ETH zürich

Deep Optimal Stopping: Solving High-dimensional Optimal Stopping Problems with Deep Learning

Timo Welti (ETH Zürich, Switzerland)

Joint works with

Sebastian Becker (Zenai AG, Zürich, Switzerland),
Patrick Cheridito (ETH Zürich, Switzerland), and
Arnulf Jentzen (ETH Zürich, Switzerland)

Consider $T > 0$, $d \in \mathbb{N}$, $\xi \in \mathbb{R}^d$, probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a normal filtration $\mathbb{F} = (\mathbb{F}_t)_{t \in [0,T]}$, and

$$\sup\left\{\mathbb{E}\big[g(\tau, X_\tau)\big] : \begin{smallmatrix} \tau\colon \Omega \to [0,T] \text{ is an} \\ \mathbb{F}^X\text{-stopping time} \end{smallmatrix}\right\} \in \mathbb{R}, \qquad (\star)$$

where

- $g\colon [0,T] \times \mathbb{R}^d \to \mathbb{R}$ is continuous and at most polynomially growing,
- $X\colon [0,T] \times \Omega \to \mathbb{R}^d$ is an $\mathbb{F}$-adapted continuous solution process of

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dW_t, \qquad X_0 = \xi, \qquad t \in [0,T],$$

- $\mu\colon \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma\colon \mathbb{R}^d \to \mathbb{R}^{d \times d}$ are Lipschitz continuous,
- $W\colon [0,T] \times \Omega \to \mathbb{R}^d$ is a Brownian motion (continuous),
- $\mathbb{F}^X = (\mathbb{F}^X_t)_{t \in [0,T]}$ is the filtration generated by X.

Optimal stopping problem: Compute $(\star)$ as well as an *optimal exercise strategy* $\tau\colon \Omega \to [0,T]$

Consider the **P**artial **D**ifferential **E**quation (PDE)

$$\max\Big\{\big(\tfrac{\partial u}{\partial t}\big)(t,x) + \tfrac{1}{2}\operatorname{trace}\big(\sigma(x)[\sigma(x)]^*(\operatorname{Hess}_x u)(t,x)\big) + \langle(\nabla_x u)(t,x), \mu(x)\rangle_{\mathbb{R}^d}, g(t,x) - u(t,x)\Big\} = 0 \qquad (\triangle)$$

and $u(T,x) = g(T,x)$, where $0 \le t \le T$, $x \in \mathbb{R}^d$, and $u \in C([0,T]\times\mathbb{R}^d,\mathbb{R})$ satisfies $u|_{[0,T)\times\mathbb{R}^d} \in C^{1,2}([0,T)\times\mathbb{R}^d,\mathbb{R})$.

Under suitable regularity assumptions it holds that

$$u(0,\xi) = \sup\Big\{\mathbb{E}\big[g(\tau, X_\tau)\big] : \substack{\tau\colon \Omega\to[0,T] \text{ is an} \\ \mathbb{F}^X\text{-stopping time}}\Big\}.$$

Approximation methods for $(\triangle)$ such as, e.g., finite differences suffer under *the curse of dimensionality*.

Approximation I: temporal discretization

Consider $N \in \mathbb{N}$ and $0 = t_0 < t_1 < \cdots < t_N = T$.

We approximate

$$X_{t_n} \approx \mathcal{X}_n,$$

where $\mathcal{X} \colon \{0, 1, \ldots, N\} \times \Omega \to \mathbb{R}^d$ satisfies $\mathcal{X}_0 = \xi$ and

$$\mathcal{X}_{n+1} = \mathcal{X}_n + \mu(\mathcal{X}_n)\,(t_{n+1} - t_n) + \sigma(\mathcal{X}_n)\,(W_{t_{n+1}} - W_{t_n})\,.$$

Next we approximate

$$\sup\left\{\mathbb{E}\big[g(\tau, X_\tau)\big] : \substack{\tau\colon \Omega\to[0,T] \text{ is an} \\ \mathbb{F}^X\text{-stopping time}}\right\} \approx \sup\left\{\mathbb{E}\big[g(t_\rho, \mathcal{X}_\rho)\big] : \substack{\rho\colon \Omega\to\{0,1,\ldots,N\} \text{ is} \\ \text{an } \mathbb{F}^{\mathcal{X}}\text{-stopping time}}\right\},$$

where $\mathbb{F}^{\mathcal{X}} = (\mathbb{F}^{\mathcal{X}}_n)_{n\in\{0,1,\ldots,N\}}$ is the filtration generated by $\mathcal{X}$.

Lemma (factorization lemma for stopping times)

Let $d, N \in \mathbb{N}$, *let* $(\Omega, \mathcal{F}, \mathbb{P})$ *be a probability space, and let* $\mathcal{X} \colon \{0, 1, \ldots, N\} \times \Omega \to \mathbb{R}^d$ *be a stochastic process.*

Then for every $\mathbb{F}^{\mathcal{X}}$*-stopping time* $\rho \colon \Omega \to \{0, 1, \ldots, N\}$ *there exist Borel measurable* $\mathbb{U}_{n,\rho} \colon (\mathbb{R}^d)^{n+1} \to \{0, 1\}$, $n \in \{0, 1, \ldots, N\}$, *such that* $\forall\, x_0, x_1, \ldots, x_N \in \mathbb{R}^d \colon \sum_{n=0}^{N} \mathbb{U}_{n,\rho}(x_0, x_1, \ldots, x_n) = 1$ *and*

$$\rho = \sum_{n=0}^{N} n\, \mathbb{U}_{n,\rho}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n).$$

Proof: Let $\rho\colon \Omega \to \{0, 1, \dots, N\}$ be an $\mathbb{F}^{\mathcal{X}}$-stopping time.

Then

$$\Omega \ni \omega \mapsto \mathbb{1}_{\{\rho=n\}}(\omega) \in \{0, 1\}$$

is $\sigma_\Omega((\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_n))/\mathcal{B}(\{0, 1\})$-measurable.

Hence there exist Borel measurable

$$\mathbb{V}_{n,\rho}\colon (\mathbb{R}^d)^{n+1} \to \{0, 1\}, n \in \{0, 1, \dots, N\},$$

such that $\forall\, \omega \in \Omega\colon \mathbb{1}_{\{\rho=n\}}(\omega) = \mathbb{V}_{n,\rho}\big(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), \dots, \mathcal{X}_n(\omega)\big)$.

Let $\mathbb{U}_{n,\rho}\colon (\mathbb{R}^d)^{n+1} \to \{0, 1\}$, $n \in \{0, 1, \dots, N\}$, be the functions such that

$$\begin{aligned} &\mathbb{U}_{n,\rho}(x_0, x_1, \dots, x_n) \\ &= \max\{\mathbb{V}_{n,\rho}(x_0, x_1, \dots, x_n), n + 1 - N\}\left[1 - \sum_{k=0}^{n-1} \mathbb{U}_{k,\rho}(x_0, x_1, \dots, x_k)\right]. \end{aligned}$$

□

Approximation II: neural network architectures for stopping times

Consider Markov process $\mathcal{X}\colon \{0,1,\ldots,N\} \times \Omega \to \mathbb{R}^d$ and

$$\sup\left\{\mathbb{E}\big[g(t_\rho, \mathcal{X}_\rho)\big] : \substack{\rho\colon \Omega \to \{0,1,\ldots,N\} \text{ is} \\ \text{an } \mathbb{F}^{\mathcal{X}}\text{-stopping time}}\right\}.$$

Idea: Replace $\mathbb{V}_{n,\rho}\colon (\mathbb{R}^d)^{n+1} \to \{0,1\}$ by artificial neural networks

$$u_{n,\theta}\colon \mathbb{R}^d \to (0,1)$$

for suitable $\theta \in \mathbb{R}^\nu$, $\nu \in \mathbb{N}$ and construct $U_{n,\theta}\colon (\mathbb{R}^d)^{n+1} \to (0,1)$ through

$$\begin{aligned} &U_{n,\theta}(x_0, x_1, \ldots, x_n) \\ &= \max\{u_{n,\theta}(x_n), n+1-N\}\left[1 - \sum_{k=0}^{n-1} U_{k,\theta}(x_0, x_1, \ldots, x_k)\right]. \end{aligned}$$

We think $U_{n,\theta}(x_0, x_1, \ldots, x_n) \approx \mathbb{U}_{n,\rho}(x_0, x_1, \ldots, x_n)$.

Approximation II: neural network architectures for stopping times

Note that

$$g(t_\rho, \mathcal{X}_\rho) = \sum_{n=0}^{N} \mathbb{1}_{\{\rho=n\}}\, g(t_n, \mathcal{X}_n) = \sum_{n=0}^{N} \mathbb{U}_{n,\rho}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n)\, g(t_n, \mathcal{X}_n)$$
$$\approx \sum_{n=0}^{N} U_{n,\theta}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n)\, g(t_n, \mathcal{X}_n).$$

We hence approximate

$$\sup\left\{ \mathbb{E}\big[g(t_\rho, \mathcal{X}_\rho)\big] : \substack{\rho\colon \Omega \to \{0,1,\ldots,N\} \text{ is} \\ \text{an } \mathbb{F}^{\mathcal{X}}\text{-stopping time}} \right\}$$
$$\approx \sup\left\{ \mathbb{E}\left[\sum_{n=0}^{N} U_{n,\theta}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n)\, g(t_n, \mathcal{X}_n) \right] : \theta \in \mathbb{R}^\nu \right\},$$

i.e., we suggest to maximize

$$\mathbb{R}^\nu \ni \theta \mapsto \mathbb{E}\left[\sum_{n=0}^{N} U_{n,\theta}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n)\, g(t_n, \mathcal{X}_n) \right] \in \mathbb{R}.$$

Approximation III: stochastic gradient ascent-type approximations

We suggest to maximize

$$\mathbb{R}^\nu \ni \theta \mapsto \mathbb{E}\left[\sum_{n=0}^{N} U_{n,\theta}(\mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_n)\, g(t_n, \mathcal{X}_n)\right] \in \mathbb{R}$$

by stochastic gradient ascent-type optimization algorithms.

This results in random approximations

$$\Theta_m = (\Theta_m^{(1)}, \ldots, \Theta_m^{(\nu)}) \colon \Omega \to \mathbb{R}^\nu$$

for $m \in \{0, 1, 2, \ldots\}$.

Let $M \in \mathbb{N}$ and consider a realization $\widehat{\Theta}_M \in \mathbb{R}^\nu$ of Θ_M.

Approximation IV: final approximations

Approximation for *optimal exercise strategy:*

$\mathbb{F}^{\mathcal{X}}$-stopping time $\tau_{\widehat{\Theta}_M} : \Omega \to \{0, 1, \ldots, N\}$ given by

$$\tau_{\widehat{\Theta}_M} = \min\left\{ n \in \{0, 1, \ldots, N\} \colon \sum_{k=0}^{n} U_{k,\widehat{\Theta}_M}(\mathcal{X}_0, \ldots, \mathcal{X}_k) \geq 1 - U_{n,\widehat{\Theta}_M}(\mathcal{X}_0, \ldots, \mathcal{X}_n) \right\}.$$

Approximation for *price*:

Monte Carlo approximation of

$$\mathbb{E}\left[g\left(t_{\tau_{\widehat{\Theta}_M}}, \mathcal{X}_{\tau_{\widehat{\Theta}_M}} \right) \right].$$

Recall: Replace $\mathbb{V}_{n,\rho} \colon (\mathbb{R}^d)^{n+1} \to \{0,1\}$ by artificial neural networks $u_{n,\theta} \colon \mathbb{R}^d \to (0,1)$ for suitable $\theta \in \mathbb{R}^\nu$, $\nu \in \mathbb{N}$.

Let $\nu = (N+1)(2d+1)(d+1)$, $\forall\, k \in \mathbb{N}$ let $\mathcal{L}_k \colon \mathbb{R}^k \to (0,1)^k$ satisfy $\forall\, x = (x_1, \ldots, x_k) \in \mathbb{R}^k$:

$$\mathcal{L}_k(x) = \left(\frac{\exp(x_1)}{\exp(x_1)+1}, \frac{\exp(x_2)}{\exp(x_2)+1}, \ldots, \frac{\exp(x_k)}{\exp(x_k)+1} \right),$$

$\forall\, \theta = (\theta_1, \ldots, \theta_\nu) \in \mathbb{R}^\nu$, $v \in \mathbb{N}_0$, $k, l \in \mathbb{N}$ with $v + k(l+1) \leq \nu$ let $A^{\theta,v}_{k,l} \colon \mathbb{R}^l \to \mathbb{R}^k$ satisfy $\forall\, x = (x_1, \ldots, x_l) \in \mathbb{R}^l$:

$$A^{\theta,v}_{k,l}(x) = \begin{pmatrix} \theta_{v+1} & \theta_{v+2} & \cdots & \theta_{v+l} \\ \theta_{v+l+1} & \theta_{v+l+2} & \cdots & \theta_{v+2l} \\ \vdots & \vdots & \vdots & \vdots \\ \theta_{v+(k-1)l+1} & \theta_{v+(k-1)l+2} & \cdots & \theta_{v+kl} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_l \end{pmatrix} + \begin{pmatrix} \theta_{v+kl+1} \\ \theta_{v+kl+2} \\ \vdots \\ \theta_{v+kl+k} \end{pmatrix},$$

and $\forall\, n \in \{0, 1, \ldots, N\}$, $\theta \in \mathbb{R}^\nu$ let $u_{n,\theta} \colon \mathbb{R}^d \to (0,1)$ satisfy

$$u_{n,\theta} = \mathcal{L}_1 \circ A^{\theta,(2nd+n)(d+1)}_{1,d} \circ \mathcal{L}_d \circ A^{\theta,(2nd+n+1)(d+1)}_{d,d} \circ \mathcal{L}_d \circ A^{\theta,((2n+1)d+n+1)(d+1)}_{d,d}.$$

Numerical simulations

Implementations in PYTHON using TENSORFLOW

GPU: NVIDIA GeForce GTX 1080 @ 1974 MHz, 2560 cuda cores,
8 GB GDDR5X @ 1809.5 MHz,

CPU: Intel Core i7-6800K @ 3.4 GHz, 6 cores, 64 GB DDR4

American basket strangle spread option in the Black Scholes model

Example from Kohler, Krzyżak, & Todorovic 2010

Consider

- $d = 5$, $r = 5\%$, $K_1 = 75$, $K_2 = 90$, $K_3 = 110$, $K_4 = 125$,
- $\Sigma = (\varsigma_1, \ldots, \varsigma_d) \in \mathbb{R}^{5\times 5}$ fixed historical volatility data,
- $T = 1$, $N = 48$, $t_n = \frac{nT}{N}$, $\xi = (100, \ldots, 100) \in \mathbb{R}^d$,
- $\mu \colon \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma \colon \mathbb{R}^d \to \mathbb{R}^{d\times d}$ satisfying $\mu(x) = r\, x$ and $\sigma(x) = \operatorname{diag}(x_1, \ldots, x_d)\, \Sigma^*$,
- $\mathcal{X} = (\mathcal{X}^{(1)}, \ldots, \mathcal{X}^{(d)}) \colon \{0, 1, \ldots, N\} \times \Omega \to \mathbb{R}^d$ satisfying

$$\mathcal{X}_n^{(i)} = X_{t_n}^{(i)} = \exp\left(\left[r - \tfrac{1}{2}\|\varsigma_i\|_{\mathbb{R}^d}^2\right] t_n + \left\langle \varsigma_i, W_{t_n} \right\rangle_{\mathbb{R}^d}\right) \xi_i,$$

- $g \colon [0, T] \times \mathbb{R}^d \to \mathbb{R}$ satisfying

$$g(t, x) = -e^{-rt} \max\left\{K_1 - \frac{1}{d}\sum_{i=1}^{d} x_i, 0\right\} + e^{-rt} \max\left\{K_2 - \frac{1}{d}\sum_{i=1}^{d} x_i, 0\right\} + e^{-rt} \max\left\{\frac{1}{d}\sum_{i=1}^{d} x_i - K_3, 0\right\} - e^{-rt} \max\left\{\frac{1}{d}\sum_{i=1}^{d} x_i - K_4, 0\right\}.$$

We approximate the price

$$\sup\left\{\mathbb{E}\left[g(\tau, X_\tau)\right] : \begin{matrix}\tau : \Omega \to [0,T] \text{ is an} \\ \mathbb{F}^X\text{-stopping time}\end{matrix}\right\}.$$

Numerical results based on 10 independent runs:

Mean	Std. dev.	Avg. runtime	Reference LB $(\star)$
11.794	0.004	46.4	11.75

$(\star)$ lower bound in Kohler, Krzyżak, & Todorovic 2010

Bermudan max-call option in the Black Scholes model

Benchmark example for $d = 5$ (Longstaff & Schwartz 2001, Rogers 2002, Broadie & Glasserman 2004, Andersen & Broadie 2004, ...)

Consider

- $r = 5\%$, $\delta = 10\%$, $\beta = 20\%$, $K = 100$,
- $T = 3$, $N = 9$, $t_n = \frac{nT}{N}$, $\xi = (100, \ldots, 100) \in \mathbb{R}^d$,
- $\mu \colon \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma \colon \mathbb{R}^d \to \mathbb{R}^{d \times d}$ satisfying
  $\mu(x) = (r - \delta)\, x$ and $\sigma(x) = \beta \operatorname{diag}(x_1, \ldots, x_d)$,
- $\mathcal{X} = (\mathcal{X}^{(1)}, \ldots, \mathcal{X}^{(d)}) \colon \{0, 1, \ldots, N\} \times \Omega \to \mathbb{R}^d$ satisfying

$$\mathcal{X}_n^{(i)} = X_{t_n}^{(i)} = \exp\!\left(\left[r - \delta - \tfrac{1}{2}\beta^2\right] t_n + \beta\, W_{t_n}^{(i)}\right) \xi_i,$$

- $g \colon [0, T] \times \mathbb{R}^d \to \mathbb{R}$ satisfying

$$g(t, x) = e^{-rt} \max\big\{\max\{x_1, \ldots, x_d\} - K, 0\big\}.$$

We approximate the price

$$\sup\left\{\mathbb{E}\left[e^{-r\tau} \max\big\{\max\{X_\tau^{(1)}, \ldots, X_\tau^{(d)}\} - K, 0\big\}\right] \colon \begin{array}{c}\tau\colon \Omega \to \{t_0, t_1, \ldots, t_N\} \text{ is an} \\ (\mathbb{F}_t^X)_{t \in \{t_0, t_1, \ldots, t_N\}}\text{-stopping time}\end{array}\right\}.$$

Numerical results based on 10 independent runs:

	Mean	Std. dev.	Avg. runtime	Reference	
$d = 3$	18.687	0.006	32.6	18.69	$(\star)$
$d = 5$	26.144	0.007	32.9	$[26.115, 26.164]$	$(\triangle)$
$d = 10$	38.278	0.010	35.5	–	
$d = 20$	51.569	0.007	42.6	–	
$d = 30$	59.521	0.009	49.2	–	
$d = 50$	69.574	0.011	63.3	–	
$d = 100$	83.386	0.008	101.2	–	
$d = 200$	97.411	0.010	179.1	–	
$d = 500$	116.249	0.013	507.4	–	

$(\star)$ binomial value in Andersen & Broadie 2004,

$(\triangle)$ 95% confidence interval in Broadie & Cao 2008

5000-dimensional Bermudan max-call option in the Black Scholes model

GPU: Nvidia Tesla P100 @ 1328 MHz, 3584 cuda cores, 16 GB HBM2 @ 1408 MHz,
CPU: 2 x Intel Xeon @ 2.4 GHz, 12 cores, 56 GB DDR4

	Price	Rel. err. (2000 steps)	Runtime (2000 steps)
$d = 5000$	165.430	0.013 %	1071.9

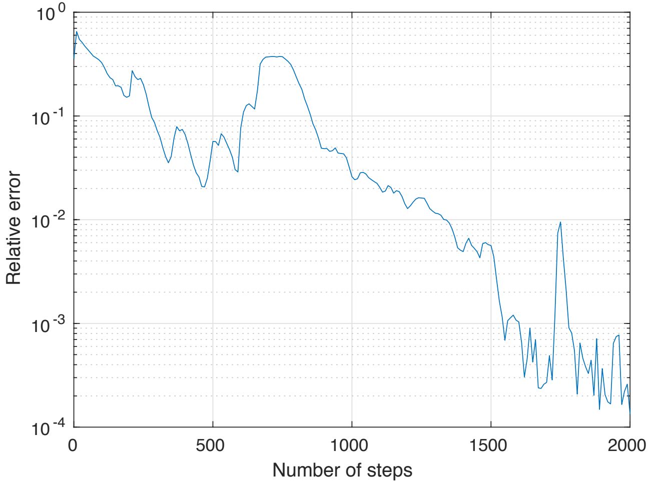

Schweizerische Studienstiftung
Fondation suisse d'études
Fondazione Svizzera degli Studi
Fundaziun svizra da studis
Swiss Study Foundation

Your questions are welcome!