CSE 574 Planning and Learning Methods in AI

Ransalu Senanayake

www.ransalu.com

Week 2

Week	**Date**	**Theory**
Week 1	08/18/2023	Introduction to decision-making
Week 2	08/25/2023	Myopic planning
Week 3	09/01/2023	Myopic planning
Week 4	09/08/2023	Non-myopic planning and reinforcement learning
Week 5	09/15/2023	Non-myopic planning and reinforcement learning
Week 6	09/22/2023	Imitation learning for decision-making
Week 7	09/29/2023	Imitation learning for decision-making
Week 8	10/06/2023	Human-in-the-loop planning
Week 9	10/13/2023	Human-in-the-loop planning
Week 10	10/20/2023	Task and motion planning
Week 11	10/27/2023	Task and motion planning
Week 12	11/03/2023	Multiagent planning and decision-making under uncertainty
~~Week 13~~	~~11/10/2023~~	Veterans Day (No class)
Week 14	11/17/2023	Classical planning and hierarchical planning/learning
~~Week 15~~	~~11/24/2023~~	Thanksgiving (No class)
Week 16	12/01/2023	Project presentations

Bayesian Optimization (BO)

- We want to find the *global minimum* of an <u>expensive</u> to evaluate function
- We don't know much about the function but we can evaluate it
  - We don't know about 1^{st} and 2nd derivates. Hence, gradient methods (e.g., Newton/quasi-newton method) is not possible. BO is *derivative free*.
  - We don't know if the function has a special structure (e.g., convex). Hence, we can't use convex optimization techniques. BO is a *black-box optimization* technique.
- Myopic planning (i.e., one-step lookahead)
- A.k.a. surrogate optimization

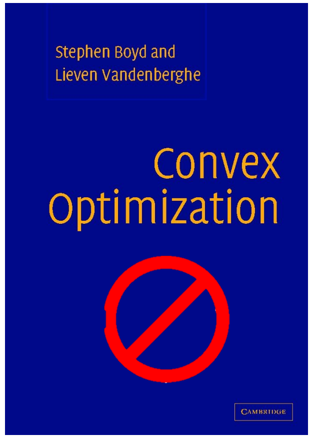

Applications

Bayesian Optimisation for Intelligent Environmental Monitoring

Roman Marchant and Fabio Ramos

***Abstract*— Environmental Monitoring (EM) is typically performed using sensor networks that collect measurements in predefined static locations. The possibility of having one or more autonomous robots to perform this task increases versatility and reduces the number of necessary sensor nodes to cover the same area. However, several problems arise when making use of autonomous moving robots for EM. The main challenges**

This paper presents a novel approach for active sampling, using *Bayesian Optimisation* (BO) [2–4]. The BO framework allows a mobile robot to choose sensing locations taking into account the uncertainty of the understanding of the phenomenon, the expected value of the studied variable (e.g. temperature, humidity, ambient light), and the cost to travel

Applications

auto-sklearn in one image

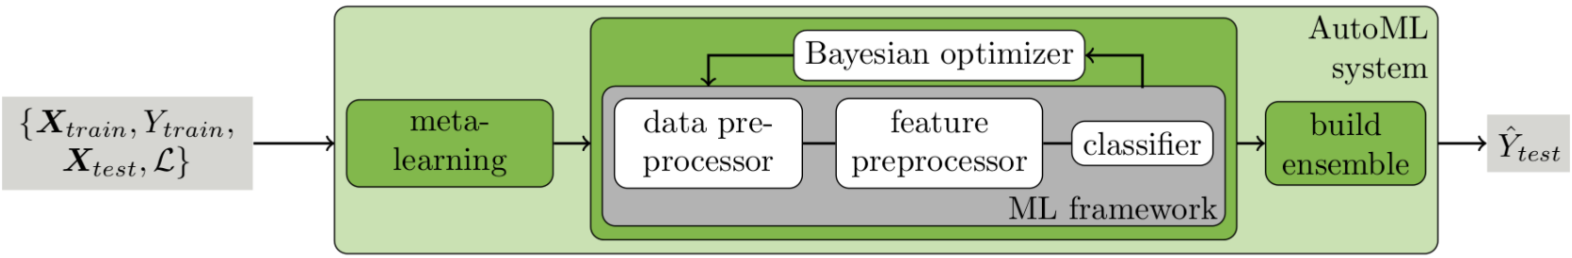

Google Cloud

Contact sales

Blog

Solutions & technology ∨ Ecosystem ∨ Developers & Practitioners Transform with Google Cloud

Google Cloud

Hyperparameter tuning in Cloud Machine Learning Engine using Bayesian Optimization

August 10, 2017

Puneith Kaul
Software Engineering Manager

Daniel Golovin
Staff Software Engineer, Google Brain

The BO algorithm

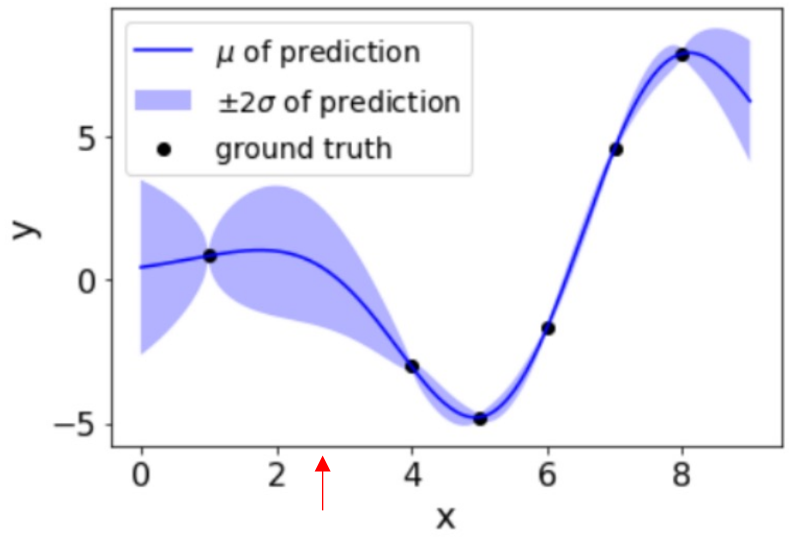

Minimize the black-box function by taking samples

1) stay closer to the current best (exploitation)
2) sample from more *uncertain* areas (exploration)

```
for N iterations
   Evaluate the point to sample using the acquisition function using the dataset
   Add the new minimum to the dataset
end
```

PI
EI
GP-UCB

Probability of Improvement (PI)

For now, assume an oracle gives us $\hat{\mu}$ and $\hat{\sigma}$. Later, we'll learn how to build the oracle using a Gaussian process.

$$u(x) = \begin{cases} 1, & f \leq f_* \\ 0, & Otherwise \end{cases}$$

$$a_{PI}(x) = \mathbb{E}[u(x)|x, \mathcal{X}] = \int_{-\infty}^{f_*} 1.\mathcal{N}(f; \hat{\mu}, \hat{\sigma})\mathrm{d}f = \Phi(f_*; \hat{\mu}, \hat{\sigma}) = \Phi\left(\frac{f_* - \hat{\mu}}{\hat{\sigma}}\right)$$

$$\mathcal{X} \leftarrow \mathcal{X} \cup \{\underset{x}{\operatorname{argmax}}\, a_{PI}(x)\}$$

- What is $a_{PI}(x)$ if $x \in \mathcal{X}$?
- What is the effect of $\hat{\mu}$ and $\hat{\sigma}$ on $a_{PI}(x)$?

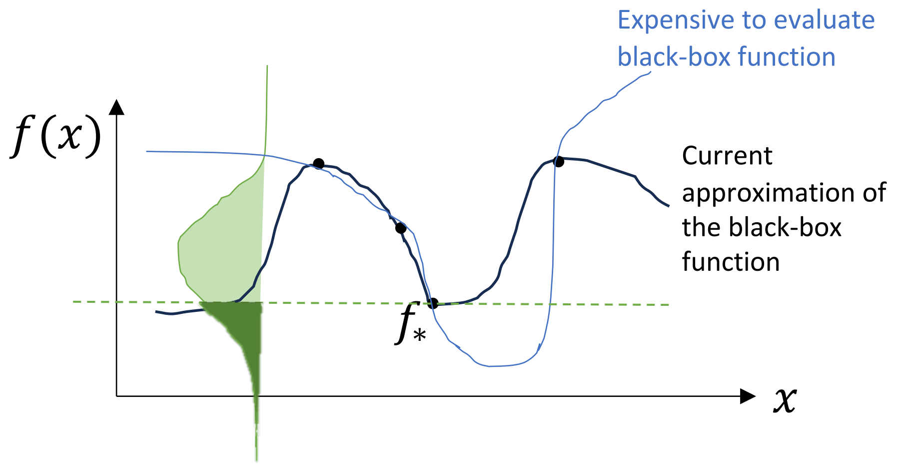

Expected Improvement (EI)

$$u(x) = \max(f_* - f, 0)$$

$$\begin{aligned} a_{EI}(x) = \mathbb{E}[u(x)|x, \mathcal{X}] &= \int_{-\infty}^{f_*} (f_* - f)\mathcal{N}(f; \hat{\mu}, \hat{\sigma})\mathrm{d}f. \\ &= (f_* - \hat{\mu})\Phi(f_*; \hat{\mu}, \hat{\sigma}) + \hat{\sigma}\mathcal{N}(f_*, \hat{\mu}, \hat{\sigma}) \\ &= (f_* - \hat{\mu})\Phi\left(\frac{f_* - \hat{\mu}}{\hat{\sigma}}\right) + \hat{\sigma}\mathcal{N}\left(\frac{f_* - \hat{\mu}}{\hat{\sigma}}\right) \end{aligned}$$

$$\mathcal{X} \leftarrow \mathcal{X} \cup \{\operatorname*{argmax}_x a_{PI}(x)\}$$

- What is $a_{EI}(x)$ if $x \in \mathcal{X}$?
- What is the effect of $\hat{\mu}$ and $\hat{\sigma}$ on $a_{EI}(x)$? Exploration vs. exploitation

Upper Confidence Bound (GP-UCB)

$$a_{GP-UCB}(x) = \hat{\mu} + \alpha\hat{\sigma}$$

$$\mathcal{X} \leftarrow \mathcal{X} \cup \{\underset{x}{\operatorname{argmax}}\, a_{GP-UCB}(x)\}$$

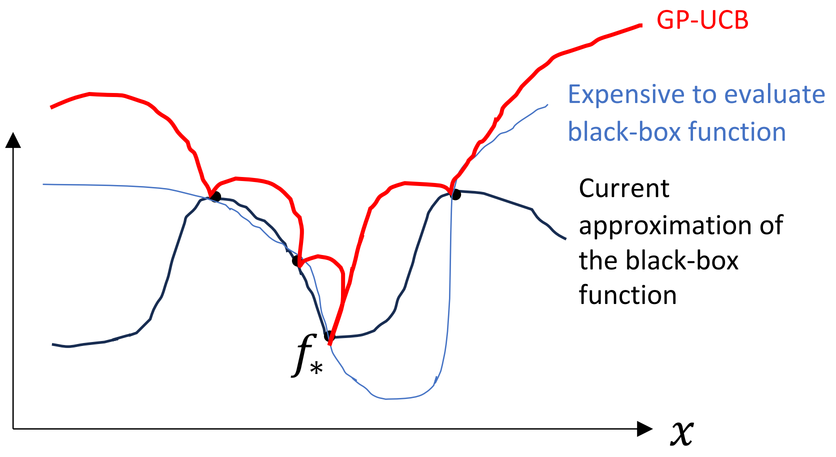

- What is $a_{GP-UCB}(x)$ when $\alpha \rightarrow 0$?
- What is $a_{GP-UCB}(x)$ when $\alpha \rightarrow \infty$

How to estimate the uncertainty?
i.e. How to build the oracle?

- Aleatoric vs. epistemic uncertainty
- Distributions over parameters
- Bayesian neural networks – feasible?
- Gaussian processes – the good, bad, and the ugly
    - Scalability with the number of points
    - Choice of kernel and learning its hyperparameters

Gaussian Process (GP) for Regression

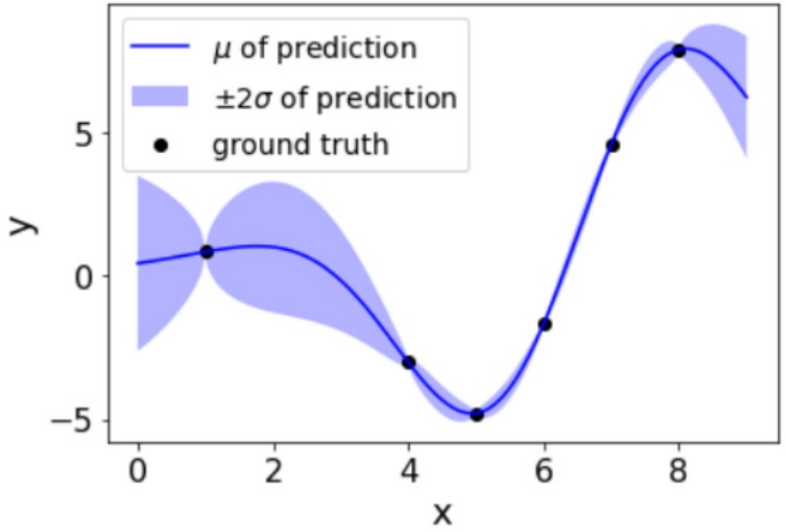

- A Bayesian non-parametric technique
- We have to learn the hyperparameters of the GP
- Can we work with high dimensional data?

Gaussian process prediction

$$\hat{\mu}(x) = k(x, X)\left(K(X, X) + \sigma_{noise}^{-1} I\right)^{-1} \mathbf{y}$$

$$\hat{\sigma}(x) = k(x, x) - k(x, X)\left(K(X, X) + \sigma_{noise}^{-1} I\right)^{-1} k(X, x)$$

Gaussian process fitting/learning

Learning a GP means, we learn it’s hyperparameters.
For instance, for a squared-exponential kernel,

$$k(x_1, x_2) = \alpha^2 \exp\left(-\frac{||x_1 - x_2||_2^2}{l^2}\right)$$

we have to learn α^2 and l^2.

The BO algorithm – In practice

```
Start with a few random points (hint: ~√d where d is the number of dimensions of x)
for N iterations
   Fit the GP with the current set of samples
   Evaluate the point to sample using the acquisition function and fitted GP's μ̂ and σ̂
   Append the dataset with the new sample/minimum
   (hint: Once in a while append the dataset with a random sample for better exploration)
end
```

- It's challenging when, say, $d > 30$
- If we know the black-box model has sharp transitions, then we can use a Mattern kernel

Other topics

- Is BO just an infinite-arm (every x) bandit?
- Multi-fidelity BO
- Multi-objective BO
- Multi-task BO
- Parallel BO
- Constrained BO
- Submodular optimization
- Theoretical aspects
  - E.g., Using GP-UCB, we can show that it can achieve the global optimum

GPyOpt

BoTorch