An obvious optimization method: blackbox query

Samuel Vaiter

Created: 2023-11-08.
Last update: 2023-11-08.
Status: inprogress.

Blackbox query

$$\min_{x \in [-1,1]^d} f(x)$$

Goal: find $\hat{x} \approx x^\star$ up to a precision $\varepsilon > 0$.

Discretization:

$$G_\varepsilon = \{k\varepsilon \ : \ k \in \{-\lfloor \varepsilon^{-1} \rfloor, \ldots, \lfloor \varepsilon^{-1} \rfloor\}\}$$

Approximate solution:

$$\hat{x} = \operatorname*{argmin}_{x \in G_\varepsilon} f(x)$$

Error (for L-Lipschitz function):

$$f(\hat{x}) - f(x^\star) \leq L\varepsilon$$

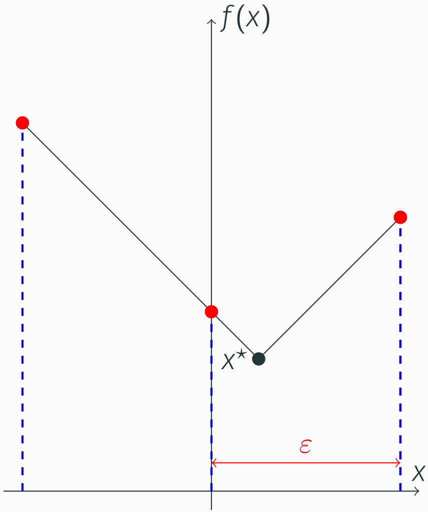

Blackbox query

$$\min_{x\in[-1,1]^d} f(x)$$

Goal: find $\hat{x} \approx x^\star$ up to a precision $\varepsilon > 0$.

Discretization:

$$G_\varepsilon = \left\{k\varepsilon \ : \ k \in \{-\lfloor \varepsilon^{-1} \rfloor, \dots, \lfloor \varepsilon^{-1} \rfloor\}\right\}$$

Approximate solution:

$$\hat{x} = \operatorname*{argmin}_{x\in G_\varepsilon} f(x)$$

Error (for L-Lipschitz function):

$$f(\hat{x}) - f(x^\star) \leq L\varepsilon$$

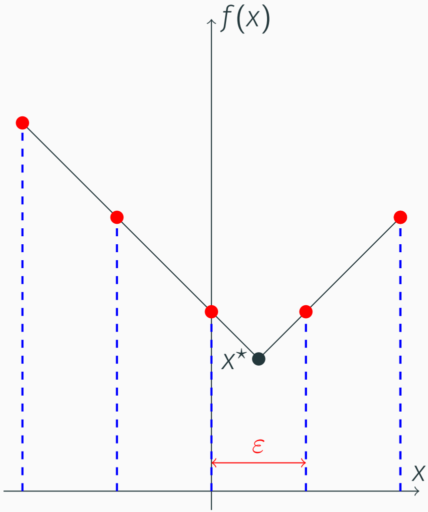

Blackbox query

$$\min_{x \in [-1,1]^d} f(x)$$

Goal: find $\hat{x} \approx x^\star$ up to a precision $\varepsilon > 0$.

Discretization:

$$G_\varepsilon = \{k\varepsilon \ : \ k \in \{-\lfloor \varepsilon^{-1} \rfloor, \dots, \lfloor \varepsilon^{-1} \rfloor\}\}$$

Approximate solution:

$$\hat{x} = \operatorname{argmin}_{x \in G_\varepsilon} f(x)$$

Error (for L-Lipschitz function):

$$f(\hat{x}) - f(x^\star) \leq L\varepsilon$$

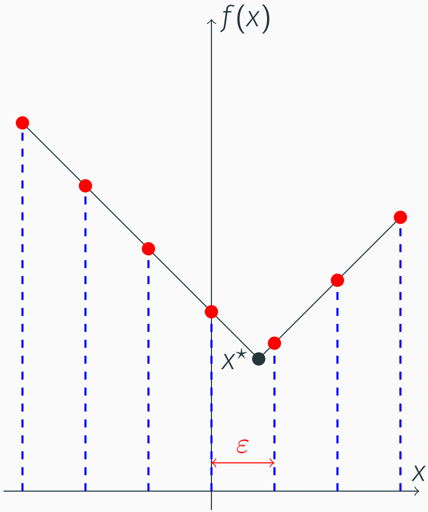

Blackbox query: what's the catch? Curse of dimensionality!

Size of G_ε to get an ε precision

- *1D*: $O(\lfloor \varepsilon^{-1} \rfloor)$ calls to f

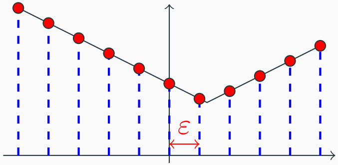

Blackbox query: what's the catch? Curse of dimensionality!

Size of G_ε to get an ε precision

- *1D*: $O(\lfloor \varepsilon^{-1} \rfloor)$ calls to f
- *2D*: $O(\lfloor \varepsilon^{-1} \rfloor)^2$ calls to f

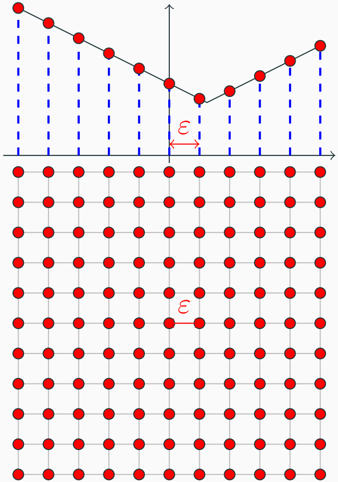

Blackbox query: what's the catch? Curse of dimensionality!

Size of G_ε to get an ε precision

- *1D*: $O(\lfloor \varepsilon^{-1} \rfloor)$ calls to f
- *2D*: $O(\lfloor \varepsilon^{-1} \rfloor)^2$ calls to f
- *3D*: $O(\lfloor \varepsilon^{-1} \rfloor)^3$ calls to f

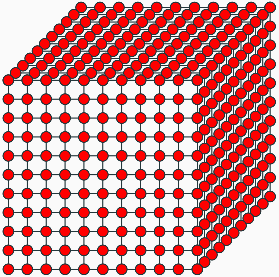

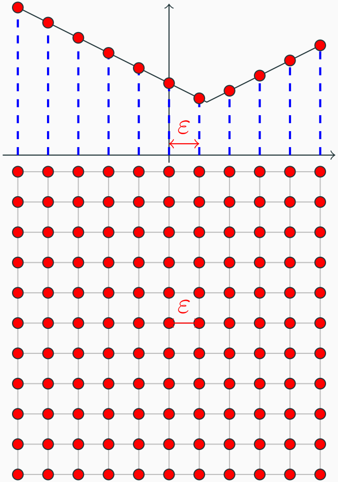

Blackbox query: what's the catch? Curse of dimensionality!

Size of G_ε to get an ε precision

- *1D*: $O(\lfloor \varepsilon^{-1} \rfloor)$ calls to f
- *2D*: $O(\lfloor \varepsilon^{-1} \rfloor)^2$ calls to f
- *3D*: $O(\lfloor \varepsilon^{-1} \rfloor)^3$ calls to f

dimension d: $O(\lfloor \varepsilon^{-1} \rfloor)^d$ calls to f