Optimistic Optimization

Lucian Buşoniu

26 May 2014

References

These methods were published in:

1. Rémi Munos, *Optimistic optimization of deterministic functions without the knowledge of its smoothness*, Advances in Neural Information Processing Systems 2011.
2. Rémi Munos, *The optimistic principle applied to games, optimization and planning: Towards foundations of Monte-Carlo Tree Search*, Foundations and Trends in Machine Learning 7, 2014, pp. 1–130.

1 is original reference, 2 is an extensive survey including applications to control

Optimization problem

$$\max_{x \in X} f(x)$$

Assumption 1

Function $f : X \to \mathbb{R}$ is Lipschitz-continuous with respect to a semimetric ℓ:

$$\left| f(x) - f(x') \right| \leq \ell(x, x')$$

Definition (Semimetric)

A function $\ell : X \times X \to \mathbb{R}$ satisfying:

- $\ell(x, x') \geq 0$
- $\ell(x, x') = \ell(x', x)$
- $\ell(x, x') = 0$ if and only if $x = x'$

Intuitively: a notion of distance

Motivation

No method with guaranteed performance for any function

Especially when the metric ℓ is unknown

DOO idea

- Explore the space X iteratively
- Always expand **optimistic** set, with largest upper bound: $b(X_i) = f(x_i) + \delta(X_i)$, diam. $\delta(X_i) = \sup_{x,x' \in X_i} \ell(x, x')$
- Until n expansions exhausted

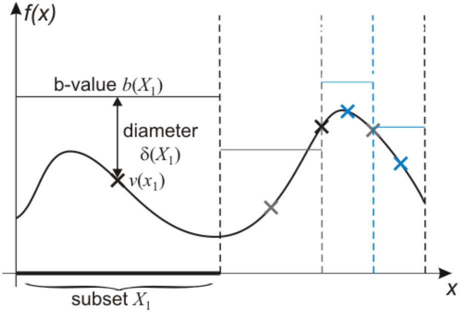

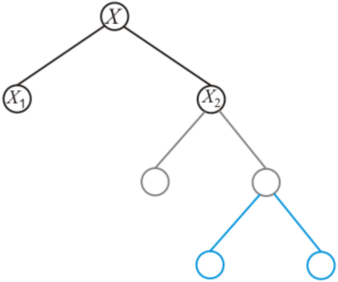

Partitioning

- In general, a hierarchical partitioning rule must be defined
- Set $X_{0,1} = X$ at depth 0 split into $X_{1,1}, \ldots, X_{1,K}$ at depth 1
- Each set $X_{d,i}$ at depth d split into K subsets at depth $d+1$

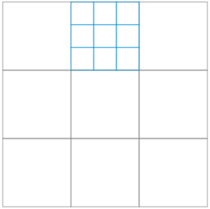

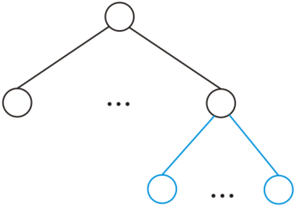

Partitioning requirements

Assumption 2

The sets $X_{d,i}$ in the hierarchical partitioning must:

a) Shrink with the depth:
$\delta(X_{d,i}) \leq \delta_d$ for any set i at d; δ_d decreases with d

b) Be well-shaped:
each $\delta(X_{d,i})$ contains a ball in the semimetric ℓ having radius proportional to δ_d, $B(x_{d,i}, \nu\delta_d)$

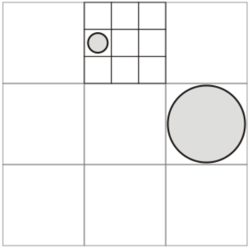

DOO algorithm

initialize tree with root $X_{0,1} = X$
for $t = 1$ to n **do**
  $X^{\dagger}_{d,i} \leftarrow \arg\max_{X_{d,i} \in \text{ leaves}} b(X_{d,i})$
  expand $X^{\dagger}_{d,i}$ (partition the set), adding children to tree
end for
output best sample $\widehat{x}^* = \arg\max_{x_{d,i} \in \text{tree}} f(x_{d,i})$

(Munos, 2011)

Examples

- Quadratic function
- Rosenbrock banana function

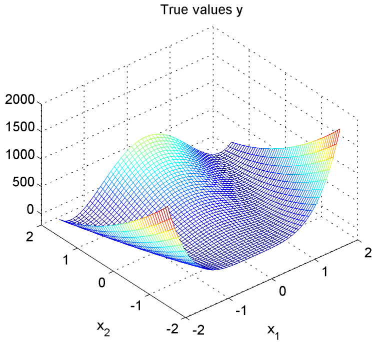

An easy near-optimality guarantee

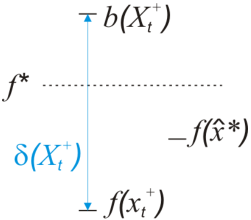

- Denote the expanded set at iteration t by $X_t^\dagger$
- $b(X_t^\dagger) \geq f^*$, otherwise it wouldn't have been selected
- $f(x_t^\dagger) \leq f^*$ by definition
- $f(\widehat{x}^*) \geq f(x_t^\dagger)$ because $\widehat{x}^*$ maximizes f on the tree
- So $f^* - f(\widehat{x}^*) \leq \delta(X_t^\dagger)$ at any t, and therefore $\leq \delta_{d^*}$, where d^* the deepest expanded depth

Near-optimality dimension

Definition (Near-optimality dimension)

Smallest β so that the near-optimal sets:

$$X_\varepsilon = \{x \in X \mid f^* - f(x) \leq \varepsilon\}$$

can be covered by (on the order of) $\varepsilon^{-\beta}$ balls of radius ε in the semimetric ℓ

β measures how closely ℓ captures the smoothness of f

DOO near-optimality

Theorem

Consider a partition with exponentially decreasing sets, $\delta_d = \gamma^d$, $\gamma < 1$. Then the solution returned by DOO satisfies:

$$f^* - f(\widehat{x}^*) \approx \begin{cases} n^{-1/\beta} & \text{if } \beta > 0 \\ \gamma^{cn} & \text{if } \beta = 0 \end{cases}$$

Example: zero dimension

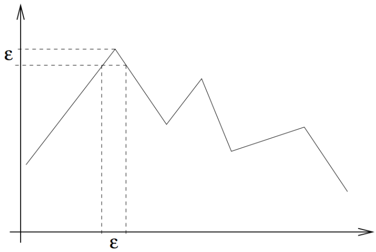

- Take $f(x^*) - f(x) \approx |x^* - x|$ and $\ell(x, x') = |x - x'|$
- $X_\varepsilon =$ an interval of length ε, which is also an ℓ-ball of size ε
- So it takes a constant = ε^0 number of balls to cover X_ε, and $\beta = 0$

(taken from Munos)

Example: positive dimension

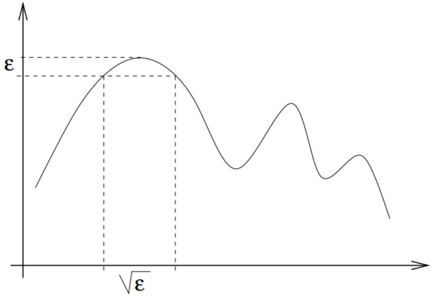

- If $f(x^*) - f(x) \approx |x^* - x|^2$, X_ε is an interval of length $\sqrt{\varepsilon}$
- When $\ell(x, x') = |x - x'|^2$, a ℓ-ball of size ε is also an interval of length $\sqrt{\varepsilon}$, and $\beta = 0$
- When $\ell(x, x') = |x - x'|$, a ℓ-ball of size ε is an interval of length ε, so it takes $\varepsilon/\sqrt(\varepsilon) = \varepsilon^{-1/2}$ balls to cover X_ε, and $\beta = 1/2$

(taken from Munos)

Example: semimetric mismatch

Influence of semimetric (mis)match for a quadratic function

SOO idea

- What if ℓ / δ unknown? (i.e., smoothness of f unknown)
- Expand **all potentially optimistic sets** $X_{d,i}$, for which:
  $f(x_{d,i}) \geq f(x_{d',j})$ for all leaves j at smaller depths $d' \leq d$

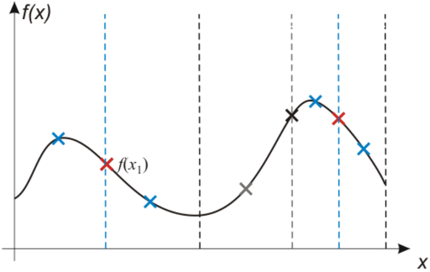

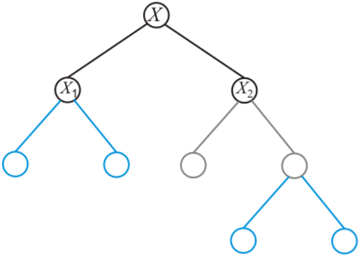

SOO algorithm

initialize tree with root $X_{0,1} = X$
repeatat each iteration $t = 1, 2, \ldots$
  for $d = 0, \ldots, \min\{\text{current tree depth}, d_{\max}(t)\}$ **do**
    $X^{\dagger}_{d,i} \leftarrow \arg\max_{X_{d,i} \in \text{ leaves at } d} f(x_{d,i})$
    if $f(x^{\dagger}_{d,i}) \geq f(x_{d',j}) \; \forall$ leaves j at $d' \leq d$ **then**
      expand $X^{\dagger}_{d,i}$
    end if
  end for
until n expansions performed
output best sample $\widehat{x}^* = \arg\max_{x_{d,i} \in \text{tree}} f(x_{d,i})$

(Munos, 2011)

Examples

- Quadratic function
- Rosenbrock banana function

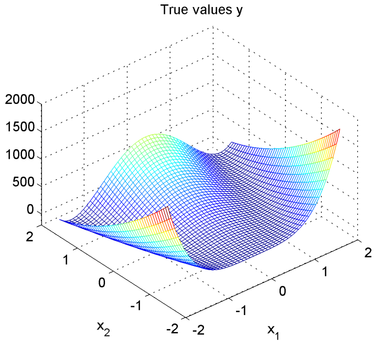

SOO near-optimality

Theorem

Consider a partition with exponentially decreasing sets, $\delta_d = \gamma^d$, $\gamma < 1$. Take $d_{\max}(t) = \sqrt{t}$, then the solution returned by SOO satisfies:

$$f^* - f(\widehat{x}^*) \approx \begin{cases} n^{-\frac{1}{2\beta}} & \text{if } \beta > 0 \\ \gamma^{c'n} & \text{if } \beta = 0 \end{cases}$$

1. Problem & motivation
2. DOO: Deterministic optimistic optimization
3. SOO: Simultaneous optimistic optimization
4. Application: Multiagent consensus

Consensus in nonlinear multiagent systems

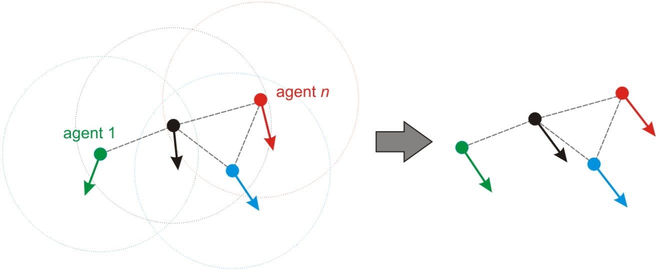

- Agents with nonlinear dynamics $x_{i,k+1} = f_i(x_{i,k}, u_{i,k})$
- Consensus problem: agents must reach agreement on (some) state variables
- Communication on an incomplete graph
- **Challenge:** No solution for general f

OO for consensus

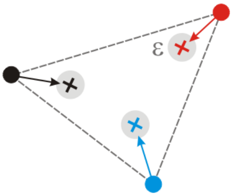

1. Design target states with a classical consensus method
2. Use DOO or SOO to optimize action sequences in order to reach within ε of target states

- **Consensus guaranteed** under conditions on f
- Tradeoff: length of action sequence must be known and small

Consensus of multiple robot arms