Myopic Posterior Sampling for Adaptive Goal Oriented Design of Experiments

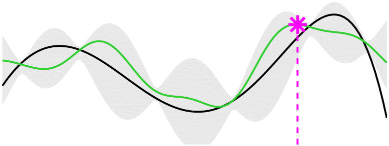

Kirthevasan Kandasamy, Willie Neiswanger, Reed Zhang,
Akshay Krishnamurthy, Jeff Schneider, Barnabás Póczos

Carnegie Mellon University   Auton Lab   Microsoft Research

ICML 2019

Example 1: Active Learning in Parametric Models

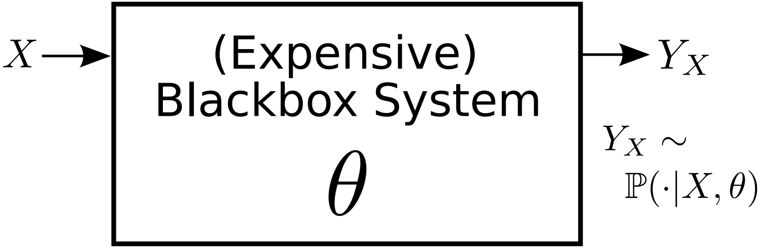

Goal: Learn parameter θ in as few experiments.

Algorithms: Active-Set-Select (Chaudhuri et al. 2015)

Example 2: Blackbox Optimisation

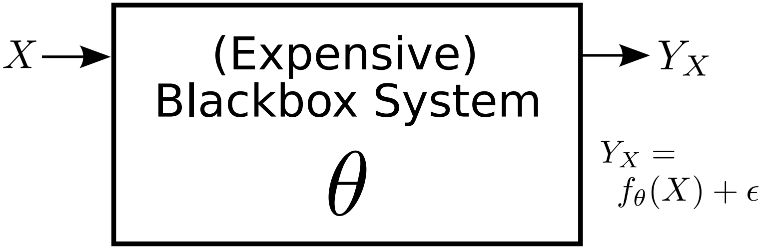

Goal: Find $\mathrm{argmax}_x f_\theta(x)$ in as few experiments.

Algorithms: UCB (Srinivas et al 2010, Auer 2002), EI (Jones et al 1998).

Adaptive Goal Oriented Design of Experiments

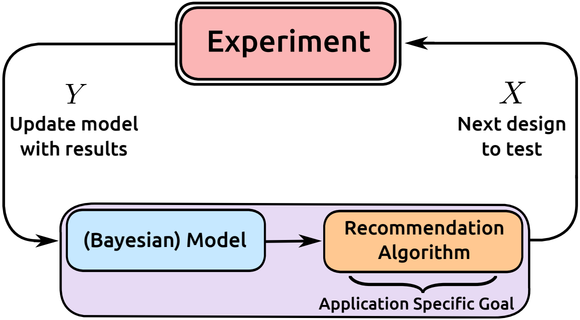

Adaptive Goal Oriented Design of Experiments

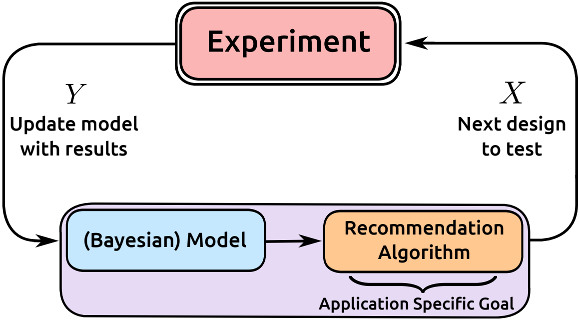

- Blackbox Optimisation
- Active Learning
- Active Quadrature (Osborne et al. 2012)
- Active Level Set Estimation (Gotovos et al. '13)
- Active Search (Ma et al. '17)
- Active Posterior Estimation (Kandasamy et al. '15)

Adaptive Goal Oriented Design of Experiments

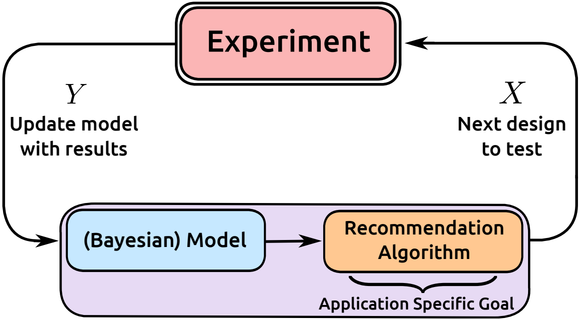

- Blackbox Optimisation
- Active Learning
- Active Quadrature (Osborne et al. 2012)
- Active Level Set Estimation (Gotovos et al. '13)
- Active Search (Ma et al. '17)
- Active Posterior Estimation (Kandasamy et al. '15)

Issues:

- New goal/setting $\implies$ New algorithm?
- Algorithms tend to depend on the model and vice versa.

Adaptive Goal Oriented Design of Experiments

1. **System:**
   - An *unknown* parameter θ completely specifies the system.
   - A prior $\mathbb{P}(\theta)$ and a likelihood $\mathbb{P}(Y|X, \theta)$.

Adaptive Goal Oriented Design of Experiments

1. **System:**
   - An *unknown* parameter θ completely specifies the system.
   - A prior $\mathbb{P}(\theta)$ and a likelihood $\mathbb{P}(Y|X,\theta)$.
2. **Goal:**
   - Collect data $D_n = \{(x_t, y_{x_t})\}_{t=1}^{n}$ to maximise a user specified reward function $\lambda(\theta, D_n)$.

Algorithm: Myopic Posterior Sampling (MPS)

Inspired by Posterior (Thompson) Sampling (Thompson 1933).

At each time step, myopically choose action by assuming that a posterior sample $\theta' \sim \mathbb{P}(\theta|\text{past-experiments})$ is the true parameter.

Algorithm: Myopic Posterior Sampling (MPS)

Inspired by Posterior (Thompson) Sampling (Thompson 1933).

At each time step, myopically choose action by assuming that a posterior sample $\theta' \sim \mathbb{P}(\theta|\text{past-experiments})$ is the true parameter.

Only requires that we can sample from the posterior.

- Many probabilistic programming tools available today.

Theory

Theorem (Informal): Under certain conditions, MPS is competitive with a *globally* optimal oracle that *knows* θ.

Proof ideas from adaptive submodularity and bandits.

Theory

Theorem (Informal): Under certain conditions, MPS is competitive with a *globally* optimal oracle that *knows* θ.

Proof ideas from adaptive submodularity and bandits.

Prior work: With adaptive submodularity, myopic *planning* algorithms are good when the reward is known a priori.

Theory

Theorem (Informal): Under certain conditions, MPS is competitive with a *globally* optimal oracle that *knows* θ.

Proof ideas from adaptive submodularity and bandits.

Prior work: With adaptive submodularity, myopic *planning* algorithms are good when the reward is known a priori.

This work:

- $\lambda(\theta, D_n)$: reward not known a priori.
- A myopic *learning+planning* algorithm is good in adaptive submodular environments.

Experiments

Active Learning

Synthetic Example

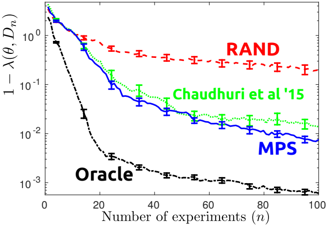

Active Level Set Estimation

Luminous Red Galaxies

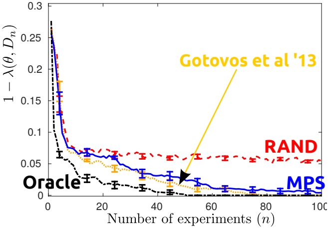

Active Posterior Estimation

Type Ia Supernova

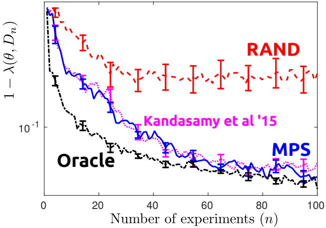

Application Specific Goal

Electrolyte Design

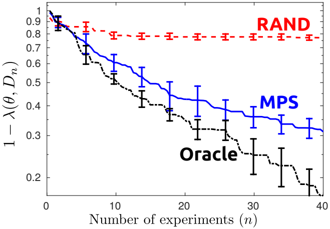

Willie

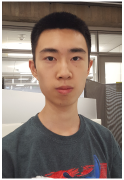

Reed

Akshay

Jeff

Barnabas

Carnegie Mellon University

Auton Lab

Microsoft Research

Code: `github.com/kirthevasank/mps`

Poster: #262