Deep Optimal Stopping

Sebastian Becker
ZENAI

Patrick Cheridito
RiskLab, ETH Zurich

Arnulf Jentzen
Universität Münster

Vienna, November 2019

The Problem

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau),$$

The Problem

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau),$$

where

- $(X_n)_{n=0}^N$ is a d-dimensional Markov process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$

The Problem

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau),$$

where

- $(X_n)_{n=0}^N$ is a d-dimensional Markov process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$
- $g\colon \{0, 1, \ldots, N\} \times \mathbb{R}^d \to \mathbb{R}$ a measurable function such that

$$\mathbb{E}|g(n, X_n)| < \infty \quad \text{for all } n = 0, \ldots, N$$

The Problem

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau),$$

where

- $(X_n)_{n=0}^N$ is a d-dimensional Markov process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$
- $g\colon \{0, 1, \ldots, N\} \times \mathbb{R}^d \to \mathbb{R}$ a measurable function such that

$$\mathbb{E}|g(n, X_n)| < \infty \quad \text{for all } n = 0, \ldots, N$$

- $\mathcal{T}$ is the set of all X-stopping times τ

  that is, $\{\tau = n\} \in \sigma(X_0, \ldots, X_n)$ for all $n = 0, 1, \ldots, N$

About the assumptions

- **Discrete time**

About the assumptions

- **Discrete time**
  - Many problems are already in discrete time

About the assumptions

- **Discrete time**
  - Many problems are already in discrete time
  - Most relevant continuous-time problems can be approximated by time-discretized versions

About the assumptions

- **Discrete time**
  - Many problems are already in discrete time
  - Most relevant continuous-time problems can be approximated by time-discretized versions

- **Markov assumption**

About the assumptions

- **Discrete time**
  - Many problems are already in discrete time
  - Most relevant continuous-time problems can be approximated by time-discretized versions

- **Markov assumption**
  - Every discrete-time process can be made Markov by including all relevant information in the current state ... by increasing the dimension of $(X_n)_{n=0}^N$

Examples

1. **Bermudan max-call options**

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$
   - a risk-free interest rate $r \in \mathbb{R}$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \dots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$
   - a risk-free interest rate $r \in \mathbb{R}$
   - dividend yields $\delta_i \in [0, \infty)$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$
   - a risk-free interest rate $r \in \mathbb{R}$
   - dividend yields $\delta_i \in [0, \infty)$
   - volatilities $\sigma_i \in (0, \infty)$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$
   - a risk-free interest rate $r \in \mathbb{R}$
   - dividend yields $\delta_i \in [0, \infty)$
   - volatilities $\sigma_i \in (0, \infty)$
   - and a d-dimensional Brownian motion W with constant correlation ρ_{ij} between increments of different components W^i and W^j

Examples

1. **Bermudan max-call options**

Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

for

- initial values $s_0^i \in (0, \infty)$
- a risk-free interest rate $r \in \mathbb{R}$
- dividend yields $\delta_i \in [0, \infty)$
- volatilities $\sigma_i \in (0, \infty)$
- and a d-dimensional Brownian motion W with constant correlation ρ_{ij} between increments of different components W^i and W^j

A Bermudan max-call option has time-t payoff $\left(\max_{1 \le i \le d} S_t^i - K\right)^+$

Examples

1. **Bermudan max-call options**

Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

for

- initial values $s_0^i \in (0, \infty)$
- a risk-free interest rate $r \in \mathbb{R}$
- dividend yields $\delta_i \in [0, \infty)$
- volatilities $\sigma_i \in (0, \infty)$
- and a d-dimensional Brownian motion W with constant correlation ρ_{ij} between increments of different components W^i and W^j

A Bermudan max-call option has time-t payoff $\left(\max_{1 \leq i \leq d} S_t^i - K\right)^+$ and can be exercised at one of finitely many times $0 = t_0 < t_1 = \frac{T}{N} < t_2 = \frac{2T}{N} < \cdots < t_N = T$

Examples

1. **Bermudan max-call options**

   Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

   for

   - initial values $s_0^i \in (0, \infty)$
   - a risk-free interest rate $r \in \mathbb{R}$
   - dividend yields $\delta_i \in [0, \infty)$
   - volatilities $\sigma_i \in (0, \infty)$
   - and a d-dimensional Brownian motion W with constant correlation ρ_{ij} between increments of different components W^i and W^j

   A Bermudan max-call option has time-t payoff $\left(\max_{1 \le i \le d} S_t^i - K\right)^+$ and can be exercised at one of finitely many times $0 = t_0 < t_1 = \frac{T}{N} < t_2 = \frac{2T}{N} < \cdots < t_N = T$

   Price: $$\sup_{\tau \in \{t_0, t_1, \ldots, T\}} \mathbb{E}\left[e^{-r\tau}\left(\max_{1 \le i \le d} S_\tau^i - K\right)^+\right]$$

Examples

1. **Bermudan max-call options**

Consider d assets with prices evolving according to a multi-dimensional Black–Scholes model

$$S_t^i = s_0^i \exp\left([r - \delta_i - \sigma_i^2/2]t + \sigma_i W_t^i\right), \quad i = 1, 2, \ldots, d,$$

for

- initial values $s_0^i \in (0, \infty)$
- a risk-free interest rate $r \in \mathbb{R}$
- dividend yields $\delta_i \in [0, \infty)$
- volatilities $\sigma_i \in (0, \infty)$
- and a d-dimensional Brownian motion W with constant correlation ρ_{ij} between increments of different components W^i and W^j

A Bermudan max-call option has time-t payoff $\left(\max_{1 \le i \le d} S_t^i - K\right)^+$ and can be exercised at one of finitely many times $0 = t_0 < t_1 = \frac{T}{N} < t_2 = \frac{2T}{N} < \cdots < t_N = T$

Price: $$\sup_{\tau \in \{t_0, t_1, \ldots, T\}} \mathbb{E}\left[e^{-r\tau}\left(\max_{1 \le i \le d} S_\tau^i - K\right)^+\right] = \sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau)$$

This problem has been studied for $d = 2, 3, 5$ (among others) by

- Longstaff and Schwartz (2001)
- Rogers (2002)
- García (2003)
- Boyle, Kolkiewicz and Tan (2003)
- Haugh and Kogan (2004)
- Broadie and Glasserman (2004)
- Andersen and Broadie (2004)
- Broadie and Cao (2008)
- Berridge and Schumacher (2008)
- Belomestny (2011, 2013)
- Jain and Oosterlee (2015)
- Lelong (2016)

Our price estimates

for $s_0^i = 100$, $\sigma_i = 20\%$, $r = 5\%$, $\delta = 10\%$, $\rho_{ij} = 0$, $K = 100$, $T = 3$, $N = 9$:

# assets	Point Est.	Comp. Time	95% Conf. Int.	Bin. Tree
2	13.899	$28.7s$	$[13.880, 13.910]$	13.902
3	18.690	$28.9s$	$[18.673, 18.699]$	18.69

Our price estimates

for $s_0^i = 100$, $\sigma_i = 20\%$, $r = 5\%$, $\delta = 10\%$, $\rho_{ij} = 0$, $K = 100$, $T = 3$, $N = 9$:

# assets	Point Est.	Comp. Time	95% Conf. Int.	Bin. Tree	Broadie–Cao 95% Conf. Int.
2	13.899	28.7s	[13.880, 13.910]	13.902	
3	18.690	28.9s	[18.673, 18.699]	18.69	
5	26.159	28.1s	[26.138, 26.174]		[26.115, 26.164]

Our price estimates

for $s_0^i = 100$, $\sigma_i = 20\%$, $r = 5\%$, $\delta = 10\%$, $\rho_{ij} = 0$, $K = 100$, $T = 3$, $N = 9$:

# Assets	Point Est.	Comp. Time	95% Conf. Int.	Bin. Tree	Broadie–Cao 95% Conf. Int.
2	13.899	$28.7s$	[13.880, 13.910]	13.902	
3	18.690	$28.9s$	[18.673, 18.699]	18.69	
5	26.159	$28.1s$	[26.138, 26.174]		[26.115, 26.164]
10	38.337	$30.5s$	[38.300, 38.367]		
20	51.668	$37.5s$	[51.549, 51.803]		
30	59.659	$45.5s$	[59.476, 59.872]		
50	69.736	$59.1s$	[69.560, 69.945]		
100	83.584	$95.9s$	[83.357, 83.862]		
200	97.612	$170.1s$	[97.381, 97.889]		
500	116.425	$493.5s$	[116.210, 116.685]		

2. Optimally stopping a fractional Brownian motion

2. **Optimally stopping a fractional Brownian motion**

   A fractional Brownian motion with Hurst parameter $H \in (0,1]$ is a continuous centered Gaussian process $(W_t^H)_{t\geq 0}$ with covariance structure

$$\mathrm{Cov}(W_t^H, W_s^H) = \frac{1}{2}\left(t^{2H} + s^{2H} - |t-s|^{2H}\right)$$

2. **Optimally stopping a fractional Brownian motion**

A fractional Brownian motion with Hurst parameter $H \in (0,1]$ is a continuous centered Gaussian process $(W_t^H)_{t\geq 0}$ with covariance structure

$$\mathrm{Cov}(W_t^H, W_s^H) = \frac{1}{2}\left(t^{2H} + s^{2H} - |t-s|^{2H}\right)$$

- For $H = 1/2$, W^H is a Brownian motion

2. **Optimally stopping a fractional Brownian motion**

A fractional Brownian motion with Hurst parameter $H \in (0, 1]$ is a continuous centered Gaussian process $(W_t^H)_{t \geq 0}$ with covariance structure

$$\mathrm{Cov}(W_t^H, W_s^H) = \frac{1}{2}\left(t^{2H} + s^{2H} - |t - s|^{2H}\right)$$

- For $H = 1/2$, W^H is a Brownian motion
- For $H > 1/2$, W^H has positively correlated increments

2 Optimally stopping a fractional Brownian motion

A fractional Brownian motion with Hurst parameter $H \in (0, 1]$ is a continuous centered Gaussian process $(W_t^H)_{t \geq 0}$ with covariance structure

$$\mathrm{Cov}(W_t^H, W_s^H) = \frac{1}{2}\left(t^{2H} + s^{2H} - |t-s|^{2H}\right)$$

- For $H = 1/2$, W^H is a Brownian motion
- For $H > 1/2$, W^H has positively correlated increments
- For $H < 1/2$, W^H has negatively correlated increments

2 Optimally stopping a fractional Brownian motion

A fractional Brownian motion with Hurst parameter $H \in (0,1]$ is a continuous centered Gaussian process $(W_t^H)_{t\geq 0}$ with covariance structure

$$\mathrm{Cov}(W_t^H, W_s^H) = \frac{1}{2}\left(t^{2H} + s^{2H} - |t-s|^{2H}\right)$$

- For $H = 1/2$, W^H is a Brownian motion
- For $H > 1/2$, W^H has positively correlated increments
- For $H < 1/2$, W^H has negatively correlated increments

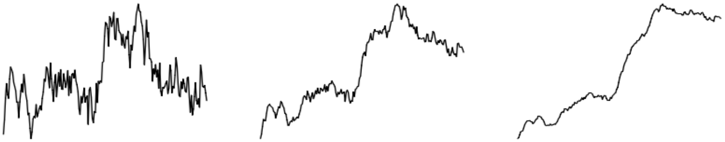

$H = 0.1$ $H = 0.5$ $H = 0.8$

Problem: $$\sup_{0\le\tau\le 1} \mathbb{E}\, W_\tau^H \qquad (*)$$

Problem: $$\sup_{0\le\tau\le 1} \mathbb{E}\, W^H_\tau \qquad (*)$$

- denote $t_n = n/100$, $n = 0, 1, 2, \ldots, 100$

Problem: $$\sup_{0\le\tau\le 1} \mathbb{E}\, W^H_\tau \qquad (*)$$

- denote $t_n = n/100$, $n = 0, 1, 2, \ldots, 100$
- introduce the 100-dimensional Markov process $(X_n)_{n=0}^{100}$ given by

$$\begin{aligned}
X_0 &= (0, 0, \ldots, 0)\\
X_1 &= (W^H_{t_1}, 0, \ldots, 0)\\
X_2 &= (W^H_{t_2}, W^H_{t_1}, 0, \ldots, 0)\\
&\vdots\\
X_{100} &= (W^H_{t_{100}}, W^H_{t_{99}}, \ldots, W^H_{t_1}).
\end{aligned}$$

Problem: $$\sup_{0\le\tau\le 1} \mathbb{E}\, W^H_\tau \qquad (*)$$

- denote $t_n = n/100$, $n = 0, 1, 2, \ldots, 100$
- introduce the 100-dimensional Markov process $(X_n)_{n=0}^{100}$ given by

$$\begin{aligned} X_0 &= (0, 0, \ldots, 0) \\ X_1 &= (W^H_{t_1}, 0, \ldots, 0) \\ X_2 &= (W^H_{t_2}, W^H_{t_1}, 0, \ldots, 0) \\ &\vdots \\ X_{100} &= (W^H_{t_{100}}, W^H_{t_{99}}, \ldots, W^H_{t_1}). \end{aligned}$$

The discretized stopping problem

$$\sup_{\tau\in\mathcal{T}} \mathbb{E}\, g(X_\tau) \quad \text{for } g(x^1, \ldots, x^{100}) = x^1,$$

approximates the continuous-time problem $(*)$ from below

Results of Kulikov and Gusyatnikov (2016)
(based on heuristic stopping rules)

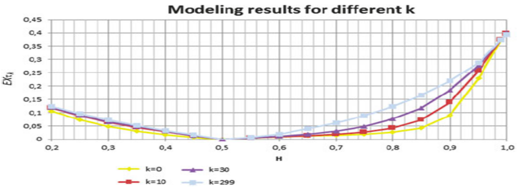

Results of Kulikov and Gusyatnikov (2016)
(based on heuristic stopping rules)

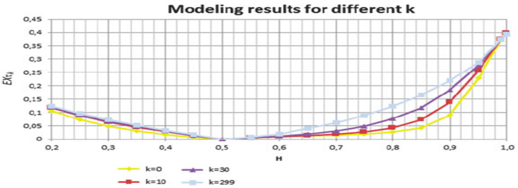

Our results

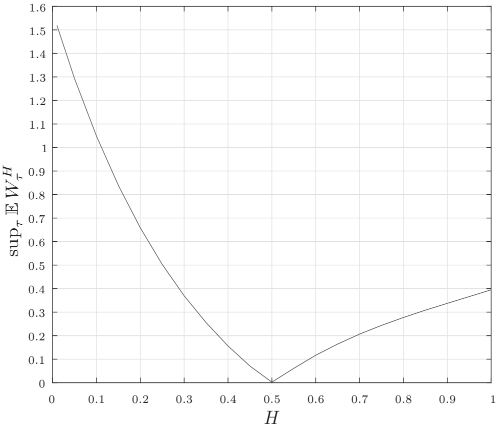

Computing a candidate optimal stopping time

- **Introduce the sequence of auxiliary stopping problems**

$$V_n = \sup_{\tau \in \mathcal{T}_n} \mathbb{E}\, g(\tau, X_\tau), \quad n = 0, 1, \ldots, N,$$

where $\mathcal{T}_n$ is the set of all stopping times $n \leq \tau \leq N$

Computing a candidate optimal stopping time

- **Introduce the sequence of auxiliary stopping problems**

$$V_n = \sup_{\tau \in \mathcal{T}_n} \mathbb{E}\, g(\tau, X_\tau), \quad n = 0, 1, \ldots, N,$$

where $\mathcal{T}_n$ is the set of all stopping times $n \leq \tau \leq N$

- **Stopping times and stopping decisions**

Let $f_n, f_{n+1}, \ldots, f_N : \mathbb{R}^d \to \{0, 1\}$ be measurable functions such that $f_N \equiv 1$. Then

$$\tau_n = \sum_{m=n}^{N} m f_m(X_m) \prod_{j=n}^{m-1} (1 - f_j(X_j)) \qquad \text{with} \quad \prod_{j=n}^{n-1} (1 - f_j(X_j)) := 1$$

is a stopping time in $\mathcal{T}_n$

Theorem

For a given $n \in \{0, 1, \ldots, N-1\}$, let τ_{n+1} be a stopping time in $\mathcal{T}_{n+1}$ of the form

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f_m(X_m) \prod_{j=n+1}^{m-1} (1 - f_j(X_j)),$$

for measurable functions $f_{n+1}, \ldots, f_N : \mathbb{R}^d \to \{0, 1\}$ with $f_N \equiv 1$.

Theorem

For a given $n \in \{0, 1, \ldots, N-1\}$, let τ_{n+1} be a stopping time in $\mathcal{T}_{n+1}$ of the form

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f_m(X_m) \prod_{j=n+1}^{m-1} (1 - f_j(X_j)),$$

for measurable functions $f_{n+1}, \ldots, f_N : \mathbb{R}^d \to \{0, 1\}$ with $f_N \equiv 1$.

Then there exists a measurable function $f_n : \mathbb{R}^d \to \{0, 1\}$ such that the stopping time

$$\tau_n = n f_n(X_n) + \tau_{n+1}(1 - f_n(X_n)) = \sum_{m=n}^{N} m f_m(X_m) \prod_{j=n}^{m-1} (1 - f_j(X_j))$$

satisfies

$$\mathbb{E}\, g(\tau_n, X_{\tau_n}) \geq V_n - \left(V_{n+1} - \mathbb{E}\, g(\tau_{n+1}, X_{\tau_{n+1}})\right)$$

Theorem

For a given $n \in \{0, 1, \dots, N-1\}$, let τ_{n+1}be a stopping time in $\mathcal{T}_{n+1}$ of the form

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f_m(X_m) \prod_{j=n+1}^{m-1} (1 - f_j(X_j)),$$

for measurable functions $f_{n+1}, \dots, f_N : \mathbb{R}^d \to \{0, 1\}$ with $f_N \equiv 1$.

Then there exists a measurable function $f_n : \mathbb{R}^d \to \{0, 1\}$ such that the stopping time

$$\tau_n = n f_n(X_n) + \tau_{n+1}(1 - f_n(X_n)) = \sum_{m=n}^{N} m f_m(X_m) \prod_{j=n}^{m-1} (1 - f_j(X_j))$$

satisfies

$$\mathbb{E}\, g(\tau_n, X_{\tau_n}) \geq V_n - \left(V_{n+1} - \mathbb{E}\, g(\tau_{n+1}, X_{\tau_{n+1}})\right)$$

Proof: Compare $g(n, X_n)$ to $\mathbb{E}\left[g(\tau_{n+1}, X_{\tau_{n+1}}) \mid X_0, X_1, \dots, X_n\right] = \mathbb{E}\left[g(\tau_{n+1}, X_{\tau_{n+1}}) \mid X_n\right]$

Neural network approximation

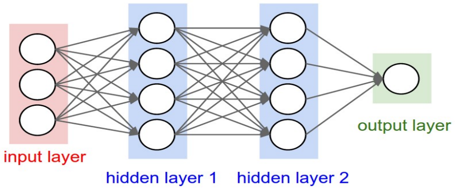

Neural network approximation

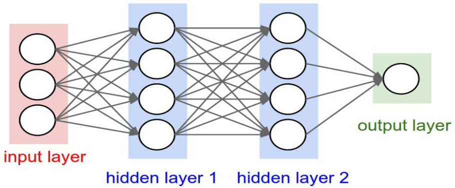

Idea Recursively approximate f_n by a neural network $f^\theta\colon \mathbb{R}^d \to \{0,1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

Neural network approximation

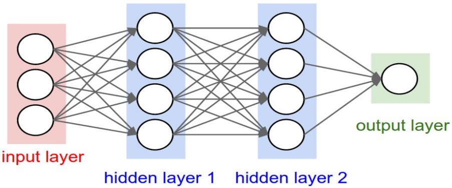

Idea Recursively approximate f_n by a neural network $f^\theta\colon \mathbb{R}^d \to \{0,1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

where

- q_1 and q_2 are positive integers specifying the number of nodes in the two hidden layers,

Neural network approximation

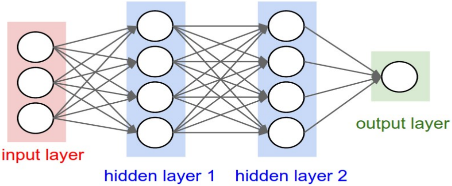

Idea Recursively approximate f_n by a neural network $f^\theta : \mathbb{R}^d \to \{0, 1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

where

- q_1 and q_2 are positive integers specifying the number of nodes in the two hidden layers,
- $a_1^\theta : \mathbb{R}^d \to \mathbb{R}^{q_1}$, $a_2^\theta : \mathbb{R}^{q_1} \to \mathbb{R}^{q_2}$ and $a_3^\theta : \mathbb{R}^{q_2} \to \mathbb{R}$ are affine functions given by

$$a_i^\theta(x) = A_i x + b_i, \; i = 1, 2, 3,$$

Neural network approximation

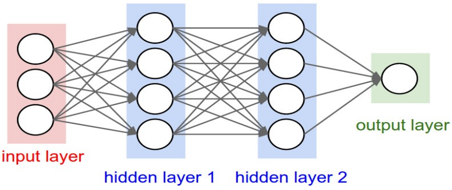

Idea Recursively approximate f_n by a neural network $f^\theta \colon \mathbb{R}^d \to \{0, 1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

where

- q_1 and q_2 are positive integers specifying the number of nodes in the two hidden layers,
- $a_1^\theta \colon \mathbb{R}^d \to \mathbb{R}^{q_1}$, $a_2^\theta \colon \mathbb{R}^{q_1} \to \mathbb{R}^{q_2}$ and $a_3^\theta \colon \mathbb{R}^{q_2} \to \mathbb{R}$ are affine functions given by
$$a_i^\theta(x) = A_i x + b_i,\ i = 1, 2, 3,$$
- for $j \in \mathbb{N}$, $\varphi_j \colon \mathbb{R}^j \to \mathbb{R}^j$ is the component-wise ReLU activation function given by $\varphi_j(x_1, \ldots, x_j) = (x_1^+, \ldots, x_j^+)$

Neural network approximation

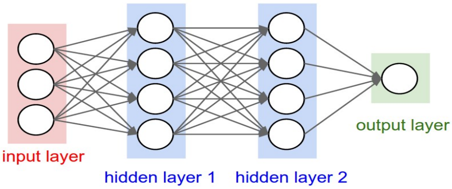

Idea Recursively approximate f_n by a neural network $f^\theta\colon \mathbb{R}^d \to \{0,1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

where

- q_1 and q_2 are positive integers specifying the number of nodes in the two hidden layers,
- $a_1^\theta\colon \mathbb{R}^d \to \mathbb{R}^{q_1}$, $a_2^\theta\colon \mathbb{R}^{q_1} \to \mathbb{R}^{q_2}$ and $a_3^\theta\colon \mathbb{R}^{q_2} \to \mathbb{R}$ are affine functions given by
$$a_i^\theta(x) = A_i x + b_i,\ i = 1,2,3,$$
- for $j \in \mathbb{N}$, $\varphi_j\colon \mathbb{R}^j \to \mathbb{R}^j$ is the component-wise ReLU activation function given by $\varphi_j(x_1,\ldots,x_j) = (x_1^+,\ldots,x_j^+)$

The components of θ consist of the entries of A_i and b_i, $i = 1,2,3$

Neural network approximation

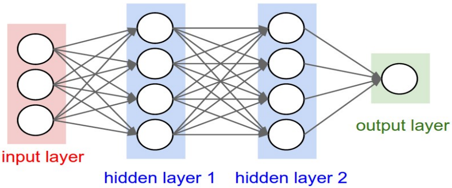

Idea Recursively approximate f_n by a neural network $f^\theta \colon \mathbb{R}^d \to \{0, 1\}$ of the form

$$f^\theta = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta,$$

where

- q_1 and q_2 are positive integers specifying the number of nodes in the two hidden layers,
- $a_1^\theta \colon \mathbb{R}^d \to \mathbb{R}^{q_1}$, $a_2^\theta \colon \mathbb{R}^{q_1} \to \mathbb{R}^{q_2}$ and $a_3^\theta \colon \mathbb{R}^{q_2} \to \mathbb{R}$ are affine functions given by
$$a_i^\theta(x) = A_i x + b_i, \ i = 1, 2, 3,$$
- for $j \in \mathbb{N}$, $\varphi_j \colon \mathbb{R}^j \to \mathbb{R}^j$ is the component-wise ReLU activation function given by $\varphi_j(x_1, \ldots, x_j) = (x_1^+, \ldots, x_j^+)$

The components of θ consist of the entries of A_i and b_i, $i = 1, 2, 3$ $\rightsquigarrow$ so # of parameters $\approx d^2$

More precisely,

- assume parameter values $\theta_{n+1}, \theta_{n+2}, \ldots, \theta_N \in \mathbb{R}^q$ have been found such that $f^{\theta_N} \equiv 1$ and the stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

  produces an expectation $\mathbb{E}\, g(\tau_{n+1}, X_{\tau_{n+1}})$ close to the optimal value V_{n+1}

More precisely,

- assume parameter values $\theta_{n+1}, \theta_{n+2}, \ldots, \theta_N \in \mathbb{R}^q$ have been found such that $f^{\theta_N} \equiv 1$ and the stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

  produces an expectation $\mathbb{E}\, g(\tau_{n+1}, X_{\tau_{n+1}})$ close to the optimal value V_{n+1}

- now try to find a maximizer $\theta_n \in \mathbb{R}^q$ of

$$\theta \mapsto \mathbb{E}\left[g(n, X_n) f^{\theta}(X_n) + g(\tau_{n+1}, X_{n+1})(1 - f^{\theta}(X_n))\right]$$

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method
- **Problem** for $x \in \mathbb{R}^d$, the θ-gradient of

$$f^\theta(x) = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta(x)$$

is 0 or does not exist

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method
- **Problem** for $x \in \mathbb{R}^d$, the θ-gradient of

$$f^\theta(x) = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta(x)$$

  is 0 or does not exist

- As an **intermediate step** consider a neural network $F^\theta \colon \mathbb{R}^d \to (0,1)$ of the form

$$F^\theta = \psi \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta \quad \text{for} \quad \psi(x) = \frac{e^x}{1+e^x}$$

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method
- **Problem** for $x \in \mathbb{R}^d$, the θ-gradient of

$$f^\theta(x) = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta(x)$$

  is 0 or does not exist

- As an **intermediate step** consider a neural network $F^\theta \colon \mathbb{R}^d \to (0,1)$ of the form

$$F^\theta = \psi \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta \quad \text{for} \quad \psi(x) = \frac{e^x}{1+e^x}$$

- Use **stochastic gradient ascent** to find an approximate optimizer $\theta_n \in \mathbb{R}^q$ of

$$\theta \mapsto \mathbb{E}\left[g(n, X_n) F^\theta(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - F^\theta(X_n))\right]$$

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method
- **Problem** for $x \in \mathbb{R}^d$, the θ-gradient of

$$f^\theta(x) = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta(x)$$

  is 0 or does not exist

- As an **intermediate step** consider a neural network $F^\theta \colon \mathbb{R}^d \to (0,1)$ of the form

$$F^\theta = \psi \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta \quad \text{for} \quad \psi(x) = \frac{e^x}{1+e^x}$$

- Use **stochastic gradient ascent** to find an approximate optimizer $\theta_n \in \mathbb{R}^q$ of

$$\theta \mapsto \mathbb{E}\left[g(n,X_n)F^\theta(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1-F^\theta(X_n))\right]$$

- **Approximate** $$f_n \approx f^{\theta_n} = 1_{[0,\infty)} \circ a_3^{\theta_n} \circ \varphi_{q_2} \circ a_2^{\theta_n} \circ \varphi_{q_1} \circ a_1^{\theta_n}$$

- **Goal** find an (approximately) optimal $\theta_n \in \mathbb{R}^q$ with a stochastic gradient ascent method
- **Problem** for $x \in \mathbb{R}^d$, the θ-gradient of

$$f^\theta(x) = 1_{[0,\infty)} \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta(x)$$

  is 0 or does not exist

- As an **intermediate step** consider a neural network $F^\theta \colon \mathbb{R}^d \to (0,1)$ of the form

$$F^\theta = \psi \circ a_3^\theta \circ \varphi_{q_2} \circ a_2^\theta \circ \varphi_{q_1} \circ a_1^\theta \quad \text{for} \quad \psi(x) = \frac{e^x}{1+e^x}$$

- Use **stochastic gradient ascent** to find an approximate optimizer $\theta_n \in \mathbb{R}^q$ of

$$\theta \mapsto \mathbb{E}\left[g(n, X_n) F^\theta(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - F^\theta(X_n))\right]$$

- **Approximate** $f_n \approx f^{\theta_n} = 1_{[0,\infty)} \circ a_3^{\theta_n} \circ \varphi_{q_2} \circ a_2^{\theta_n} \circ \varphi_{q_1} \circ a_1^{\theta_n}$
- **Repeat the same steps** at times $n-1, n-2, \ldots, 0$

Proposition

Let $n \in \{0, 1, \ldots, N-1\}$ and fix a stopping time $\tau_{n+1} \in \mathcal{T}_{n+1}$. Then, for every constant $\varepsilon > 0$, there exist numbers of hidden nodes q_1 and q_2 such that

$$\begin{aligned} &\sup_{\theta \in \mathbb{R}^q} \mathbb{E}\left[g(n, X_n) f^\theta(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f^\theta(X_n))\right] \\ &\geq \sup_{f \in \mathcal{D}} \mathbb{E}\left[g(n, X_n) f(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f(X_n))\right] - \varepsilon, \end{aligned}$$

where $\mathcal{D}$ is the set of all measurable functions $f : \mathbb{R}^d \to \{0, 1\}$.

Proposition

Let $n \in \{0, 1, \ldots, N-1\}$ and fix a stopping time $\tau_{n+1} \in \mathcal{T}_{n+1}$. Then, for every constant $\varepsilon > 0$, there exist numbers of hidden nodes q_1 and q_2 such that

$$\begin{aligned} &\sup_{\theta \in \mathbb{R}^q} \mathbb{E}\left[g(n, X_n) f^{\theta}(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f^{\theta}(X_n))\right] \\ &\geq \sup_{f \in \mathcal{D}} \mathbb{E}\left[g(n, X_n) f(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f(X_n))\right] - \varepsilon, \end{aligned}$$

where $\mathcal{D}$ is the set of all measurable functions $f : \mathbb{R}^d \to \{0, 1\}$.

Proof

1. Every measurable set $A \subseteq \mathbb{R}^d$ can be approximated in measure by compact sets $K \subseteq A$

Proposition

Let $n \in \{0, 1, \ldots, N-1\}$ and fix a stopping time $\tau_{n+1} \in \mathcal{T}_{n+1}$. Then, for every constant $\varepsilon > 0$, there exist numbers of hidden nodes q_1 and q_2 such that

$$\begin{aligned}
&\sup_{\theta \in \mathbb{R}^q} \mathbb{E}\left[g(n, X_n) f^\theta(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f^\theta(X_n))\right] \\
&\geq \sup_{f \in \mathcal{D}} \mathbb{E}\left[g(n, X_n) f(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f(X_n))\right] - \varepsilon,
\end{aligned}$$

where $\mathcal{D}$ is the set of all measurable functions $f : \mathbb{R}^d \to \{0, 1\}$.

Proof

1. Every measurable set $A \subseteq \mathbb{R}^d$ can be approximated in measure by compact sets $K \subseteq A$
2. $1_K - 1_{K^c}$ can be approximated by continuous functions k_j

Proposition

Let $n \in \{0, 1, \dots, N-1\}$ and fix a stopping time $\tau_{n+1} \in \mathcal{T}_{n+1}$. Then, for every constant $\varepsilon > 0$, there exist numbers of hidden nodes q_1 and q_2 such that

$$\begin{aligned} &\sup_{\theta \in \mathbb{R}^q} \mathbb{E}\left[g(n, X_n) f^{\theta}(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f^{\theta}(X_n))\right] \\ &\geq \sup_{f \in \mathcal{D}} \mathbb{E}\left[g(n, X_n) f(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f(X_n))\right] - \varepsilon, \end{aligned}$$

where $\mathcal{D}$ is the set of all measurable functions $f : \mathbb{R}^d \to \{0, 1\}$.

Proof

1. Every measurable set $A \subseteq \mathbb{R}^d$ can be approximated in measure by compact sets $K \subseteq A$
2. $1_K - 1_{K^c}$ can be approximated by continuous functions k_j
3. k_j can be approximated uniformly on compacts by functions of the form

$$h(x) = \sum_{i=1}^{r} (v_i^T x + c_i)^+ - \sum_{i=1}^{s} (w_i^T x + d_i)^+ \quad \text{(Leshno–Lin–Pinkus–Schocken, 1993)}$$

Proposition

Let $n \in \{0, 1, \ldots, N-1\}$ and fix a stopping time $\tau_{n+1} \in \mathcal{T}_{n+1}$. Then, for every constant $\varepsilon > 0$, there exist numbers of hidden nodes q_1 and q_2 such that

$$\sup_{\theta \in \mathbb{R}^q} \mathbb{E}\left[g(n, X_n) f^{\theta}(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f^{\theta}(X_n))\right]$$
$$\geq \sup_{f \in \mathcal{D}} \mathbb{E}\left[g(n, X_n) f(X_n) + g(\tau_{n+1}, X_{\tau_{n+1}})(1 - f(X_n))\right] - \varepsilon,$$

where $\mathcal{D}$ is the set of all measurable functions $f : \mathbb{R}^d \to \{0, 1\}$.

Proof

1. Every measurable set $A \subseteq \mathbb{R}^d$ can be approximated in measure by compact sets $K \subseteq A$
2. $1_K - 1_{K^c}$ can be approximated by continuous functions k_j
3. k_j can be approximated uniformly on compacts by functions of the form

$$h(x) = \sum_{i=1}^{r} (v_i^T x + c_i)^+ - \sum_{i=1}^{s} (w_i^T x + d_i)^+ \quad \text{(Leshno–Lin–Pinkus–Schocken, 1993)}$$

4. $1_{[0,\infty)} \circ h$ can be written as a neural network of the form $f^{\theta} = 1_{[0,\infty)} \circ a_3^{\theta} \circ \varphi_{q_2} \circ a_2^{\theta} \circ \varphi_{q_1} \circ a_1^{\theta}$

Corollary

For a given optimal stopping problem of the form

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau)$$

and a constant $\varepsilon > 0$,

Corollary

For a given optimal stopping problem of the form

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau)$$

and a constant $\varepsilon > 0$, there exist

- numbers of hidden nodes q_1, q_2 and
- functions $f^{\theta_0}, f^{\theta_1}, \ldots, f^{\theta_N} : \mathbb{R}^d \to \{0, 1\}$ of the form

$$f^{\theta_n} = 1_{[0,\infty)} \circ a_3^{\theta_n} \circ \varphi_{q_2} \circ a_2^{\theta_n} \circ \varphi_{q_1} \circ a_1^{\theta_n}$$

such that $f^{\theta_N} \equiv 1$ and the stopping time

$$\tau^\Theta = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

satisfies $\mathbb{E}\, g(\tau^\Theta, X_{\tau^\Theta}) \geq \sup_{\tau \in \mathcal{T}} \mathbb{E}\, g(\tau, X_\tau) - \varepsilon$

Training the networks

- Let $(x_n^k)_{n=0}^N$, $k = 1, 2, \ldots$ be independent simulations of $(X_n)_{n=0}^N$

Training the networks

- Let $(x_n^k)_{n=0}^N$, $k = 1, 2, \ldots$ be independent simulations of $(X_n)_{n=0}^N$
- Let $\theta_{n+1}, \ldots, \theta_N \in \mathbb{R}^q$ be given, and consider the corresponding stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

Training the networks

- Let $(x_n^k)_{n=0}^N$, $k = 1, 2, \ldots$ be independent simulations of $(X_n)_{n=0}^N$
- Let $\theta_{n+1}, \ldots, \theta_N \in \mathbb{R}^q$ be given, and consider the corresponding stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

- τ_{n+1} is of the form $\tau_{n+1} = l_{n+1}(X_{n+1}, \ldots, X_{N-1})$ for a measurable function

$$l_{n+1} : \mathbb{R}^{d(N-n-1)} \to \{n+1, n+2, \ldots, N\}$$

Training the networks

- Let $(x_n^k)_{n=0}^N$, $k = 1, 2, \ldots$ be independent simulations of $(X_n)_{n=0}^N$
- Let $\theta_{n+1}, \ldots, \theta_N \in \mathbb{R}^q$ be given, and consider the corresponding stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

- τ_{n+1} is of the form $\tau_{n+1} = l_{n+1}(X_{n+1}, \ldots, X_{N-1})$ for a measurable function

$$l_{n+1} : \mathbb{R}^{d(N-n-1)} \to \{n+1, n+2, \ldots, N\}$$

- Denote

$$l_{n+1}^k = \begin{cases} N & \text{if } n = N-1 \\ l_{n+1}(x_{n+1}^k, \ldots, x_{N-1}^k) & \text{if } n \leq N-2 \end{cases}$$

Training the networks

- Let $(x_n^k)_{n=0}^N$, $k = 1, 2, \ldots$ be independent simulations of $(X_n)_{n=0}^N$
- Let $\theta_{n+1}, \ldots, \theta_N \in \mathbb{R}^q$ be given, and consider the corresponding stopping time

$$\tau_{n+1} = \sum_{m=n+1}^{N} m f^{\theta_m}(X_m) \prod_{j=n+1}^{m-1} (1 - f^{\theta_j}(X_j))$$

- τ_{n+1} is of the form $\tau_{n+1} = l_{n+1}(X_{n+1}, \ldots, X_{N-1})$ for a measurable function

$$l_{n+1} : \mathbb{R}^{d(N-n-1)} \to \{n+1, n+2, \ldots, N\}$$

- Denote

$$l_{n+1}^k = \begin{cases} N & \text{if } n = N-1 \\ l_{n+1}(x_{n+1}^k, \ldots, x_{N-1}^k) & \text{if } n \leq N-2 \end{cases}$$

- The realized reward

$$r_n^k(\theta) = g(n, x_n^k) F^\theta(x_n^k) + g(l_{n+1}^k, x_{l_{n+1}^k}^k)(1 - F^\theta(x_n^k))$$

is continuous and almost everywhere differentiable in θ

Stochastic Gradient Ascent

- **Initialize** $\theta_{n,0}$ typically random; e.g. Xavier initialization

Stochastic Gradient Ascent

- **Initialize** $\theta_{n,0}$ typically random; e.g. Xavier initialization
- **Standard updating** $\theta_{n,k+1} = \theta_{n,k} + \eta \nabla r_n^k(\theta_{n,k})$

Stochastic Gradient Ascent

- **Initialize** $\theta_{n,0}$ typically random; e.g. Xavier initialization
- **Standard updating** $\theta_{n,k+1} = \theta_{n,k} + \eta \nabla r_n^k(\theta_{n,k})$
- **Variants**
  - Mini-batches
  - Batch normalization
  - Momentum
  - Adagrad
  - RMSProp
  - AdaDelta
  - ADAM
  - Decoupling weight decay
  - Warm restarts
  - ...

Stochastic Gradient Ascent

- **Initialize** $\theta_{n,0}$ typically random; e.g. Xavier initialization
- **Standard updating** $\theta_{n,m+1} = \theta_{n,m} + \eta \nabla r_n^m(\theta_{n,m})$
- **Variants**
  - Mini-batches
  - Batch normalization
  - Momentum
  - Adagrad
  - RMSProp
  - AdaDelta
  - ADAM
  - Decoupling weight decay
  - Warm restarts
  - ...

Lower bound

- The candidate optimal stopping time

$$\tau^{\Theta} = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

  yields a lower bound

$$L = \mathbb{E}\, g(\tau^{\Theta}, X_{\tau^{\Theta}}) \quad \text{for the optimal value} \quad V_0 = \sup_{\tau} \mathbb{E}\, g(\tau, X_{\tau})$$

Lower bound

- The candidate optimal stopping time

$$\tau^{\Theta} = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

  yields a lower bound

$$L = \mathbb{E}\, g(\tau^{\Theta}, X_{\tau^{\Theta}}) \quad \text{for the optimal value} \quad V_0 = \sup_{\tau} \mathbb{E}\, g(\tau, X_{\tau})$$

- Let $(y_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_L$, be a new set of independent simulations of $(X_n)_{n=0}^N$

Lower bound

- The candidate optimal stopping time

$$\tau^{\Theta} = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

  yields a lower bound

$$L = \mathbb{E}\, g(\tau^{\Theta}, X_{\tau^{\Theta}}) \quad \text{for the optimal value} \quad V_0 = \sup_{\tau} \mathbb{E}\, g(\tau, X_{\tau})$$

- Let $(y_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_L$, be a new set of independent simulations of $(X_n)_{n=0}^N$
- τ^{Θ} can be written as $\tau^{\Theta} = l(X_0, \ldots, X_{N-1})$ for a measurable function $l : \mathbb{R}^{dN} \to \{0, 1, \ldots, N\}$

Lower bound

- The candidate optimal stopping time

$$\tau^{\Theta} = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

  yields a lower bound

$$L = \mathbb{E}\, g(\tau^{\Theta}, X_{\tau^{\Theta}}) \quad \text{for the optimal value} \quad V_0 = \sup_{\tau} \mathbb{E}\, g(\tau, X_{\tau})$$

- Let $(y_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_L$, be a new set of independent simulations of $(X_n)_{n=0}^N$
- τ^{Θ} can be written as $\tau^{\Theta} = l(X_0, \ldots, X_{N-1})$ for a measurable function $l : \mathbb{R}^{dN} \to \{0, 1, \ldots, N\}$
- Denote $l^k = l(y_0^k, \ldots, y_{N-1}^k)$

Lower bound

- The candidate optimal stopping time

$$\tau^{\Theta} = \sum_{n=1}^{N} n f^{\theta_n}(X_n) \prod_{j=0}^{n-1} (1 - f^{\theta_j}(X_j))$$

yields a lower bound

$$L = \mathbb{E}\, g(\tau^{\Theta}, X_{\tau^{\Theta}}) \quad \text{for the optimal value} \quad V_0 = \sup_{\tau} \mathbb{E}\, g(\tau, X_\tau)$$

- Let $(y_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_L$, be a new set of independent simulations of $(X_n)_{n=0}^N$
- τ^{Θ} can be written as $\tau^{\Theta} = l(X_0, \ldots, X_{N-1})$ for a measurable function $l : \mathbb{R}^{dN} \to \{0, 1, \ldots, N\}$
- Denote $l^k = l(y_0^k, \ldots, y_{N-1}^k)$
- Use the Monte Carlo approximation

$$\hat{L} = \frac{1}{K_L} \sum_{k=1}^{K_L} g(l^k, y_{l^k}^k) \quad \text{as an estimate for} \quad L$$

Lower confidence bounds

- Assume $\mathbb{E}\left[g(n,X_n)^2\right] < \infty$ for all $n = 0, 1, \ldots, N$

Lower confidence bounds

- Assume $\mathbb{E}\left[g(n, X_n)^2\right] < \infty$ for all $n = 0, 1, \ldots, N$
- Consider the sample variance

$$\hat{\sigma}_L^2 = \frac{1}{K_L - 1} \sum_{k=1}^{K_L} \left(g(l^k, y_{l^k}^k) - \hat{L}\right)^2$$

Lower confidence bounds

- Assume $\mathbb{E}\left[g(n, X_n)^2\right] < \infty$ for all $n = 0, 1, \ldots, N$
- Consider the sample variance

$$\hat{\sigma}_L^2 = \frac{1}{K_L - 1} \sum_{k=1}^{K_L} \left(g(l^k, y_{l^k}^k) - \hat{L}\right)^2$$

- By the CLT,

$$\left[\hat{L} - z_\alpha \frac{\hat{\sigma}_L}{\sqrt{K_L}}, \infty\right)$$

is an asymptotically valid $1 - \alpha$ confidence interval for L

where z_α is the $1 - \alpha$ quantile of the standard normal distribution

Lower confidence bounds

- Assume $\mathbb{E}\left[g(n, X_n)^2\right] < \infty$ for all $n = 0, 1, \ldots, N$
- Consider the sample variance

$$\hat{\sigma}_L^2 = \frac{1}{K_L - 1} \sum_{k=1}^{K_L} \left(g(l^k, y_{l^k}^k) - \hat{L}\right)^2$$

- By the CLT,

$$\left[\hat{L} - z_\alpha \frac{\hat{\sigma}_L}{\sqrt{K_L}}, \infty\right)$$

  is an asymptotically valid $1 - \alpha$ confidence interval for L

  where z_α is the $1 - \alpha$ quantile of the standard normal distribution

- Therefore,

$$\mathbb{P}\left[V_0 \geq \hat{L} - z_\alpha \frac{\hat{\sigma}_L}{\sqrt{K_L}}\right] \geq \mathbb{P}\left[L \geq \hat{L} - z_\alpha \frac{\hat{\sigma}_L}{\sqrt{K_L}}\right] \approx 1 - \alpha$$

Upper bound

Let (H_n) be the Snell envelope of $G_n = g(n, X_n)$, $n = 0, 1, \ldots, N$,

Upper bound

Let (H_n) be the Snell envelope of $G_n = g(n, X_n)$, $n = 0, 1, \dots, N$, with Doob decomposition $H_n = H_0 + M_n^H - A_n^H$

Upper bound

Let (H_n) be the Snell envelope of $G_n = g(n, X_n)$, $n = 0, 1, \ldots, N$, with Doob decomposition $H_n = H_0 + M_n^H - A_n^H$

The following is a variant of the dual formulation of Rogers (2002), Haugh–Kogan (2004) and Andersen–Broadie (2004)

Upper bound

Let (H_n) be the Snell envelope of $G_n = g(n, X_n)$, $n = 0, 1, \ldots, N$,
with Doob decomposition $H_n = H_0 + M_n^H - A_n^H$

The following is a variant of the dual formulation of Rogers (2002), Haugh–Kogan (2004) and Andersen–Broadie (2004)

Proposition

For every $(\mathcal{F}_n^X)$-martingale (M_n) with $M_0 = 0$ and estimation errors (ε_n) satisfying $\mathbb{E}[\varepsilon_n \mid \mathcal{F}_n^X] = 0$, one has

$$V_0 \leq \mathbb{E}\left[\max_{0 \leq n \leq N} (G_n - M_n - \varepsilon_n)\right]$$

Upper bound

Let (H_n) be the Snell envelope of $G_n = g(n, X_n)$, $n = 0, 1, \ldots, N$,
with Doob decomposition $H_n = H_0 + M_n^H - A_n^H$

The following is a variant of the dual formulation of Rogers (2002), Haugh–Kogan (2004) and Andersen–Broadie (2004)

Proposition

For every $(\mathcal{F}_n^X)$-martingale (M_n) with $M_0 = 0$ and estimation errors (ε_n) satisfying $\mathbb{E}[\varepsilon_n \mid \mathcal{F}_n^X] = 0$, one has

$$V_0 \leq \mathbb{E}\left[\max_{0 \leq n \leq N} (G_n - M_n - \varepsilon_n)\right]$$

On the other hand,

$$V_0 = \mathbb{E}\left[\max_{0 \leq n \leq N} \left(G_n - M_n^H\right)\right]$$

Estimating a good dual martingale

- Approximate H_n by $H_n^{\Theta} = \mathbb{E}\left[G_{\tau_n^{\Theta}} \mid \mathcal{F}_n^X\right]$

Estimating a good dual martingale

- Approximate H_n by $H_n^\Theta = \mathbb{E}\left[G_{\tau_n^\Theta} \mid \mathcal{F}_n^X\right]$
- and $\Delta M_n^H = H_n - \mathbb{E}\left[H_n \mid \mathcal{F}_{n-1}^X\right]$ by

$$\Delta M_n^\Theta = H_n^\Theta - \mathbb{E}\left[H_n^\Theta \mid \mathcal{F}_{n-1}\right] = f^{\theta_n}(X_n)G_n + (1 - f^{\theta_n}(X_n))C_n^\Theta - C_{n-1}^\Theta$$

for the continuation values

$$C_n^\Theta = \mathbb{E}\left[G_{\tau_{n+1}^\Theta} \mid \mathcal{F}_n^X\right]$$

Estimating a good dual martingale

- Approximate H_n by $H_n^\Theta = \mathbb{E}\left[G_{\tau_n^\Theta} \mid \mathcal{F}_n^X\right]$
- and $\Delta M_n^H = H_n - \mathbb{E}\left[H_n \mid \mathcal{F}_{n-1}^X\right]$ by

$$\Delta M_n^\Theta = H_n^\Theta - \mathbb{E}\left[H_n^\Theta \mid \mathcal{F}_{n-1}\right] = f^{\theta_n}(X_n)G_n + (1 - f^{\theta_n}(X_n))C_n^\Theta - C_{n-1}^\Theta$$

for the continuation values

$$C_n^\Theta = \mathbb{E}\left[G_{\tau_{n+1}^\Theta} \mid \mathcal{F}_n^X\right]$$

- Let $(z_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_U$, be a third set of independent simulations of $(X_n)_{n=0}^N$

Estimating a good dual martingale

- Approximate H_n by $H_n^\Theta = \mathbb{E}\left[G_{\tau_n^\Theta} \mid \mathcal{F}_n^X\right]$
- and $\Delta M_n^H = H_n - \mathbb{E}\left[H_n \mid \mathcal{F}_{n-1}^X\right]$ by

$$\Delta M_n^\Theta = H_n^\Theta - \mathbb{E}\left[H_n^\Theta \mid \mathcal{F}_{n-1}\right] = f^{\theta_n}(X_n)G_n + (1 - f^{\theta_n}(X_n))C_n^\Theta - C_{n-1}^\Theta$$

for the continuation values

$$C_n^\Theta = \mathbb{E}\left[G_{\tau_{n+1}^\Theta} \mid \mathcal{F}_n^X\right]$$

- Let $(z_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_U$, be a third set of independent simulations of $(X_n)_{n=0}^N$
- For all z_n^k, simulate J independent continuation paths $\tilde{z}_{n+1}^{k,j}, \ldots, \tilde{z}_N^{k,j}$

Estimating a good dual martingale

- Approximate H_n by $H_n^\Theta = \mathbb{E}\left[G_{\tau_n^\Theta} \mid \mathcal{F}_n^X\right]$
- and $\Delta M_n^H = H_n - \mathbb{E}\left[H_n \mid \mathcal{F}_{n-1}^X\right]$ by

$$\Delta M_n^\Theta = H_n^\Theta - \mathbb{E}\left[H_n^\Theta \mid \mathcal{F}_{n-1}\right] = f^{\theta_n}(X_n)G_n + (1 - f^{\theta_n}(X_n))C_n^\Theta - C_{n-1}^\Theta$$

for the continuation values

$$C_n^\Theta = \mathbb{E}\left[G_{\tau_{n+1}^\Theta} \mid \mathcal{F}_n^X\right]$$

- Let $(z_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_U$, be a third set of independent simulations of $(X_n)_{n=0}^N$
- For all z_n^k, simulate J independent continuation paths $\tilde{z}_{n+1}^{k,j}, \ldots, \tilde{z}_N^{k,j}$
- $$C_n^k = \frac{1}{J}\sum_{j=1}^J g\left(\tau_{n+1}^{k,j}, \tilde{z}^{k,j}_{\tau_{n+1}^{k,j}}\right)$$

can be understood as realizations of $C_n^\Theta + \tilde{\varepsilon}_n$

Estimating a good dual martingale

- Approximate H_n by $H_n^{\Theta} = \mathbb{E}\left[G_{\tau_n^{\Theta}} \mid \mathcal{F}_n^X\right]$
- and $\Delta M_n^H = H_n - \mathbb{E}\left[H_n \mid \mathcal{F}_{n-1}^X\right]$ by

$$\Delta M_n^{\Theta} = H_n^{\Theta} - \mathbb{E}\left[H_n^{\Theta} \mid \mathcal{F}_{n-1}\right] = f^{\theta_n}(X_n)G_n + (1 - f^{\theta_n}(X_n))C_n^{\Theta} - C_{n-1}^{\Theta}$$

  for the continuation values

$$C_n^{\Theta} = \mathbb{E}\left[G_{\tau_{n+1}^{\Theta}} \mid \mathcal{F}_n^X\right]$$

- Let $(z_n^k)_{n=0}^N$, $k = 1, 2, \ldots, K_U$, be a third set of independent simulations of $(X_n)_{n=0}^N$
- For all z_n^k, simulate J independent continuation paths $\tilde{z}_{n+1}^{k,j}, \ldots, \tilde{z}_N^{k,j}$
-

$$C_n^k = \frac{1}{J} \sum_{j=1}^{J} g\left(\tau_{n+1}^{k,j}, \tilde{z}^{k,j}_{\tau_{n+1}^{k,j}}\right)$$

  can be understood as realizations of $C_n^{\Theta} + \tilde{\varepsilon}_n$
- This gives realizations M_n^k of $M_n^{\Theta} + \varepsilon_n$

Estimating an upper bound

- $$U = \mathbb{E}\left[\max_{0 \le n \le N}\left(G_n - M_n^{\Theta} - \varepsilon_n\right)\right]$$ is an upper bound for V_0

Estimating an upper bound

- $$U = \mathbb{E}\left[\max_{0\leq n\leq N}\left(G_n - M_n^{\Theta} - \varepsilon_n\right)\right] \quad \text{is an upper bound for} \quad V_0$$

- Use the Monte Carlo approximation

$$\hat{U} = \frac{1}{K_U}\sum_{k=1}^{K_U}\max_{0\leq n\leq N}\left(g(n, z_n^k) - M_n^k\right) \quad \text{as an estimate for} \quad U$$

Estimating an upper bound

- $$U = \mathbb{E}\left[\max_{0\leq n\leq N}\left(G_n - M_n^{\Theta} - \varepsilon_n\right)\right] \quad \text{is an upper bound for} \quad V_0$$

- Use the Monte Carlo approximation

$$\hat{U} = \frac{1}{K_U}\sum_{k=1}^{K_U}\max_{0\leq n\leq N}\left(g(n, z_n^k) - M_n^k\right) \quad \text{as an estimate for} \quad U$$

Our point estimate of V_0: $\dfrac{\hat{L} + \hat{U}}{2}$

Confidence intervals for V_0

- By the CLT,

$$\left(-\infty\,,\ \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right]$$

is an asymptotically valid $1-\alpha$ confidence interval for U, where $\hat{\sigma}_U$ is the corresponding sample standard deviation

Confidence intervals for V_0

- By the CLT,
$$\left(-\infty\,,\;\hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right]$$
is an asymptotically valid $1-\alpha$ confidence interval for U, where $\hat{\sigma}_U$ is the corresponding sample standard deviation

- One has
$$\mathbb{P}\left[V_0 \le \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right] \ge \mathbb{P}\left[U \le \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right] \approx 1-\alpha.$$

Confidence intervals for V_0

- By the CLT,

$$\left(-\infty\,,\ \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right]$$

  is an asymptotically valid $1-\alpha$ confidence interval for U, where $\hat{\sigma}_U$ is the corresponding sample standard deviation

- One has

$$\mathbb{P}\left[V_0 \leq \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right] \geq \mathbb{P}\left[U \leq \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right] \approx 1-\alpha.$$

- So

$$\left[\hat{L} - z_\alpha \frac{\hat{\sigma}_L}{\sqrt{K_L}}\,,\ \hat{U} + z_\alpha \frac{\hat{\sigma}_U}{\sqrt{K_U}}\right]$$

  is an asymptotically valid $1-2\alpha$ confidence interval

Thank You!