the automatic LQR tuning is to vary the parameters θ such as to minimize the cost

$$J = J(\boldsymbol{\theta}).$$

Optimization problem

$$\arg\min J(\boldsymbol{\theta}) \quad \text{s.t. } \boldsymbol{\theta} \in \mathcal{D}$$

Consider the approximate one of quadratic cost function

$$\hat{J} = \frac{1}{K}\left[\sum_{k=0}^{K} \boldsymbol{x}_k^T \boldsymbol{Q} \boldsymbol{x}_k + \boldsymbol{u}_k^T \boldsymbol{R} \boldsymbol{u}_k\right]$$

LQR TUNING WITH ENTROPY SEARCH

uncertainty over the objective function J is represented by a probability measure p(J), typically a Gaussian process (GP)

prior knowledge about the cost function J as the GP

$$J(\boldsymbol{\theta}) \sim \mathcal{GP}\left(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}_*)\right)$$

with mean function $\mu(\theta)$and covariance function $k(\theta, \theta_*)$

squared exponential (SE) covariance function

$$k_{\mathrm{SE}}(\boldsymbol{\theta}, \boldsymbol{\theta}_*) = \sigma^2 \exp\left[-\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}_*)^{\mathrm{T}} \boldsymbol{S}(\boldsymbol{\theta} - \boldsymbol{\theta}_*)\right]$$

LQR TUNING WITH ENTROPY SEARCH

noisy evaluations of cost function can be modeled as

$$\hat{J} = J(\boldsymbol{\theta}) + \varepsilon$$

Conditioning the GP on the data $\{y, \Theta\}$ then yields another GP with posterior mean $\bar{\mu}(\theta)$ and a posterior variance $\bar{k}(\theta, \theta_*)$.

➔ shape of the mean is adjusted to fit the data points, and the uncertainty (standard deviation) is reduced around the evaluations points.

LQR TUNING WITH ENTROPY SEARCH

ES is one out of several popular formulations of Bayesian Optimization

Aims to **reduce the uncertainty in the location** $\theta_* \; by \; selecting \; points$ $that \; are$ expected to cause **the largest reduction in entropy**(➔ uncertainty) of distribution, J(θ)

Bayesian

Optimization

Bayesian Optimization and How to Scale It Up

Zi Wang
CS Colloquium, US

Blackbox Function Optimization

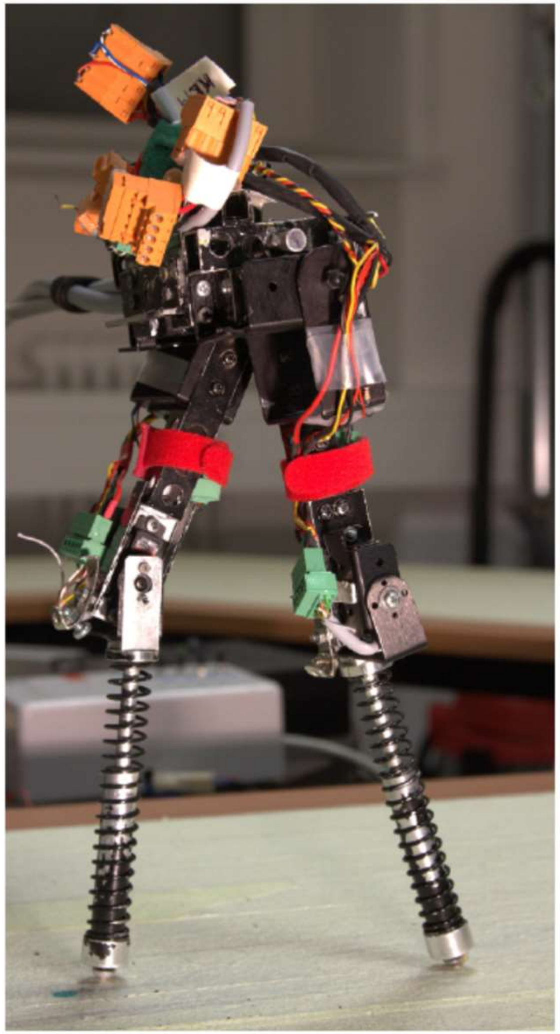

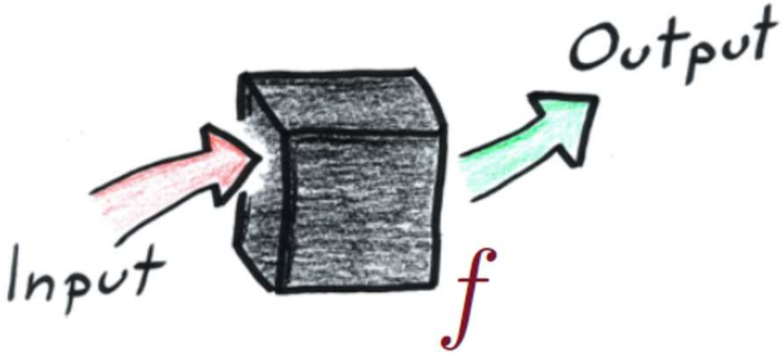

Goal:

$$x_* = \underset{\mathcal{X} \subset \mathbb{R}^d}{\operatorname{argmax}} \; f(x)$$

(Calandra et al., 2015)

Bayesian Optimization

Idea: build a probabilistic model of the function f

LOOP

- choose new query point(s) to evaluate
  decision criterion: acquisition function $\alpha_t(\cdot)$
- update model

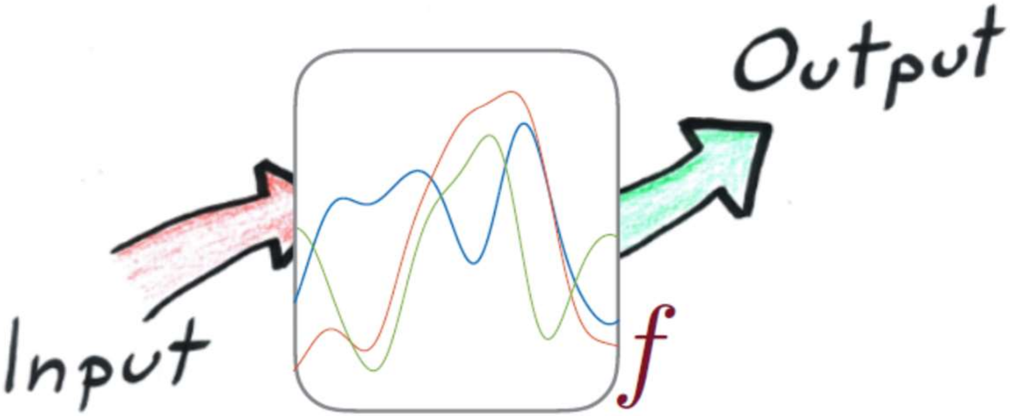

$$x_* = \underset{\mathcal{X} \subset \mathbb{R}^d}{\operatorname{argmax}} f(x)$$

$$x_t = \underset{\mathcal{X} \subset \mathbb{R}^d}{\operatorname{argmax}} \alpha_t(x)$$

$$t = 1, \cdots, T$$

Gaussian Processes (GPs)

- probability distribution over functions
- any finite set of function values is a multi-variate Gaussian

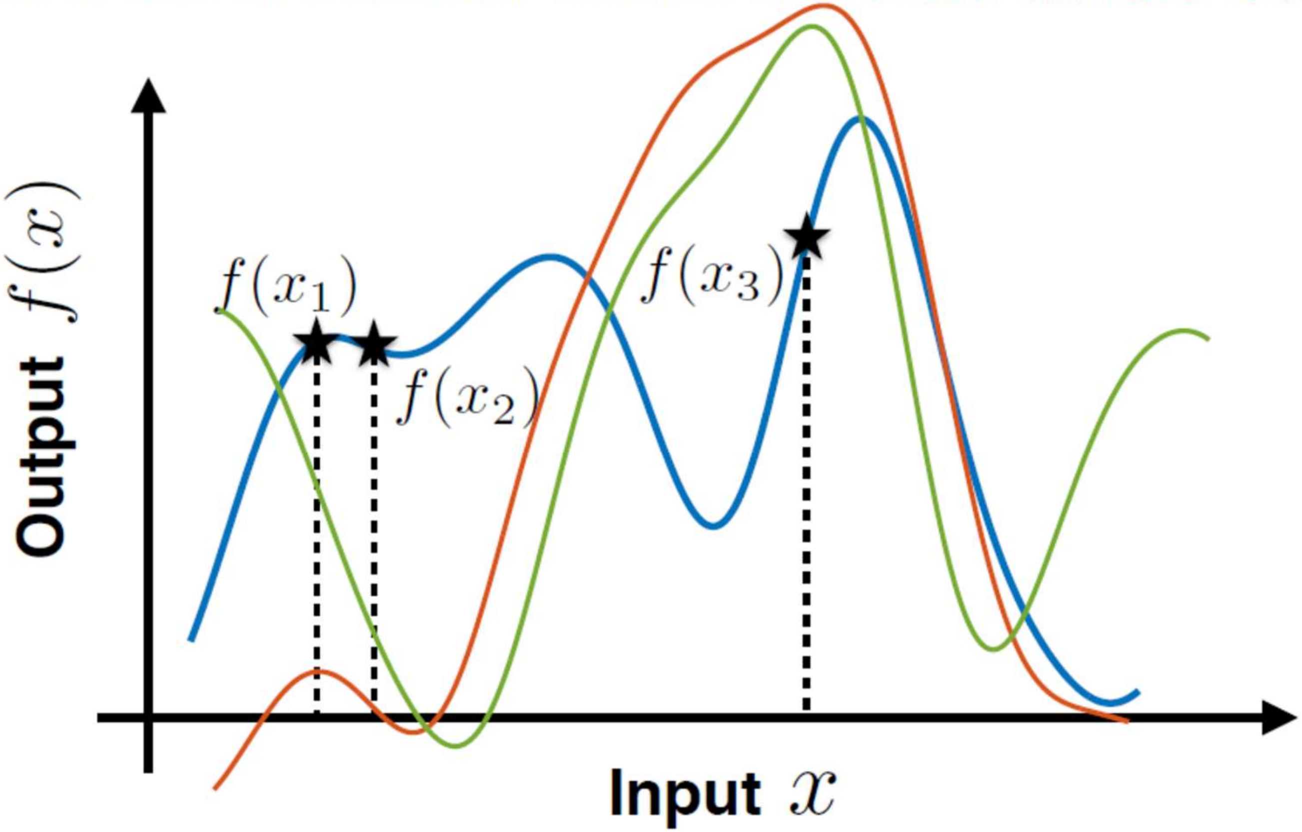

Gaussian Processes (GPs)

- probability distribution over functions
- any finite set of function values is a multi-variate Gaussian
- kernel function $k(\cdot,\cdot)$; mean function $\mu(\cdot)$

$$\begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu(x_1) \\ \vdots \\ \mu(x_n) \end{bmatrix}, \begin{bmatrix} k(x_1,x_1), & \cdots, & k(x_1,x_n) \\ \vdots, & & \vdots \\ k(x_n,x_1), & \cdots, & k(x_n,x_n) \end{bmatrix} \right)$$

- function $f \sim GP(\mu, k)$; observe noisy output at x_τ

$$y_\tau = f(x_\tau) + \epsilon,\ \epsilon \sim \mathcal{N}(0, \sigma^2)$$

Gaussian Processes (GPs)

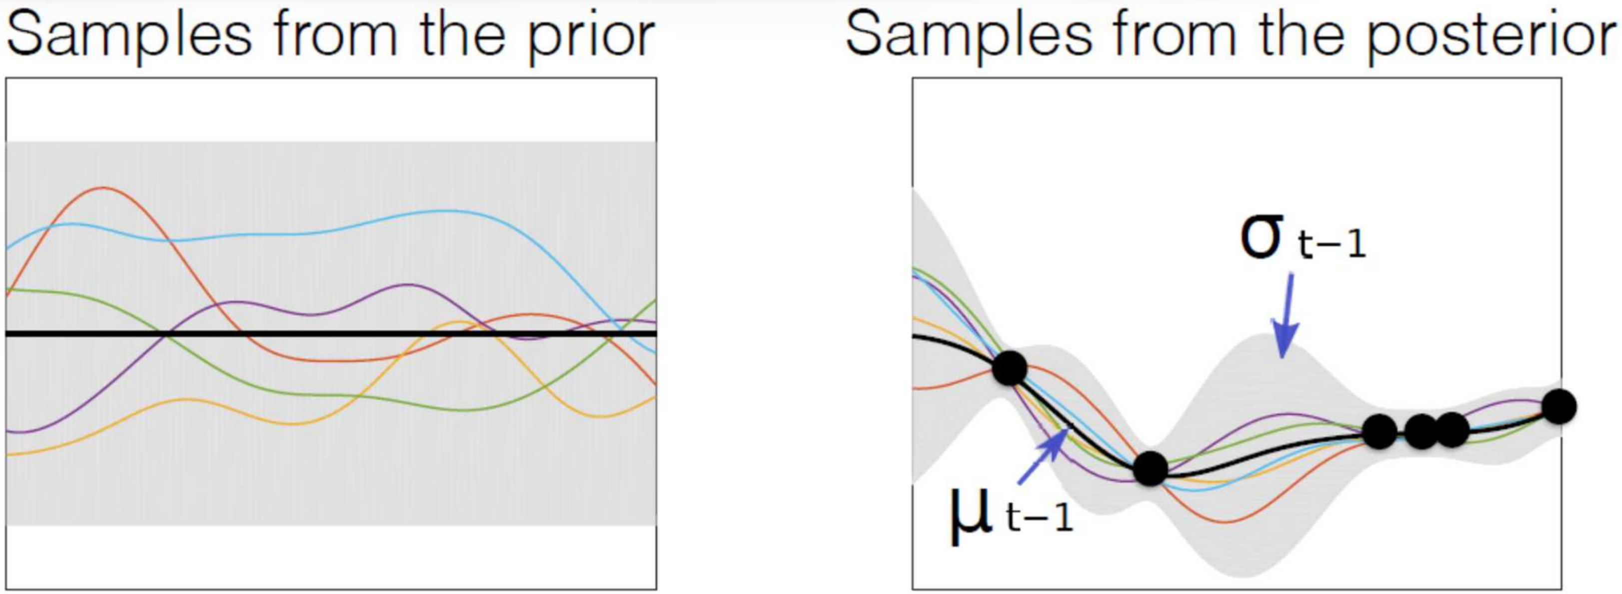

Given observations $D_t = \{(x_\tau, y_\tau)\}_{\tau=1}^{t-1}$, predict posterior mean and variance in closed form via conditional Gaussian

$$\mu_{t-1}(x) = k_{t-1}(x)^{\mathrm{T}}(K_{t-1} + \sigma^2 I)^{-1} y_{t-1}$$

$$\sigma_{t-1}(x)^2 = k(x,x) - k_{t-1}(x)^{\mathrm{T}}(K_{t-1} + \sigma^2 I)^{-1} k_{t-1}(x)$$

Bayesian Optimization

Idea: build a probabilistic model of the function f

LOOP

- choose new query point(s) to evaluate
  decision criterion: acquisition function $\alpha_t(\cdot)$
- update model

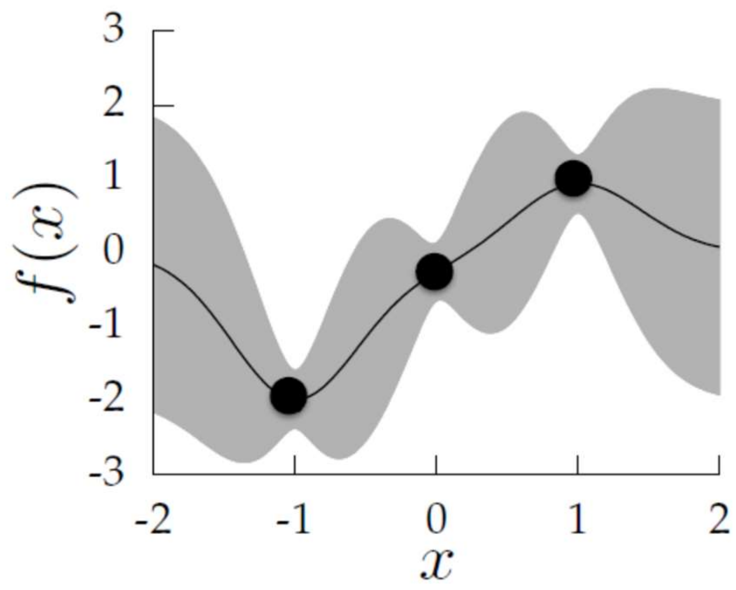

$$x_* = \underset{\mathcal{X} \subset \mathbb{R}^d}{\text{argmin}} \; f(x)$$

$$x_t = \underset{\mathcal{X} \subset \mathbb{R}^d}{\text{argmax}} \; \alpha_t(x)$$

$$t = 1, \cdots, T$$

Entropy Search and Predictive Entropy Search

$$\underset{x_t \in \mathfrak{X}}{\text{maximize}}\ \alpha_t(x_t) \qquad t = 1, \cdots, T$$

$$I(a;b) = H(a) - H(a|b)$$
$$= H(b) - H(b|a)$$

Point to query

Location of global optimum

Observed Data

$$\alpha_t(x) = I(\{x, y\}; x_* \mid D_t)$$

$$= H(x_*|D_t) - H(x_*|D_t U(x, y))$$

Information gain, I: Mutual information for an observed data

➔ Reduction of uncertainty in the location x_* by selecting points (x, y) that are expected to cause the largest reduction in entropy of distribution $H(x_*|D_t)$

Entropy Search and Predictive Entropy Search

$$p_{\min}(x) \equiv p[x = \arg\min f(x)]$$

Probability of Improvement

$$u_{\mathrm{PI}}(x) = p[f(x) < \eta] = \int_{-\infty}^{\eta} \mathcal{N}(f(x); \mu(x), \sigma(x)^2)\, \mathrm{d}f(x) = \Phi\left(\frac{\eta - \mu(x)}{\sigma(x)}\right)$$

$$\eta : \mathit{current\ best\ guess}$$

$$\Phi(z) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{z}{\sqrt{2}}\right)\right]$$

: Gaussian Cumulative Density Ft'n

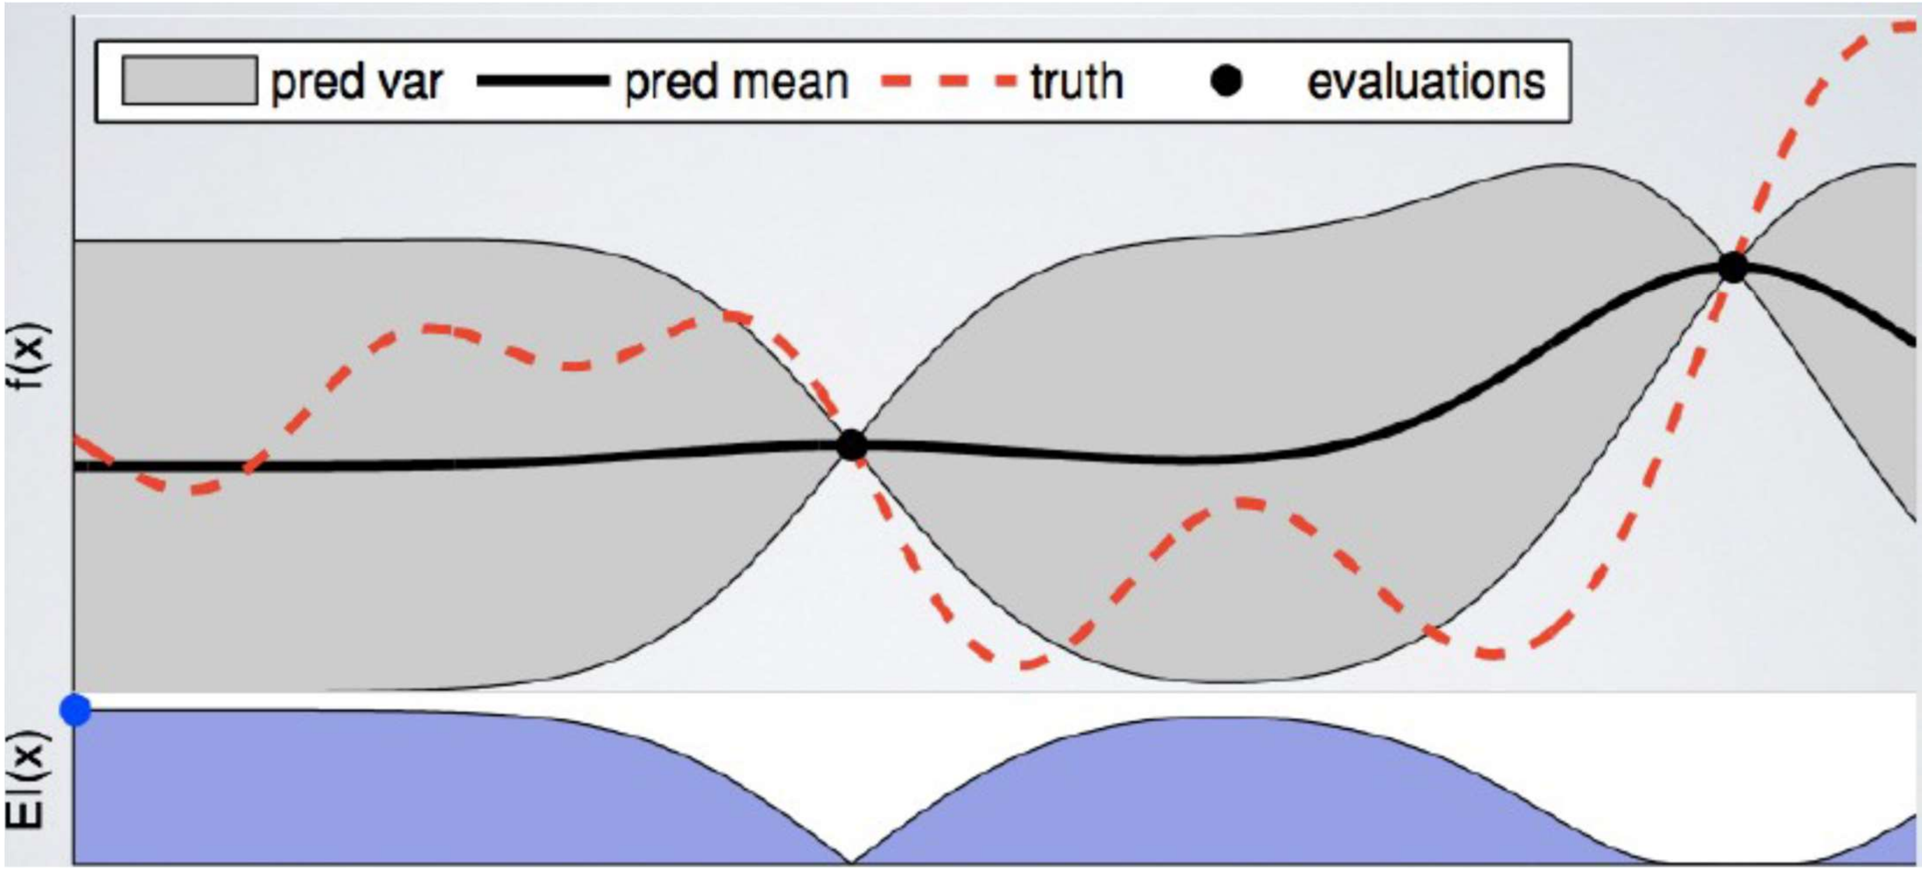
pred var
pred mean
truth
evaluations
f(x)
EI(x)

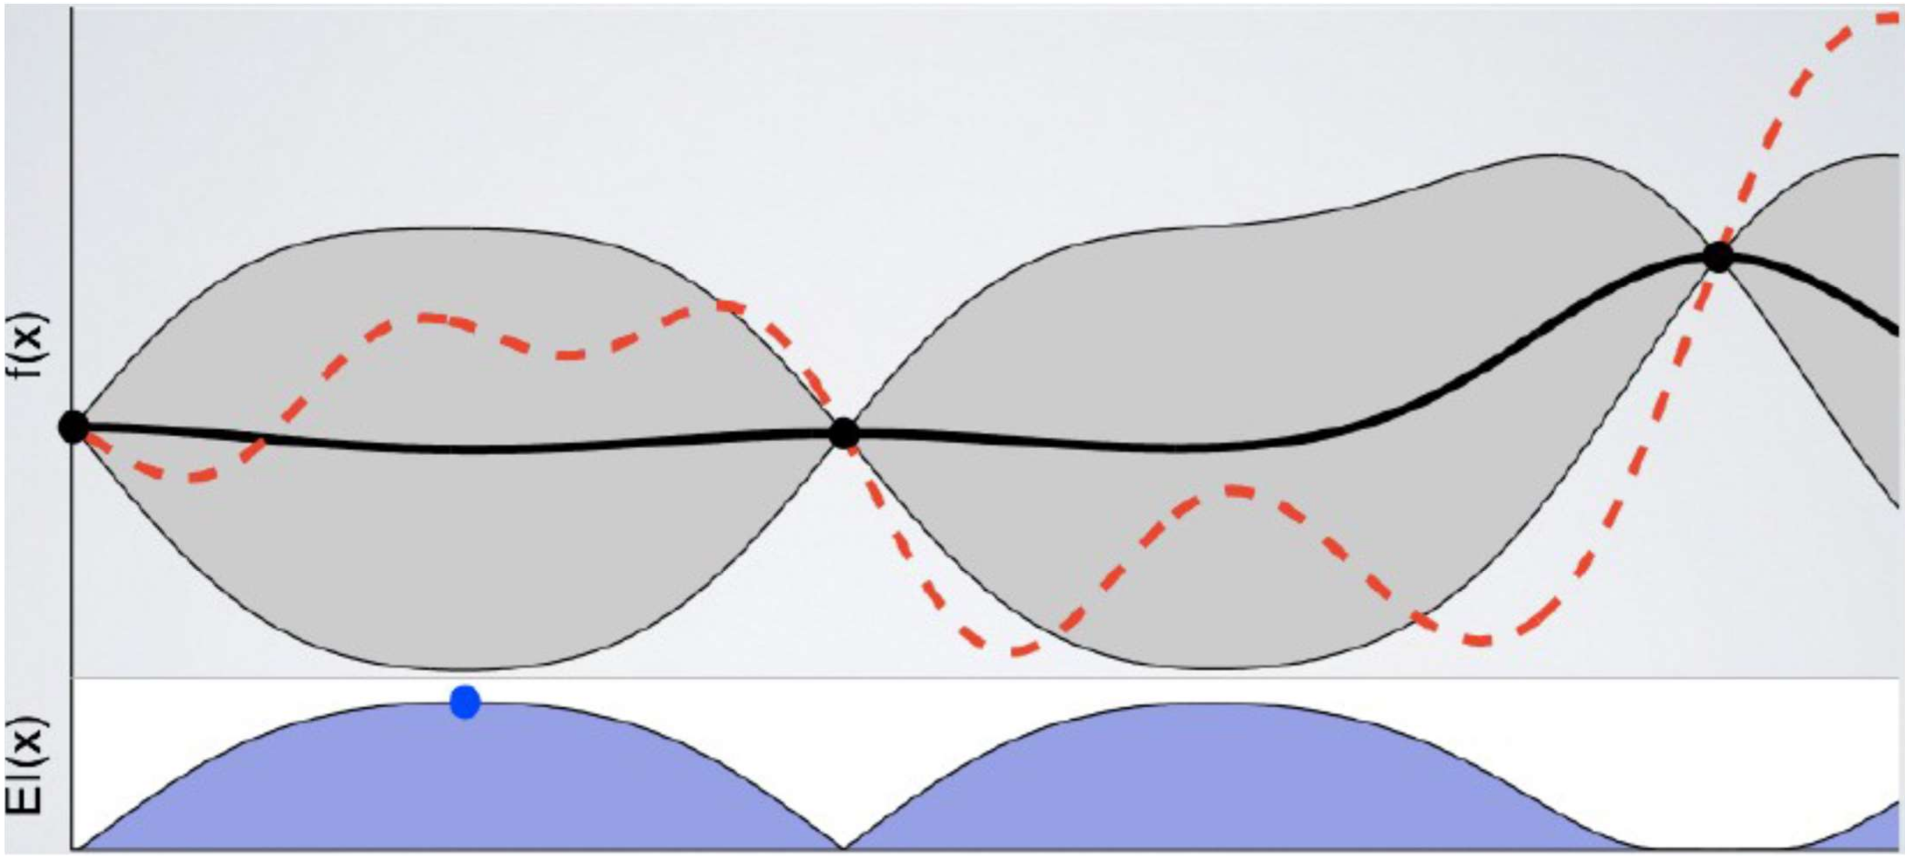
f(x)
EI(x)

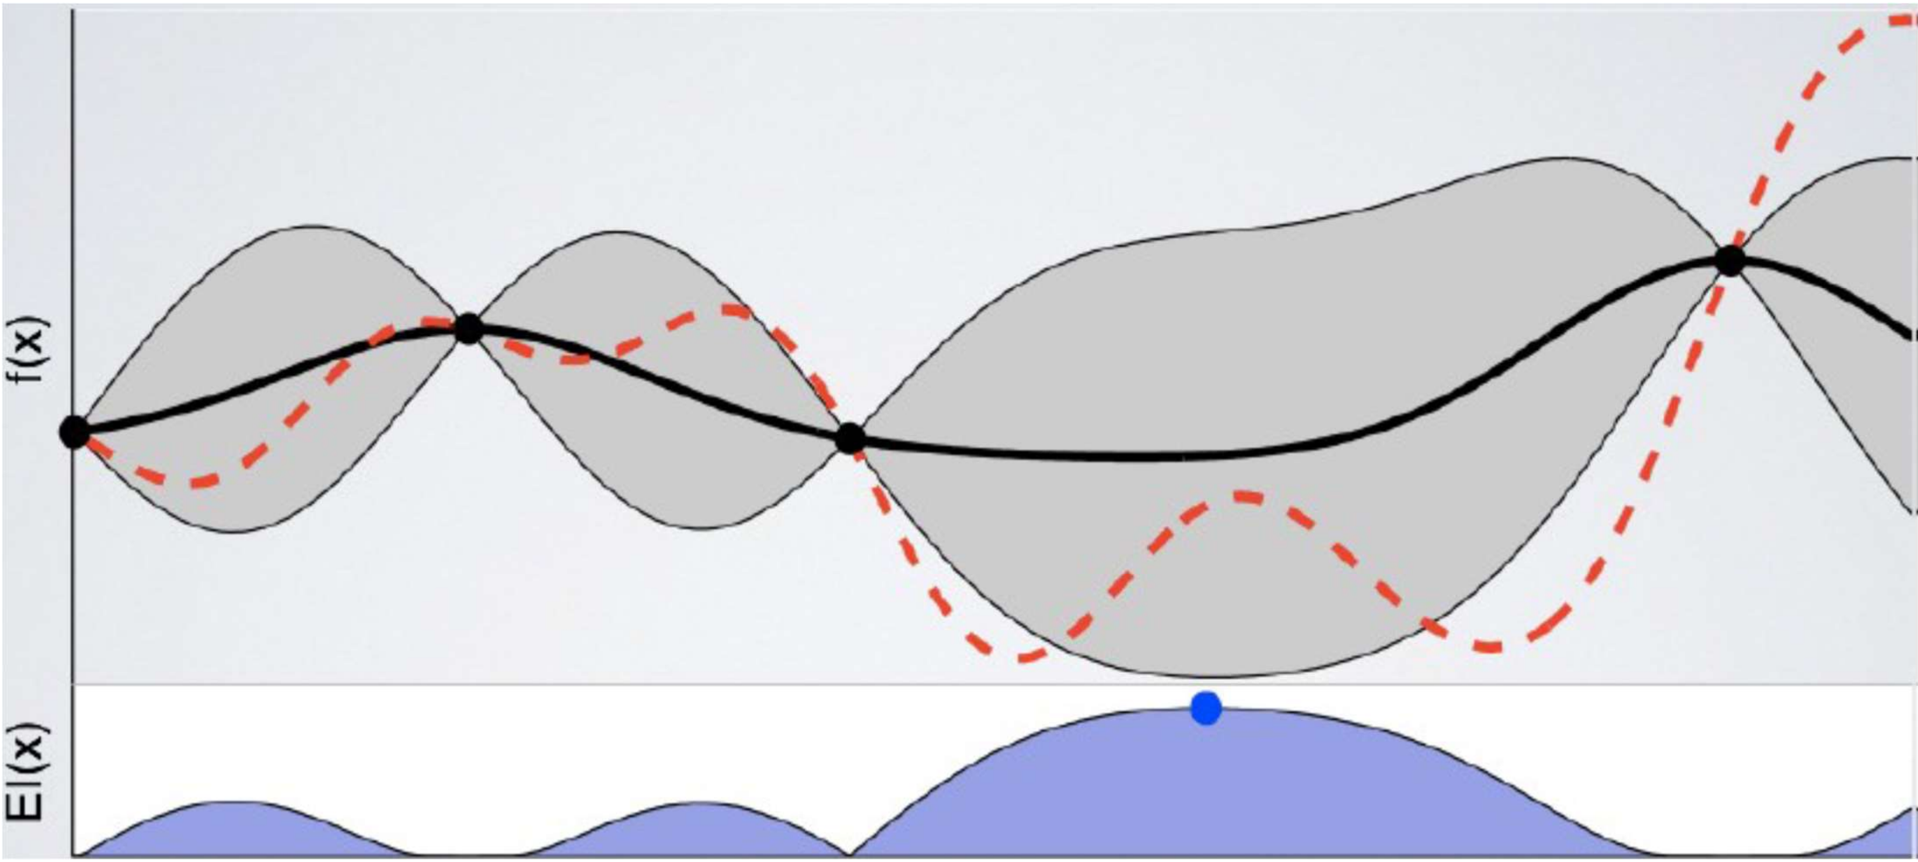
f(x)
EI(x)

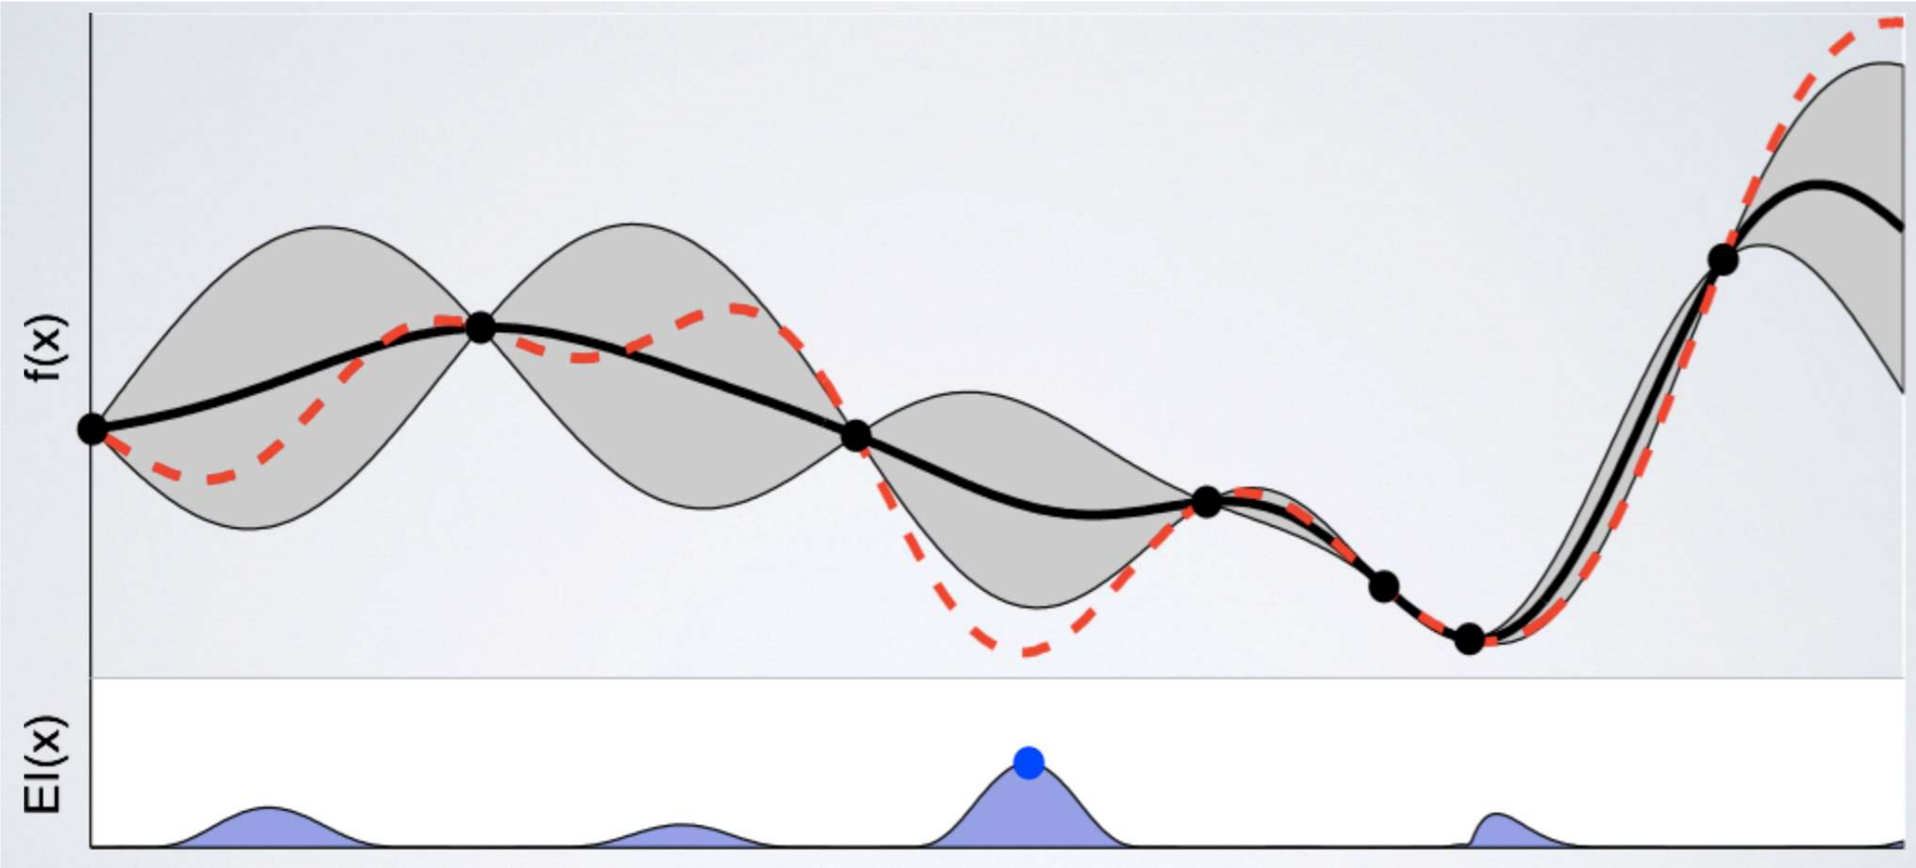
f(x)
EI(x)

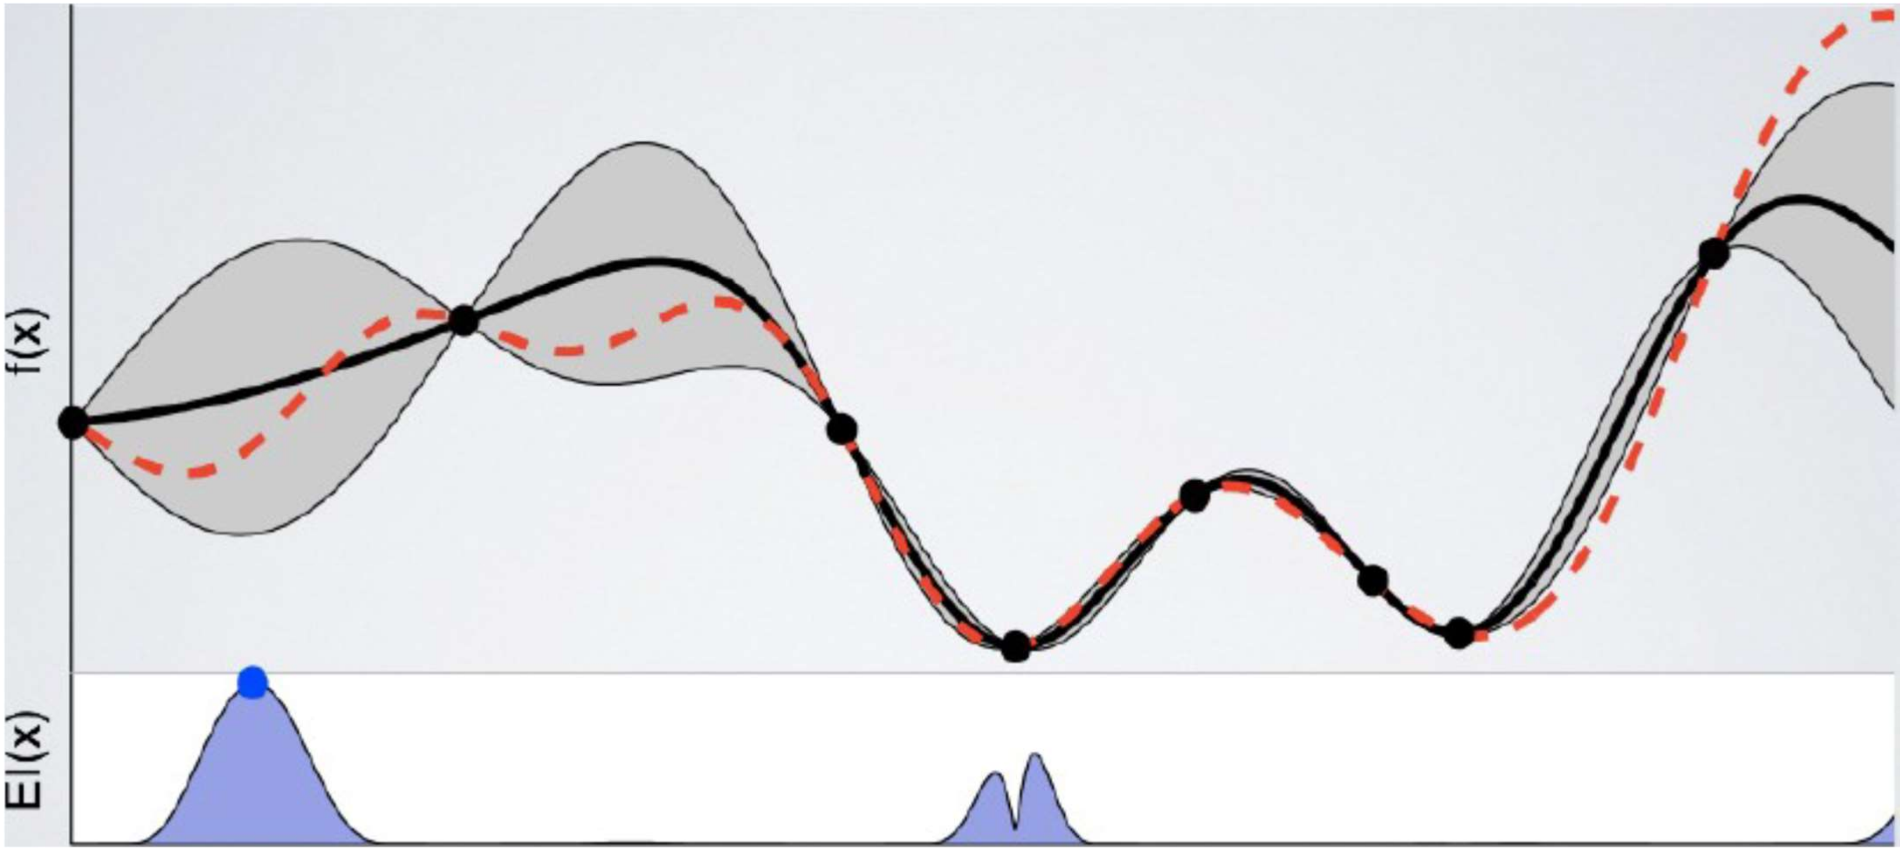

f(x)
EI(x)

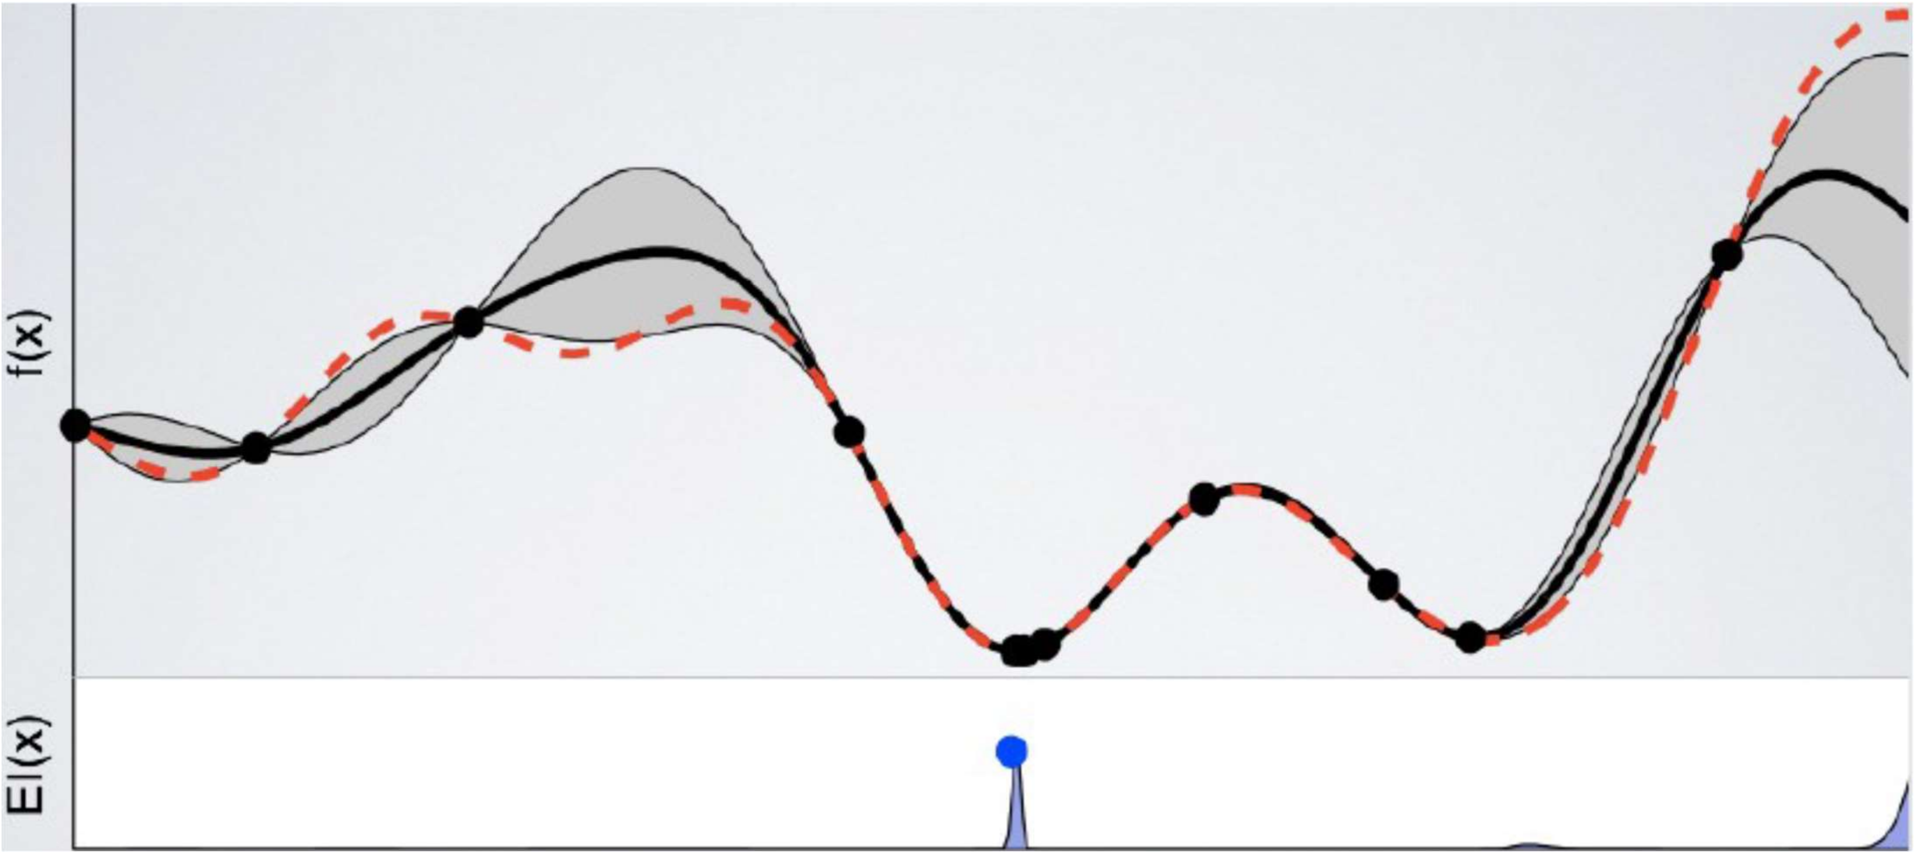

f(x)
EI(x)

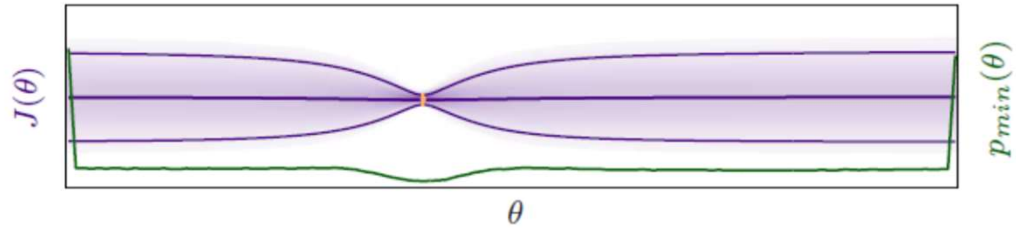

(a) 1 evaluation

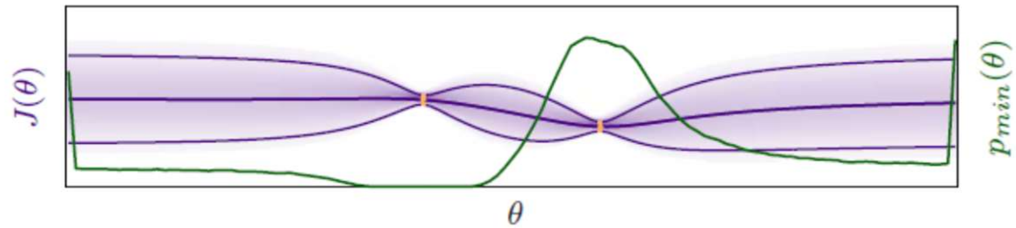

(b) 2 evaluations

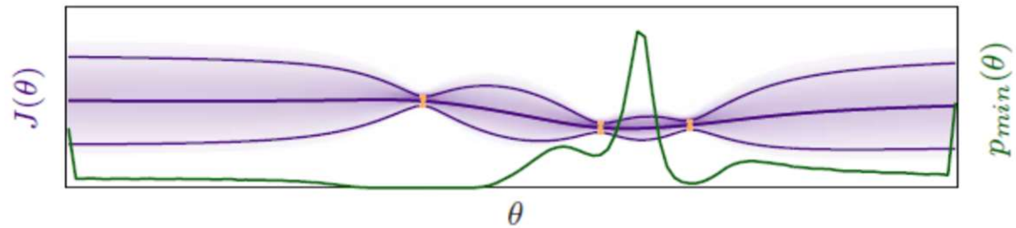

(c) 3 evaluations

Evolution of an example Gaussian process for three successive function evaluations (orange dots)

Approximated probability distribution over the location of the minimum $p_{min}(\theta)$ in green