# Longstaff Schwartz algorithm and Neural Network regression

#### Jérôme Lelong (joint work with B. Lapeyre)

Univ. Grenoble Alpes

Advances in Financial Mathematics 2020

# Introduction

- Computing an American option involving a large number of assets remains numerically challenging.
- A hope: Neural Network (NN) can (may) help to reduce the computational burden.
- Some previous works using NN for optimal stopping (not LS algorithm though)
  - Michael Kohler, Adam Krzyżak, and Nebojsa Todorovic. Pricing of high-dimensional american options by neural networks.
     Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics, 20(3):383–410, 2010
  - S. Becker, P. Cheridito, and A. Jentzen. Deep optimal stopping. Journal of Machine Learning Research, 20(74):1–25, 2019

### Computing Bermudan options prices

- ► A discrete time (discounted) payoff process  $(Z_{T_k})_{0 \le k \le N}$  adapted to  $(\mathcal{F}_{T_k})_{0 \le k \le N}$ .  $\max_{0 \le k \le N} |Z_{T_k}| \in L^2$ .
- The time- $T_k$  discounted value of the Bermudan option is given by

$$U_{T_k} = \operatorname{esssup}_{\tau \in \mathcal{T}_{T_k}} \mathbb{E}[Z_{\tau} | \mathcal{F}_{T_k}]$$

where  $\mathcal{T}_t$  is the set of all  $\mathcal{F}$ - stopping times with values in  $\{T_k, T_{k+1}, ..., T\}$ .

► From the Snell enveloppe theory, we derive the standard dynamic programming algorithm (→ "Tsistsiklis-Van Roy" type algorithms).

(1) 
$$\begin{cases} U_{T_N} = Z_{T_N} \\ U_{T_k} = \max\left(Z_{T_k}, \mathbb{E}[U_{T_{k+1}} | \mathcal{F}_{T_k}]\right) \end{cases}$$

### The policy iteration approach...

Let  $\tau_k$  be the smallest optimal stopping time after  $T_k$ .

(2) 
$$\begin{cases} \tau_N = T_N \\ \tau_k = T_k \mathbf{1}_{\left\{Z_{T_k} \geq \mathbb{E}[Z_{\tau_{k+1}} | \mathcal{F}_{T_k}]\right\}} + \tau_{k+1} \mathbf{1}_{\left\{Z_{T_k} < \mathbb{E}[Z_{\tau_{k+1}} | \mathcal{F}_{T_k}]\right\}}. \end{cases}$$

This is a dynamic programming principle on the policy not on the value function  $\rightarrow$  "Longstaff-Schwartz" algorithm. This approach has the practitioners' favour for its robustness.

Difficulty: how to compute the conditional expectations?

### ... in a Markovian context

▶ Markovian context:  $(X_t)_{0 \le t \le T}$  is a Markov process and  $Z_{T_k} = \phi_k(X_{T_k})$ .

$$\mathbb{E}[Z_{\tau_{k+1}}|\mathcal{F}_{T_k}] = \mathbb{E}[Z_{\tau_{k+1}}|X_{T_k}] = \psi_k(X_{T_k})$$

where  $\psi_k$  is a measurable function.

Because of the  $L^2$  assumption,  $\psi_k$  can be computed by a least-square problem

$$\inf_{\psi \in L^2(\mathcal{L}(X_{T_k}))} \mathbb{E}\left[\left|Z_{ au_{k+1}} - \psi(X_{T_k})
ight|^2
ight]$$

### Different numerical strategies

- The standard numerical (LS) approach: approximate the space L<sup>2</sup> by a finite dimensional vector space (polynomials, ...)
- We investigate the use of Neural Networks to approximate  $\psi_k$ .
- Kohler et al. [2010]: neural networks but in a different context (approximation of the value function Tsitsiklis and Roy [2001], equation (1)) and re-simulation of the paths at each time steps.

The Longstaff Schwartz algorithm

### LS: truncation step

Longstaff-Schwartz type algorithms rely on direct approximation of *stopping times* and use of *the same simulated paths* for all time steps (obvious and large computational gains).

- $(g_k, k \ge 1)$  is an  $L^2(\mathcal{L}(X))$  basis and  $\Phi_p(X, \theta) = \sum_{k=1}^p \theta_k g_k(X)$ .
- Backward approximation of iteration policy using (2),

$$\begin{cases} \widehat{\tau}_N^{p,} = T_N \\ \widehat{\tau}_n^p = T_n \, \mathbf{1}_{\left\{Z_{T_n} \ge \Phi_p(X_{T_n}; \widehat{\theta}_n^p)\right\}} + \widehat{\tau}_{n+1}^p \, \mathbf{1}_{\left\{Z_{T_n} < \Phi_p(X_{T_n}; \widehat{\theta}_n^p)\right\}} \end{cases}$$

 with conditional expectation computed using a Monte Carlo minimization problem: θ<sup>p</sup><sub>n</sub> is a minimizer of

$$\inf_{\theta} \mathbb{E} \left( \left| \Phi_p(X_{T_n}; \theta) - Z_{\widehat{\tau}_{n+1}^p} \right|^2 \right).$$

• Price approximation:  $U_0^p = \max\left(Z_0, \mathbb{E}\left(Z_{\widehat{\tau}_1^p}\right)\right).$ 

The Longstaff Schwartz algorithm 0000000000

### The LS algorithm

- (g<sub>k</sub>, k ≥ 1) is an L<sup>2</sup>(L(X)) basis and Φ<sub>p</sub>(X, θ) = Σ<sup>p</sup><sub>k=1</sub> θ<sub>k</sub> g<sub>k</sub>(X).
   Paths X<sup>(m)</sup><sub>T0</sub>, X<sup>(m)</sup><sub>T1</sub>,..., X<sup>(m)</sup><sub>TN</sub> and payoff paths Z<sup>(m)</sup><sub>T0</sub>, Z<sup>(m)</sup><sub>T1</sub>,..., Z<sup>(m)</sup><sub>TN</sub>, m = 1,..., M.
- Backward approximation of iteration policy,

$$\begin{cases} \widehat{\tau}_{N}^{p,(m)} = T_{N} \\ \widehat{\tau}_{n}^{p,(m)} = T_{n} \mathbf{1}_{\left\{ Z_{T_{n}}^{(m)} \ge \Phi_{p}(X_{T_{n}}^{(m)};\widehat{\theta}_{n}^{p,M}) \right\}} + \widehat{\tau}_{n+1}^{p,(m)} \mathbf{1}_{\left\{ Z_{T_{n}}^{(m)} < \Phi_{p}(X_{T_{n}}^{(m)};\widehat{\theta}_{n}^{p,M}) \right\}} \end{cases}$$

with conditional expectation computed using a Monte Carlo minimization problem: 
 *θ*<sub>n</sub><sup>p,M</sup> is a minimizer of

$$\inf_{\theta} \frac{1}{M} \sum_{m=1}^{M} \left| \Phi_p(X_{T_n}^{(m)}; \theta) - Z_{\tau_{n+1}^{p,(m)}}^{(m)} \right|^2.$$

• Price approximation: 
$$U_0^{p,M} = \max\left(Z_0, \frac{1}{M}\sum_{m=1}^M Z_{\widehat{\tau}_1^{p,(m)}}^{(m)}\right)$$

# Reference papers

#### • Description of the algorithm:

F.A. Longstaff and R.S. Schwartz. Valuing American options by simulation : A simple least-square approach. *Review of Financial Studies*, 14:113–147, 2001.

#### Rigorous approach:

Emmanuelle Clément, Damien Lamberton, and Philip Protter. An analysis of a least squares regression method for american option pricing.

Finance and Stochastics, 6(4):449-471, 2002.

- $U^p_0$  converge to  $U_0, p 
  ightarrow +\infty$
- $U_0^{p,M}$  converge to  $U_0^p, M \to +\infty$  a.s.
- "almost" a central limit theorem

# The modified algorithm

- In LS algorithm replace the approximation on a Hilbert basis Φ<sub>p</sub>(.; θ) by a Neural Network. This is not a vector space approximation (non linear).
- ▶ The optimization problem is non linear, non convex, ...
- Aim: extending the proof of (a.s.) convergence results

### A quick view of Neural Networks

- ▶ In short, a NN:  $x \to \Phi_p(x, \theta) \in \mathbb{R}$ , with  $\theta \in \mathbb{R}^d$ , d large
- $\bullet \ \Phi_p = A_L \circ \sigma_a \circ A_{L-1} \circ \cdots \circ \sigma_a \circ A_1, L \ge 2$
- $A_l(x_l) = w_l x_l + \beta_l$  (affine functions)
- ► L 2 "number of hidden layers"
- $\triangleright$  *p* "maximum number of neurons per layer" (i.e. sizes of the  $w_l$  matrix)
- $\sigma_a$  a fixed non linear (called *activation function*) applied component wise
- $\theta := (w_l, \beta_l)_{l=1,...,L}$  parameters of all the layers
- Restriction to a compact set Θ<sub>p</sub> = {θ : |θ| ≤ γ<sub>p</sub>} and assume lim<sub>p→∞</sub> γ<sub>p</sub> = ∞. → use the USLLN.

$$\blacktriangleright \mathcal{NN}_p = \{ \Phi_p(\cdot, \theta) : \theta \in \Theta_p \} \text{ and } \mathcal{NN}_\infty = \cup_{p \in \mathbb{N}} \mathcal{NN}_p$$

The Longstaff Schwartz algorithm

# Hypothesis H

For every p, there exists  $q \ge 1$ 

$$\forall \theta \in \Theta_p, \quad |\Phi_p(x,\theta)| \le \kappa_q (1+|x|^q)$$

a.s. the random function  $\theta \in \Theta_p \longmapsto \Phi_p(X_{T_n}, \theta)$  are continuous.

$$\inf_{ heta \in \Theta_p} \mathbb{E} \left( \left| \Phi_p(X_{T_n}; heta) - Z_{\widehat{ au}_{n+1}}^p \right|^2 
ight),$$

then  $\Phi_p(x, \theta_1) = \Phi_p(x, \theta_2)$  for almost all x No need of a unique minimizer but only of the represented function.

### The result

### **Theorem 1**

#### Under hypothesis H

Convergence of the Neural network approximation

$$\lim_{p\to\infty}\mathbb{E}[Z_{\tau^p_n}|\mathcal{F}_{T_n}]=\mathbb{E}[Z_{\tau_n}|\mathcal{F}_{T_n}] \text{ in } L^2(\Omega) \quad (i.e. \ U_0^p\to U_0).$$

SLLN: for every 
$$k = 1, \ldots, N$$
,

$$\lim_{M \to \infty} \frac{1}{M} \sum_{m=1}^{M} Z_{\widehat{\tau}_{k}^{p,(m)}}^{(m)} = \mathbb{E} \left[ Z_{\tau_{k}^{p}} \right] \quad a.s. \quad (i.e. \ U_{0}^{p,M} \to U_{0}^{p})$$

## Convergence of the NN approximation

A simple consequence of Hornik [1991].

Also known as the "Universal Approximation Theorem".

### **Theorem 2 (Hornik)**

Assume that the function  $\sigma_a$  is non constant and bounded. Let  $\mu$  denote a probability measure on  $\mathbb{R}^r$ , then  $\mathcal{NN}_{\infty}$  is dense in  $L^2(\mathbb{R}^r, \mu)$ .

► Corollary: If for every  $p, \alpha_p \in \Theta_p$  is a minimizer of

$$\inf_{\theta \in \Theta_p} \mathbb{E}[|\Phi_p(X;\theta) - Y|^2],$$

 $(\Phi_p(X; \alpha_p))_p$  converges to  $\mathbb{E}[Y|X]$  in  $L^2(\Omega)$  when  $p \to \infty$ .

▶ proof of the convergence of the "non-linear approximation"  $\Phi_p(X; \theta)$ .

### Convergence of Monte-Carlo approximation

- ▶ *p* is fixed,  $N \to +\infty$
- Now, minimisation problems are non linear, need more abstract arguments to prove convergence
- Two ingredients (quite "old" results)
- First result: approximation of minimization problems

### Lemma 3 (Rubinstein and Shapiro [1993])

▶ 
$$(f_n)_n$$
 defined on a compact set  $K \subset \mathbb{R}^d$ .  $v_n = \inf_{x \in K} f_n(x)$ 

•  $x_n$  a sequence of minimizers  $f_n(x_n) = \inf_{x \in K} f_n(x)$ .

• 
$$v^* = \inf_{x \in K} f(x)$$
 and  $\mathcal{S}^* = \{x \in K : f(x) = v^*\}.$ 

If  $(f_n)_n$  converges uniformly on K to a continuous function f, then  $v_n \to v^*$  and  $d(x_n, S^*) \to 0$  a.s.

### Convergence of Monte-Carlo approximation

Second result: SLLN in Banach spaces (Ledoux and Talagrand [1991], goes back to Mourier [1953]).

#### Lemma 4

Let 
$$(\xi_i)_{i\geq 1}$$
 i.i.d.  $\mathbb{R}^m$ -valued,  $h : \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}$ . If  
•  $a.s., \theta \in \mathbb{R}^d \mapsto h(\theta, \xi_1)$  is continuous,  
•  $\forall K > 0, \mathbb{E}\left[\sup_{|\theta| \leq K} |h(\theta, \xi_1)|\right] < +\infty.$   
Then  

$$\lim_{n \to \infty} \sup_{|\theta| \leq K} \left| \frac{1}{n} \sum_{i=1}^n h(\theta, \xi_i) - \mathbb{E}[h(\theta, \xi_1)] \right| = 0 \quad a.s.$$

## Convergence of Monte-Carlo approximation

Combining these two results with the backward iteration introduced by Clément et al. [2002], we get

### **Proposition**

Under hypothesis **H**, for every n = 1, ..., N,  $\Phi(X_{T_n}^{(1)}; \hat{\theta}_n^{p,M})$  converges to  $\Phi_p(X_{T_n}^{(1)}; \theta_n^p)$  a.s. when  $M \to \infty$ .

### Implementation details

- Python code with TensorFlow.
- We use ADAM algorithm to fit the neural network at each time step.
- We use the same NN through all time steps: take  $\hat{\theta}_{n+1}^{p,M}$ , as the starting point of the training algorithm at time time *n*.
- ► No use of setting *epochs*> 1 for *n* < *N* − 1. This allows for huge computational time savings.
- We only use the in-the-money paths

$$\inf_{\theta\in\Theta_p}\mathbb{E}\left[\left|\Phi_p(X_{T_n};\theta)-Z_{\tau_{n+1}^p}\right|^2\mathbf{1}_{\{Z_{T_n}\}}>0\right].$$

### Put Basket option in the Black Scholes model

L	$d_l$	epochs=1	epochs=5	epochs=10
2	32	$4.08~(\pm 0.031)$	$4.1 (\pm 0.034)$	4.11 (± 0.029)
2	128	$4.08~(\pm 0.036)$	$4.09 (\pm 0.034)$	4.1 (± 0.032)
2	512	$4.07~(\pm 0.034)$	4.09 (± 0.036)	4.1 (± 0.033)
4	32	$4.07~(\pm 0.034)$	4.09 (± 0.033)	4.1 (± 0.032)
4	128	$4.06~(\pm 0.039)$	$4.09~(\pm 0.04)$	$4.1~(\pm 0.037)$
4	512	$4.05~(\pm 0.037)$	$4.08~(\pm 0.034)$	$4.09 (\pm 0.031)$
8	32	$4.07~(\pm 0.034)$	$4.09~(\pm 0.037)$	4.1 (± 0.035)
8	128	$4.06~(\pm 0.039)$	$4.09~(\pm 0.032)$	$4.1~(\pm 0.035)$
8	512	$4.04~(\pm 0.066)$	$4.07~(\pm 0.069)$	$4.08~(\pm 0.063)$

Table: Basket option with r = 0.05, d = 5,  $\sigma^i = 0.2$ ,  $\omega^i = 1/d$ ,  $S_0^i = 100$ ,  $\rho = 0.2$ , K = 100, N = 20 and M = 100,000. The standard Longstaff Schwartz algorithm yields  $4.11 \pm 0.03$  (resp.  $4.04 \pm 0.034$ ) for an order 3 (resp. 1) polynomial regression.

### Put option in the Heston model

L	$d_l$	epochs=1	epochs=5	epochs=10
2	32	$1.69 (\pm 0.017)$	$1.7~(\pm 0.017)$	$1.7~(\pm 0.016)$
2	128	$1.69 (\pm 0.017)$	$1.7~(\pm 0.019)$	$1.7~(\pm 0.019)$
2	512	$1.69 (\pm 0.019)$	$1.69~(\pm 0.019)$	$1.69~(\pm 0.018)$
4	32	$1.69 (\pm 0.022)$	$1.69~(\pm 0.017)$	$1.7~(\pm 0.018)$
4	128	$1.69 (\pm 0.024)$	$1.69~(\pm 0.02)$	$1.7~(\pm 0.016)$
4	512	$1.68 (\pm 0.025)$	$1.69 (\pm 0.022)$	$1.69 (\pm 0.022)$
8	32	$1.69 (\pm 0.023)$	$1.69~(\pm 0.02)$	$1.69~(\pm 0.019)$
8	128	$1.68~(\pm 0.03)$	$1.69 (\pm 0.022)$	$1.69~(\pm 0.02)$
8	512	$1.68~(\pm 0.03)$	$1.68~(\pm 0.041)$	$1.68~(\pm 0.053)$

Table: Prices for put option in the Heston model with parameters the geometric basket put option with parameters with  $S_0 = K = 100$ , T = 1,  $\sigma_0 = 0.01$ ,  $\xi = 0.2$ ,  $\theta = 0.01$ ,  $\kappa = 2$ ,  $\rho = -0.3$ , r = 0.1, N = 10 and M = 100,000. The standard Longstaff Schwartz algorithm yields  $1.70 \pm 0.008$  (resp.  $1.675 \pm 0.005$ ) for an order 6 (resp. 1) polynomial regression.

# Conclusion

- Learn the continuation value using a NN instead of a polynomial regression
- NN do not help much for low dimensional problems but do scale far better
- Relatively small NN provide very accurate results (a few hundred neurons with 1 or 2 hidden layers)
- Setting epochs = 1 is fine for all dates but the last one.
- NN have proved to be a very versatile and efficient tool to compute Bermudan option prices...
- ... but keep in mind that using large NN is not green!

# Bibliography (1)

- S. Becker, P. Cheridito, and A. Jentzen. Deep optimal stopping. Journal of Machine Learning Research, 20(74):1–25, 2019.
- Emmanuelle Clément, Damien Lamberton, and Philip Protter. An analysis of a least squares regression method for american option pricing. *Finance and Stochastics*, 6(4):449–471, 2002.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. Neural Networks, 4(2):251 – 257, 1991.
- Côme Huré, Huyên Pham, Achref Bachouch, and Nicolas Langrené. Deep neural networks algorithms for stochastic control problems on finite horizon, part i: convergence analysis. *arXiv preprint arXiv:1812.04300*, 2018.
- Michael Kohler, Adam Krzyżak, and Nebojsa Todorovic. Pricing of high-dimensional american options by neural networks. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 20(3):383–410, 2010.
- Michel Ledoux and Michel Talagrand. Probability in Banach spaces, volume 23 of Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1991. ISBN 3-540-52013-9. Isoperimetry and processes.
- F.A. Longstaff and R.S. Schwartz. Valuing American options by simulation : A simple least-square approach. *Review of Financial Studies*, 14:113–147, 2001.

# Bibliography (2)

- Edith Mourier. Eléments aléatoires dans un espace de Banach. Ann. Inst. H. Poincaré, 13: 161–244, 1953.
- Reuven Y. Rubinstein and Alexander Shapiro. Discrete event systems. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, 1993. ISBN 0-471-93419-4. Sensitivity analysis and stochastic optimization by the score function method.
- J.N. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *IEEE Trans. Neural Netw.*, 12(4):694–703, 2001.