Optimization as Estimation with Gaussian Processes in Bandit Settings

Zi Wang, Bolei Zhou, Stefanie Jegelka

ziw, bzhou, stefje@csail.mit.edu

May 9, 2016

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x}\in\mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

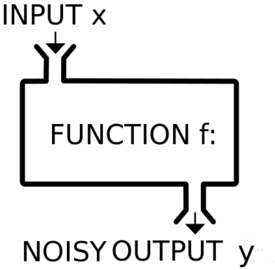

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x}\in\mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.

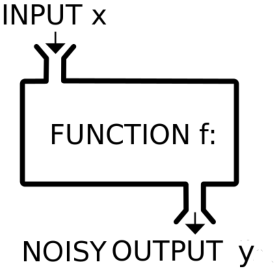

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x}\in\mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.
- Sequential queries.

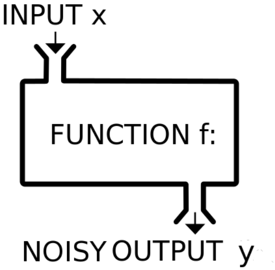

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x} \in \mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.
- Sequential queries.

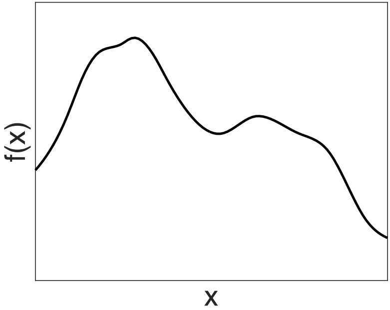

At round t,

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x}\in\mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.
- Sequential queries.

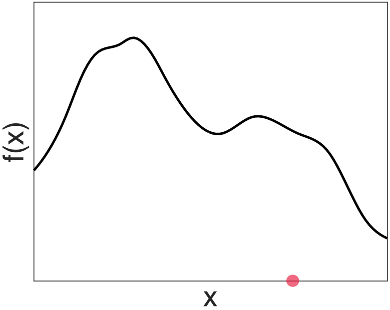

At round t,

- Choose $\boldsymbol{x}_t$;

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x}\in\mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.
- Sequential queries.

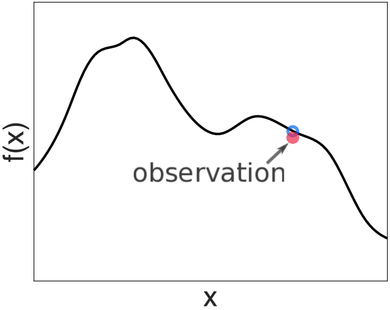

At round t,

- Choose $\boldsymbol{x}_t$;
- Observe $y_t = f(\boldsymbol{x}_t) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2)$;

Black-box function optimization in the bandit setting

$$\underset{\boldsymbol{x} \in \mathfrak{X}}{\text{maximize}} \quad f(\boldsymbol{x})$$

- f is expensive to evaluate.
- Sequential queries.

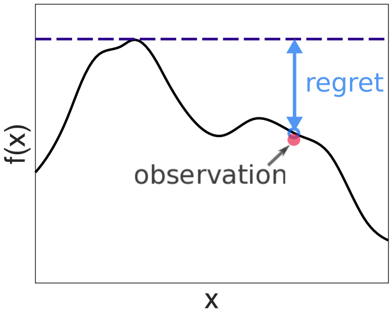

At round t,

- Choose $\boldsymbol{x}_t$;
- Observe $y_t = f(\boldsymbol{x}_t) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2)$;

Goal: Minimize cumulative regret $R_T = \sum_{t=1}^{T} \left(\max_{\boldsymbol{x} \in \mathfrak{X}} f(\boldsymbol{x}) - f(\boldsymbol{x}_t)\right)$

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

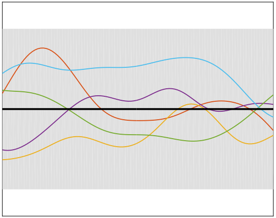

Prior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$

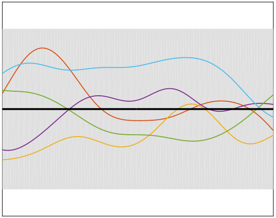

Prior distribution

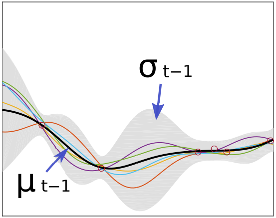

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$

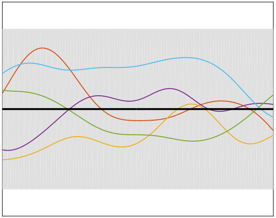

Prior distribution

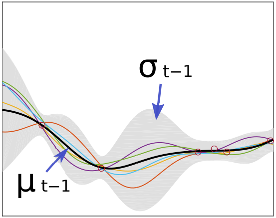

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$

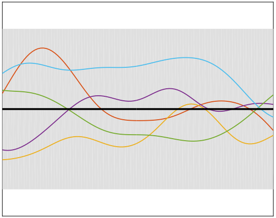

Prior distribution

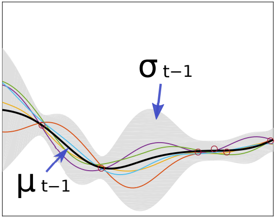

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$ by optimizing an **acquisition function**

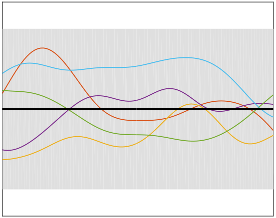

Prior distribution

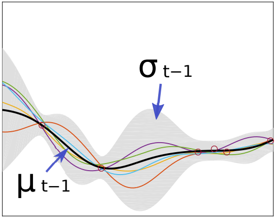

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$ by optimizing an **acquisition function**

Examples:

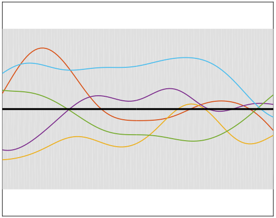

Prior distribution

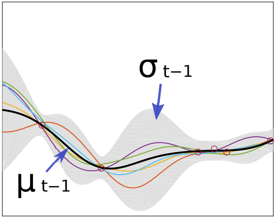

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma_{t-1}^2(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$ by optimizing an **acquisition function**

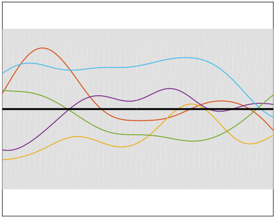

Prior distribution

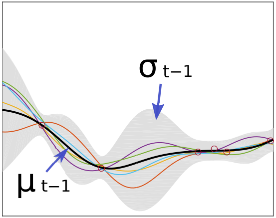

Posterior distribution

Examples:

- $\mathrm{PI}(\boldsymbol{x}) = \Pr[f(\boldsymbol{x}) > \theta_t]$ (Kushner, 1964)

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma^2_{t-1}(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$ by optimizing an **acquisition function**

Examples:

- $\mathrm{PI}(\boldsymbol{x}) = \Pr[f(\boldsymbol{x}) > \theta_t]$ (Kushner, 1964)
- $\mathrm{EI}(\boldsymbol{x}) = \mathbb{E}[(f(\boldsymbol{x}) - \theta_t)_+]$ (Močkus, 1974)

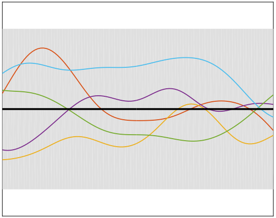

Prior distribution

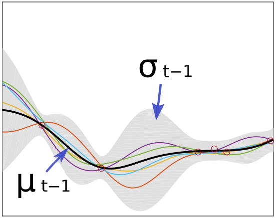

Posterior distribution

Gaussian process optimization

Assume $f \sim GP(\mu, k)$.

At round t,

- Predict $\mu_{t-1}(\boldsymbol{x})$ and $\sigma_{t-1}^2(\boldsymbol{x})$
- Pick an input $\boldsymbol{x}_t$ by optimizing an **acquisition function**

Examples:

- $\mathrm{PI}(\boldsymbol{x}) = \Pr[f(\boldsymbol{x}) > \theta_t]$ (Kushner, 1964)
- $\mathrm{EI}(\boldsymbol{x}) = \mathbb{E}[(f(\boldsymbol{x}) - \theta_t)_+]$ (Močkus, 1974)
- $\mathrm{UCB}(\boldsymbol{x}) = \mu_{t-1}(\boldsymbol{x}) + \lambda_t \sigma_{t-1}(\boldsymbol{x})$ (Srinivas et al., 2010)

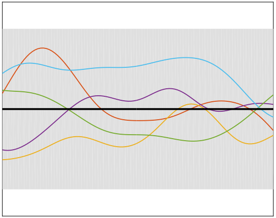

Prior distribution

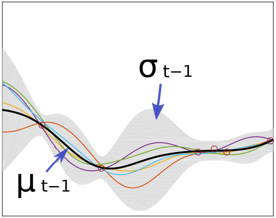

Posterior distribution

Existing acquisition functions

Upper Confidence Bound (GP-UCB) (Srinivas et al., 2010)

$$\boldsymbol{x}_t = \arg\max_{\boldsymbol{x}\in\mathfrak{X}} \mu_{t-1}(\boldsymbol{x}) + \lambda_t \sigma_{t-1}(\boldsymbol{x})$$

Can set λ_t that guarantees high-probability sub-linear regret in theory.

Existing acquisition functions

Upper Confidence Bound (GP-UCB) (Srinivas et al., 2010)

$$\boldsymbol{x}_t = \arg\max_{\boldsymbol{x} \in \mathfrak{X}} \mu_{t-1}(\boldsymbol{x}) + \lambda_t \sigma_{t-1}(\boldsymbol{x})$$

Can set λ_t that guarantees high-probability sub-linear regret in theory.

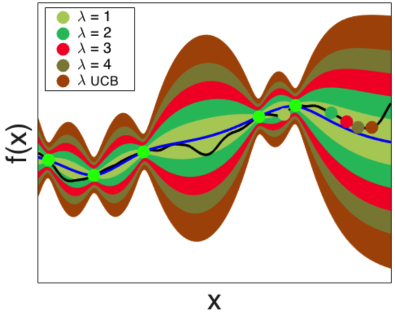

Existing acquisition functions

Upper Confidence Bound (GP-UCB) (Srinivas et al., 2010)

$$\boldsymbol{x}_t = \arg\max_{\boldsymbol{x}\in\mathfrak{X}} \mu_{t-1}(\boldsymbol{x}) + \lambda_t \sigma_{t-1}(\boldsymbol{x})$$

Can set λ_t that guarantees high-probability sub-linear regret in theory.

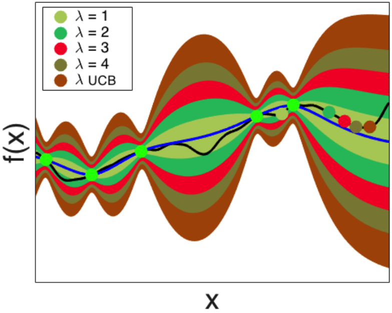

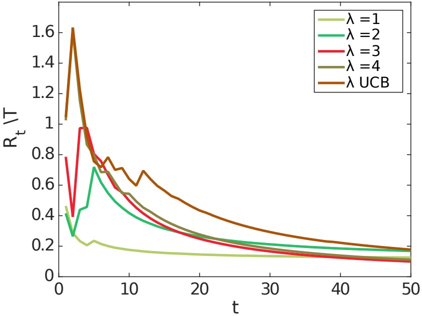

A new method: query the most likely arg max

Given the observations, what is the most likely arg max of the function?

A new method: query the most likely arg max

Given the observations, what is the most likely arg max of the function?

Notice that, for any $\boldsymbol{x} \in \mathfrak{X}$, $f(\boldsymbol{x})$ has a Gaussian distribution.

A new method: query the most likely arg max

Given the observations, what is the most likely arg max of the function?

Notice that, for any $\boldsymbol{x} \in \mathfrak{X}$, $f(\boldsymbol{x})$ has a Gaussian distribution.

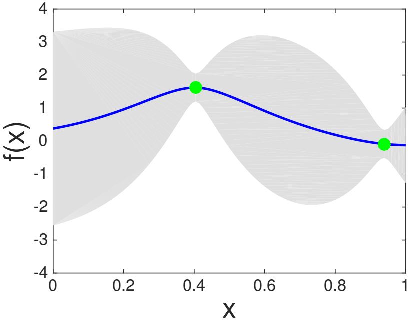

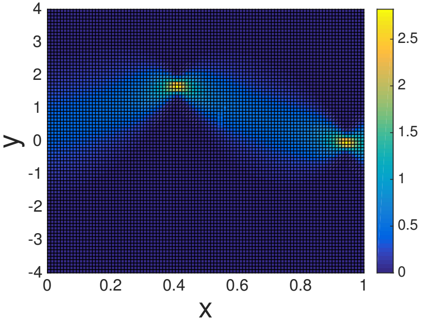

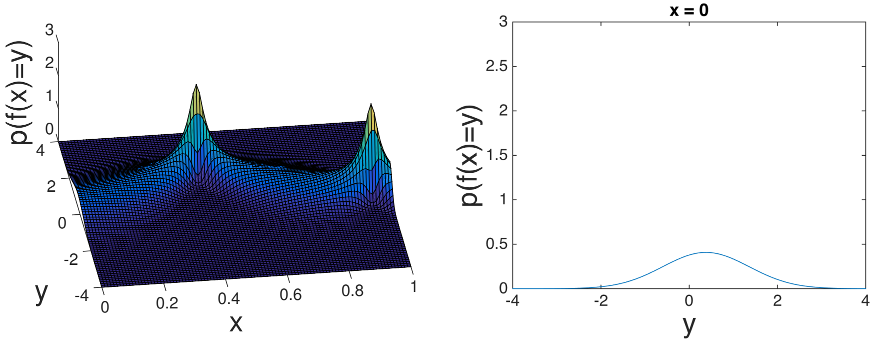
p(f(x)=y)
3
2
1
0
4
2
0
-2
-4
y
0
0.2
0.4
0.6
0.8
1
x
x = 0
p(f(x)=y)
3
2.5
2
1.5
1
0.5
0
-4
-2
0
2
4
y

Optimization as estimation

EST: estimate the $\arg\max$ of the function f.

Optimization as estimation

EST: estimate the $\arg\max$ of the function f.

1. What is the function maximum?

Optimization as estimation

EST: estimate the $\arg\max$ of the function f.

1. What is the function maximum?
2. How likely is $f(\boldsymbol{x})$ the maximum?

Step 1: Estimate the function maximum

1. What is the function maximum?

Step 1: Estimate the function maximum

1. What is the function maximum?
   Consider discrete $\mathfrak{X}$ and negligible noise,

$$\hat{m} = \mathbb{E}[\max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x})] = \max_{\tau\in[1,t-1]} y_\tau + \int_{\max_{\tau\in[1,t-1]} y_\tau}^{\infty} \Pr[\max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x}) > w] \,\mathrm{d}w$$

Step 1: Estimate the function maximum

1. What is the function maximum?
   Consider discrete $\mathfrak{X}$ and negligible noise,

$$\hat{m} = \mathbb{E}[\max_{\boldsymbol{x} \in \mathfrak{X}} f(\boldsymbol{x})] = \max_{\tau \in [1, t-1]} y_\tau + \int_{\max_{\tau \in [1,t-1]} y_\tau}^{\infty} \Pr[\max_{\boldsymbol{x} \in \mathfrak{X}} f(\boldsymbol{x}) > w] \, \mathrm{d}w$$

   - Approximate the joint Gaussian with independent Gaussians

$$g(w) = 1 - \Pr[f(\boldsymbol{x}) \leq w, \forall \boldsymbol{x} \in \mathfrak{X}] \approx 1 - \prod_{\boldsymbol{x} \in \mathfrak{X}} \Phi\left(\frac{w - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right)$$

Step 1: Estimate the function maximum

1. What is the function maximum?
   Consider discrete $\mathfrak{X}$ and negligible noise,

$$\hat{m} = \mathbb{E}[\max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x})] = \max_{\tau\in[1,t-1]} y_\tau + \int_{\max_{\tau\in[1,t-1]} y_\tau}^{\infty} \Pr[\max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x}) > w] \, \mathrm{d}w$$

- Approximate the joint Gaussian with independent Gaussians

$$g(w) = 1 - \Pr[f(\boldsymbol{x}) \leq w, \forall \boldsymbol{x} \in \mathfrak{X}] \approx 1 - \prod_{\boldsymbol{x}\in\mathfrak{X}} \Phi\left(\frac{w - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right)$$

- Integrate numerically (ESTn) or approximately (ESTa)

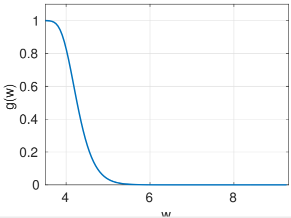

Step 2: calculate the probability that $\boldsymbol{x}$ is the $\arg\max$

2. How likely is $f(\boldsymbol{x})$ the maximum?

Step 2: calculate the probability that $\boldsymbol{x}$ is the arg max

2. How likely is $f(\boldsymbol{x})$ the maximum?

$$\Pr[f(\boldsymbol{x}) \text{ is the maximum} | \hat{m}] \approx Q\left(\frac{\hat{m} - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right) \prod_{\boldsymbol{x}' \neq \boldsymbol{x}} \Phi\left(\frac{\hat{m} - \mu(\boldsymbol{x}')}{\sigma(\boldsymbol{x}')}\right)$$

Step 2: calculate the probability that $\boldsymbol{x}$ is the arg max

2. How likely is $f(\boldsymbol{x})$ the maximum?

$$\Pr[f(\boldsymbol{x}) \text{ is the maximum} | \hat{m}] \approx \overbrace{Q\left(\frac{\hat{m} - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right)}^{\Pr[f(\boldsymbol{x}) \geq \hat{m}]} \prod_{\boldsymbol{x}' \neq \boldsymbol{x}} \Phi\left(\frac{\hat{m} - \mu(\boldsymbol{x}')}{\sigma(\boldsymbol{x}')}\right)$$

Step 2: calculate the probability that $\boldsymbol{x}$ is the $\arg\max$

2. How likely is $f(\boldsymbol{x})$ the maximum?

$$\Pr[f(\boldsymbol{x}) \text{ is the maximum} | \hat{m}] \approx \overbrace{Q\left(\frac{\hat{m} - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right)}^{\Pr[f(\boldsymbol{x}) \geq \hat{m}]} \underbrace{\prod_{\boldsymbol{x}' \neq \boldsymbol{x}} \Phi\left(\frac{\hat{m} - \mu(\boldsymbol{x}')}{\sigma(\boldsymbol{x}')}\right)}_{\Pr[\forall \boldsymbol{x}' \neq \boldsymbol{x}, f(\boldsymbol{x}') < \hat{m}]}$$

Step 2: calculate the probability that $\boldsymbol{x}$ is the arg max

2. How likely is $f(\boldsymbol{x})$ the maximum?

$$\Pr[f(\boldsymbol{x}) \text{ is the maximum} | \hat{m}] \approx \overbrace{Q\left(\frac{\hat{m} - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}\right)}^{\Pr[f(\boldsymbol{x}) \geq \hat{m}]} \underbrace{\prod_{\boldsymbol{x}' \neq \boldsymbol{x}} \Phi\left(\frac{\hat{m} - \mu(\boldsymbol{x}')}{\sigma(\boldsymbol{x}')}\right)}_{\Pr[\forall \boldsymbol{x}' \neq \boldsymbol{x}, f(\boldsymbol{x}') < \hat{m}]}$$

$$\arg\max_{\boldsymbol{x} \in \mathfrak{X}} \Pr[f(\boldsymbol{x}) \text{ is the maximum} | \hat{m}] = \arg\min_{\boldsymbol{x} \in \mathfrak{X}} \frac{\hat{m} - \mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}$$

Connections to GP-UCB and PI

EST

$$\Pr[\boldsymbol{x} = \arg\max f(\boldsymbol{x}) | \hat{m}]$$

Connections to GP-UCB and PI

EST

$$\Pr[\boldsymbol{x} = \arg\max f(\boldsymbol{x}) | \hat{m}]$$

$\theta = \hat{m}$

PI

$$\mathrm{PI}(\boldsymbol{x}) = \Pr[f(\boldsymbol{x}) > \theta]$$

Connections to GP-UCB and PI

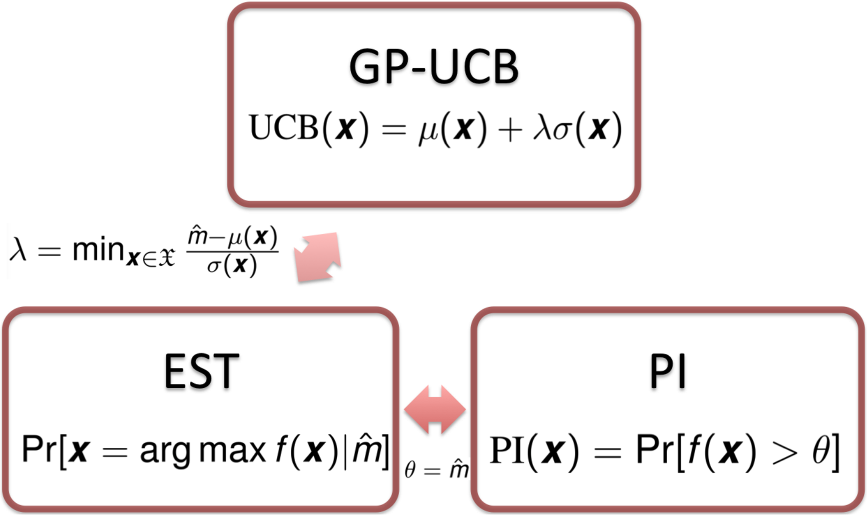

Connections to GP-UCB and PI

GP-UCB

$$\mathrm{UCB}(\boldsymbol{x}) = \mu(\boldsymbol{x}) + \lambda\sigma(\boldsymbol{x})$$

$$\lambda = \min_{\boldsymbol{x}\in\mathfrak{X}} \frac{\hat{m}-\mu(\boldsymbol{x})}{\sigma(\boldsymbol{x})}$$

$$\theta = \max_{\boldsymbol{x}\in\mathfrak{X}} \mu(\boldsymbol{x}) + \lambda\sigma(\boldsymbol{x})$$

EST

$$\Pr[\boldsymbol{x} = \arg\max f(\boldsymbol{x}) | \hat{m}]$$

$$\theta = \hat{m}$$

PI

$$\mathrm{PI}(\boldsymbol{x}) = \Pr[f(\boldsymbol{x}) > \theta]$$

Connections to GP-UCB and PI

- At round t, pick the input that is most likely to reach a target value.

$$\hat{m}_t = \begin{cases} \max_{\boldsymbol{x}\in\mathfrak{X}} \mu_{t-1}(\boldsymbol{x}) + \lambda_t \sigma_{t-1}(\boldsymbol{x}) & \text{GP-UCB} \\ \theta_t & \text{PI} \\ \mathbb{E}[\max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x})] & \text{EST} \end{cases}$$

$$\boldsymbol{x}_t \leftarrow \arg\min_{\boldsymbol{x}\in\mathfrak{X}} \frac{\hat{m}_t - \mu_{t-1}(\boldsymbol{x})}{\sigma_{t-1}(\boldsymbol{x})}$$

Regret bounds

Theorem (Regret bounds for EST)

Assume $\hat{m}_t \geq \max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x}), \forall t \in [1, T]$. *Then,*

$$\mathbb{E}[R_T] \leq \nu_{t^*}\sqrt{CT\gamma_T}.$$

With probability at least $1-\delta$*,*

$$R_T \leq (\nu_{t^*} + \zeta_T)\sqrt{CT\gamma_T},$$

$C = \frac{2}{\log(1+\sigma^{-2})}$, $\nu_t \triangleq \min_{\boldsymbol{x}\in\mathfrak{X}} \frac{\hat{m}_t - \mu_{t-1}(\boldsymbol{x})}{\sigma_{t-1}(\boldsymbol{x})}$, $t^* = \arg\max_t \nu_t$.
$k(\boldsymbol{x}, \boldsymbol{x}') \leq 1$, $\gamma_T = \max_{A\subseteq\mathfrak{X}, |A|\leq T} I(\boldsymbol{y}_A, \boldsymbol{f}_A)$, $\zeta_T = (2\log(\frac{T}{2\delta}))^{\frac{1}{2}}$.

Regret bounds

Theorem (Regret bounds for EST)

Assume $\hat{m}_t \geq \max_{\boldsymbol{x}\in\mathfrak{X}} f(\boldsymbol{x}), \forall t \in [1, T]$. *Then,*

$$\mathbb{E}[R_T] \leq \nu_{t^*}\sqrt{CT\gamma_T}.$$

With probability at least $1 - \delta$*,*

$$R_T \leq (\nu_{t^*} + \zeta_T)\sqrt{CT\gamma_T},$$

$C = \frac{2}{\log(1+\sigma^{-2})}$, $\nu_t \triangleq \min_{\boldsymbol{x}\in\mathfrak{X}} \frac{\hat{m}_t - \mu_{t-1}(\boldsymbol{x})}{\sigma_{t-1}(\boldsymbol{x})}$, $t^* = \arg\max_t \nu_t$.
$k(\boldsymbol{x}, \boldsymbol{x}') \leq 1$, $\gamma_T = \max_{A\subseteq\mathfrak{X}, |A|\leq T} I(\boldsymbol{y}_A, \boldsymbol{f}_A)$, $\zeta_T = (2\log(\frac{T}{2\delta}))^{\frac{1}{2}}$.

Estimating an upper bound on the function maximum

Slepian's Comparison Lemma (Slepian, 1962; Massart, 2007)

Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$ be two multivariate Gaussian random vectors with the same mean and variance, such that

$$\mathbb{E}[\boldsymbol{v}_i \boldsymbol{v}_j] \leq \mathbb{E}[\boldsymbol{u}_i \boldsymbol{u}_j], \forall i, j.$$

Then,

$$\mathbb{E}[\sup_{i \in [1,n]} \boldsymbol{v}_i] \geq \mathbb{E}[\sup_{i \in [1,n]} \boldsymbol{u}_i].$$

Estimating an upper bound on the function maximum

Slepian's Comparison Lemma (Slepian, 1962; Massart, 2007)

Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$ be two multivariate Gaussian random vectors with the same mean and variance, such that

$$\mathbb{E}[\boldsymbol{v}_i \boldsymbol{v}_j] \leq \mathbb{E}[\boldsymbol{u}_i \boldsymbol{u}_j], \forall i, j.$$

Then,

$$\mathbb{E}[\sup_{i \in [1,n]} \boldsymbol{v}_i] \geq \mathbb{E}[\sup_{i \in [1,n]} \boldsymbol{u}_i].$$

Ignoring positive covariance gives higher expected maximum.

Experiments

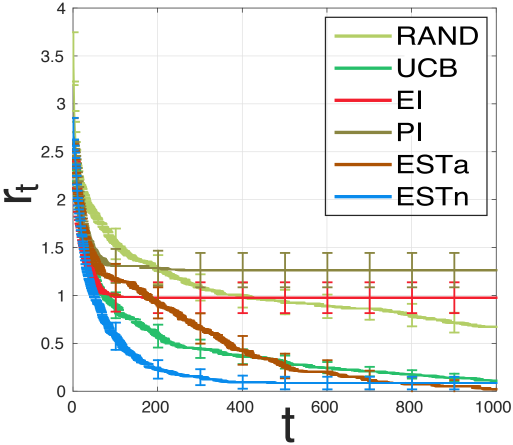

Experiments

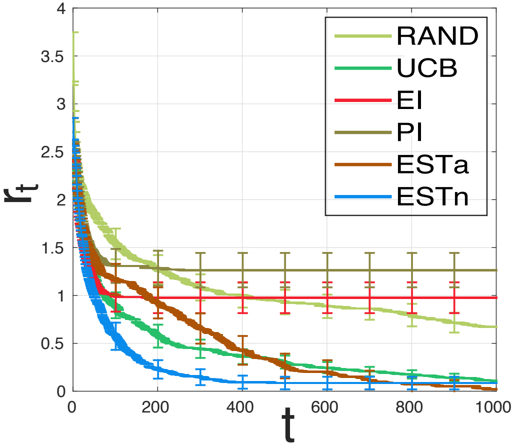

Robotics

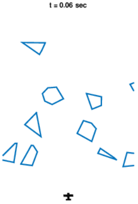

Vision

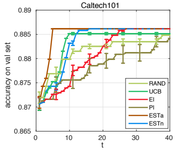

More results at Session 2 Poster 47

Summary: Optimization as Estimation

Summary: Optimization as Estimation

- A new BO strategy from the viewpoint of estimating $\arg\max$.
- Adaptively tuning λ and θ in GP-UCB and PI.

Summary: Optimization as Estimation

- A new BO strategy from the viewpoint of estimating $\arg\max$.
- Adaptively tuning λ and θ in GP-UCB and PI.
- Sub-linear regret bounds and good empirical results.

Summary: Optimization as Estimation

- A new BO strategy from the viewpoint of estimating arg max.
- Adaptively tuning λ and θ in GP-UCB and PI.
- Sub-linear regret bounds and good empirical results.
- Source code: `https://github.com/zi-w/GP-EST`

Summary: Optimization as Estimation

- A new BO strategy from the viewpoint of estimating $\arg\max$.
- Adaptively tuning λ and θ in GP-UCB and PI.
- Sub-linear regret bounds and good empirical results.
- Source code: `https://github.com/zi-w/GP-EST`

Session 2 (May 10)
Poster 47