LisMath Seminar 2020
Spring Semester

Machine Learning Driven Optimal Stopping

Afonso Moniz Moreira

Instituto Superior Técnico

22th April

1. Classic Optimal Stopping Problems
   - General Problem and Free-Boundary Solution
   - Example: Perpetual American Call

2. Machine Learning Optimal Stopping
   - Deep Optimal Stopping - DOS

Intuition

What is an Stopping Problem ?

- One assumes the world behaves in a certain way.
  - Dynamics of the object of interest
  - Performance Measure

Intuition

What is an Stopping Problem ?

- One assumes the world behaves in a certain way.
  - Dynamics of the object of interest
  - Performance Measure
- The right time to take an action backed by mathematics.

Intuition

What is an Stopping Problem ?

- One assumes the world behaves in a certain way.
  - Dynamics of the object of interest
  - Performance Measure
- The right time to take an action backed by mathematics.
- Before we take an action we have an option.

Intuition

What is an Stopping Problem ?

- One assumes the world behaves in a certain way.
  - Dynamics of the object of interest
  - Performance Measure
- The right time to take an action backed by mathematics.
- Before we take an action we have an option.
- In Real Options Valuation Theory:
  - This option has a specific value.

Intuition

What is an Stopping Problem ?

- One assumes the world behaves in a certain way.
  - Dynamics of the object of interest
  - Performance Measure
- The right time to take an action backed by mathematics.
- Before we take an action we have an option.
- In Real Options Valuation Theory:
  - This option has a specific value.
  - The correct choice has to take this into account.

Object Dynamics

Fundamental Concepts

- We consider a Filtered Probability Space:

$$(\Omega, (\mathcal{F})_{t\geq 0}, \mathbb{P}_x) \tag{1}$$

- And a given measurable space:

$$(\mathbb{R}, \mathcal{B}(\mathbb{R})) \tag{2}$$

- Following Oksendal [2013] (p.10):

Definition 1

A Stochastic Process is a parametrized collection of random variables $\{X_t\}_{t\in T}$ defined on the filtered probability space (1) and taking values on the measurable space (18).

$$X : \Omega \times [0, +\infty) \longrightarrow \mathbb{R} \tag{3}$$

Object Dynamics

Fundamental Concepts

- Geometric Browian Motion - Itô Diffusion

$$\begin{cases} dX_t = \mu X_t dt + \sigma X_t dW_t \\ X_0 = x \end{cases} \iff \begin{cases} \frac{dX_t}{dt} = \mu X_t + \sigma X_t \underbrace{N_t}_{\frac{dW_t}{dt}} \\ X_0 = x \end{cases} \tag{4}$$

- μ, $\sigma \in \mathbb{R}^+$ - Fixed Parameters
- W_t - Standard Wiener Process - Standard Brownian Motion
- N_t - White Noise
- Time-Homogeneous Markov Process which means:
  - Initial Value
  - Transition Density

Performance Measure

The Gain Function

- In with Fleming and Soner [2006] and Peskir and Shiryaev [2006]:

$$G(X_\tau) = \underbrace{\int_0^\tau e^{-\mu t} L(X_t) dt}_{\text{Running Revenue}} + \underbrace{e^{-\mu \tau} M(X_\tau)}_{\text{Terminal Revenue}} \tag{5}$$

- The gain function can either reflect a reward or a cost.
- Each component of (5) is "killed" at the rate r - why?
  - Time value of money
  - Feynman-Kac formula eases calculations

Optimal Stopping Problems

Problem Formulation

- Formally one wants to solve the following problem - Peskir and Shiryaev [2006]:

$$V(x) = \sup_{\tau \in [0,+\infty)} \left\{ \mathbb{E}\left[\underbrace{\underbrace{\int_0^{\tau} e^{-\mu t} L(X_t) dt}_{\text{Running Revenue}} + \underbrace{e^{-\mu \tau} M(X_\tau) \Big| X_0 = x}_{\text{Terminal Revenue}}}_{G(X_\tau)} \right] \right\} \tag{6}$$

- If the optimal stopping time τ^* exists then (6) is given by:

$$V(x) = \mathbb{E}\left[\int_0^{\tau^*} e^{-\mu t} L(X_t) dt + e^{-\mu \tau^*} M(X_{\tau^*}) \Big| X_0 = x \right] \tag{7}$$

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$
  - Dividing the state space $\mathbb{R}$ into continuation and stopping region:

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$
  - Dividing the state space $\mathbb{R}$ into continuation and stopping region:

$$C = \{x \in \mathbb{R} : V(x) > G(x)\} \tag{8}$$

$$D = \{x \in \mathbb{R} : V(x) = G(x)\} \tag{9}$$

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$
  - Dividing the state space $\mathbb{R}$ into continuation and stopping region:

$$C = \{x \in \mathbb{R} : V(x) > G(x)\} \tag{8}$$

$$D = \{x \in \mathbb{R} : V(x) = G(x)\} \tag{9}$$

- The optimal stopping time τ^* is defined as such:

$$\tau^* = \inf\{t \in [0, +\infty] : X_t \in D\} = \inf\{t \in [0, +\infty] : X_t \notin C\} \tag{10}$$

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$
  - Dividing the state space $\mathbb{R}$ into continuation and stopping region:

$$C = \{x \in \mathbb{R} : V(x) > G(x)\} \tag{8}$$

$$D = \{x \in \mathbb{R} : V(x) = G(x)\} \tag{9}$$

- The optimal stopping time τ^* is defined as such:

$$\tau^* = \inf\{t \in [0, +\infty] : X_t \in D\} = \inf\{t \in [0, +\infty] : X_t \notin C\} \tag{10}$$

- The optimal stopping time τ^* comes free of charge

Optimal Stopping Problems

Solving the Problem

- Solving problem (6) is doing two things:
  - Finding the value function $V(x)$
  - Dividing the state space $\mathbb{R}$ into continuation and stopping region:

$$C = \{x \in \mathbb{R} : V(x) > G(x)\} \tag{8}$$

$$D = \{x \in \mathbb{R} : V(x) = G(x)\} \tag{9}$$

- The optimal stopping time τ^* is defined as such:

$$\tau^* = \inf\{t \in [0, +\infty] : X_t \in D\} = \inf\{t \in [0, +\infty] : X_t \notin C\} \tag{10}$$

- The optimal stopping time τ^* comes free of charge
- $\implies$ stopping time τ is random whereas τ^* is not

Optimal Stopping Problems

Free-Boundary Approach

- In line with Peskir and Shiryaev [2006], for problem formulation (6):

$$V(x) = \mathbb{E}_x\left[\int_0^{\tau^*} e^{-\mu t} L(X_t) dt + e^{-\mu \tau^*} M(X_{\tau^*}) \right]$$

- The value function $V(x)$ solves the following FB problem:

$$\mathbb{L}V(x) - rV(x) + L(x) = 0, \ \ \forall x \in C \tag{11}$$

$$V(x) = M(x), \ \ \forall x \in D \tag{12}$$

$$\left.\frac{\partial V(x)}{\partial x}\right|_{x=\partial C} = \left.\frac{\partial M(x)}{\partial x}\right|_{x=\partial C} \tag{13}$$

Optimal Stopping Problems

Free-Boundary Approach

- Where $\mathbb{L}V(x)$ represents the infinitesimal generator of process X_t:

$$\mathbb{L}V(x) = \lim_{t \to 0} \left[\frac{\mathbb{E}[V(X_t)|X_0 = x] - V(x)}{t} \right] = \frac{\partial V(x)}{\partial x} \mu x + \frac{1}{2} \frac{\partial^2 V(x)}{\partial x^2} \sigma^2 x^2 \quad (14)$$

Optimal Stopping Problems

Free-Boundary Approach

- Where $\mathbb{L}V(x)$ represents the infinitesimal generator of process X_t:

$$\mathbb{L}V(x) = \lim_{t \to 0} \left[\frac{\mathbb{E}[V(X_t)|X_0 = x] - V(x)}{t} \right] = \frac{\partial V(x)}{\partial x} \mu x + \frac{1}{2} \frac{\partial^2 V(x)}{\partial x^2} \sigma^2 x^2 \tag{14}$$

- $\max\{\mathbb{L}V(x) - \mu V(x) + L(x), V(x) - M(x)\}$ - HJB

Optimal Stopping Problems

Free-Boundary Approach

- Where $\mathbb{L}V(x)$ represents the infinitesimal generator of process X_t:

$$\mathbb{L}V(x) = \lim_{t \to 0} \left[\frac{\mathbb{E}[V(X_t)|X_0 = x] - V(x)}{t} \right] = \frac{\partial V(x)}{\partial x} \mu x + \frac{1}{2} \frac{\partial^2 V(x)}{\partial x^2} \sigma^2 x^2 \tag{14}$$

- $\max\{\mathbb{L}V(x) - \mu V(x) + L(x), V(x) - M(x)\}$ - HJB
- $V(x) = M(x)$ - Value Matching Condition

Optimal Stopping Problems

Free-Boundary Approach

- Where $\mathbb{L}V(x)$ represents the infinitesimal generator of process X_t:

$$\mathbb{L}V(x) = \lim_{t \to 0} \left[\frac{\mathbb{E}[V(X_t)|X_0 = x] - V(x)}{t} \right] = \frac{\partial V(x)}{\partial x} \mu x + \frac{1}{2} \frac{\partial^2 V(x)}{\partial x^2} \sigma^2 x^2 \tag{14}$$

- $\max\{\mathbb{L}V(x) - \mu V(x) + L(x), V(x) - M(x)\}$ - HJB
- $V(x) = M(x)$ - Value Matching Condition
- $\left.\frac{\partial V(x)}{\partial x}\right|_{x=\partial C} = \left.\frac{\partial M(x)}{\partial x}\right|_{x=\partial C}$ - Smooth Pasting Condition

Perpetual American Call Option

The Standard Example

- I have the right to **buy** at a given K price - *Strike* Price

Perpetual American Call Option

The Standard Example

- I have the right to **buy** at a given K price - *Strike* Price
- I can take the action to **buy** anywhere in time.

Perpetual American Call Option

The Standard Example

- I have the right to **buy** at a given K price - *Strike* Price
- I can take the action to **buy** anywhere in time.
- We want to answer the eternal question in Finance Theory...

Perpetual American Call Option

The Standard Example

- I have the right to **buy** at a given K price - *Strike* Price
- I can take the action to **buy** anywhere in time.
- We want to answer the eternal question in Finance Theory...
- What is the fair value of such a contract/derivative product/contingent claim

Perpetual American Call Option

The Standard Example

- I have the right to **buy** at a given K price - *Strike* Price
- I can take the action to **buy** anywhere in time.
- We want to answer the eternal question in Finance Theory...
- What is the fair value of such a contract/derivative product/contingent claim
- A rational investor is willing to pay up until the discounted value of the maximum he is expecting to profit...

Perpetual American Call Option

Problem Formalization

- Hence our measure of performance is given by:

$$V(x) = \sup_{\tau \in [0,+\infty)} \left\{ \mathbb{E}\left[\underbrace{\underbrace{e^{-r\tau} \max\{X_\tau - K, 0\} | X_0 = x}_{M(X_\tau) - \text{Terminal Revenue}}}_{G(X_\tau)} \right] \right\} \tag{15}$$

- The dynamics of the stock price X_t is given by the GBM:

$$\begin{cases} dX_t = \mu X_t dt + \sigma X_t dW_t \\ X_0 = x \end{cases} \tag{16}$$

Perpetual American Call Option

Problem Solution

Drawing I

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier
    - Numeric Solutions - Die in $\mathbb{R}^5$ according to Becker et al. [2019b]

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier
    - Numeric Solutions - Die in $\mathbb{R}^5$ according to Becker et al. [2019b]
    - For d-dimensional processes, in principle, one has to account for correlation, otherwise it is just d identical problems to solve.

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier
    - Numeric Solutions - Die in $\mathbb{R}^5$ according to Becker et al. [2019b]
    - For d-dimensional processes, in principle, one has to account for correlation, otherwise it is just d identical problems to solve.
    - Even with a numerical approach - How to guess the upper boundary ?

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier
    - Numeric Solutions - Die in $\mathbb{R}^5$ according to Becker et al. [2019b]
    - For d-dimensional processes, in principle, one has to account for correlation, otherwise it is just d identical problems to solve.
    - Even with a numerical approach - How to guess the upper boundary ?
  - The stopping time τ is a random variable - A function

Deep Optimal Stopping

Why do we need it ?

- Why should we use Machine Learning for Optimal Stopping?
  - The Dimensionality Monster
    - Explicit Solutions - Die in $\mathbb{R}$, with non-homogenous Markov Processes and variable upper frontier
    - Numeric Solutions - Die in $\mathbb{R}^5$ according to Becker et al. [2019b]
    - For d-dimensional processes, in principle, one has to account for correlation, otherwise it is just d identical problems to solve.
    - Even with a numerical approach - How to guess the upper boundary ?
  - The stopping time τ is a random variable - A function
  - According to Goodfellow et al. [2016] - Deep Neural Networks can be perceived as function approximation machines designed to achieve statistical generalization.

Deep Optimal Stopping

Main Goal

- Transform the classical problem into a deep learning problem.

Deep Optimal Stopping

Main Goal

- Transform the classical problem into a deep learning problem.
- Consider several sources of risk - $\mathbb{R}^d$ processes

Deep Optimal Stopping

Main Goal

- Transform the classical problem into a deep learning problem.
- Consider several sources of risk - $\mathbb{R}^d$ processes
- We follow Becker et al. [2019a] and Becker et al. [2019b].

Deep Optimal Stopping

Performance Measure

- Same Filtered Probability Space:

$$(\Omega, (\mathcal{F})_{t\geq 0}, \mathbb{P}_x) \tag{17}$$

- But in a d-dimensional State Space - d sources of risk

$$(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d)) \tag{18}$$

Deep Optimal Stopping

Performance Measure

- Same Filtered Probability Space:

$$(\Omega, (\mathcal{F})_{t\geq 0}, \mathbb{P}_x) \tag{17}$$

- But in a d-dimensional State Space - d sources of risk

$$(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d)) \tag{18}$$

- We face the same optimization problem as (6):

$$V(x) = \sup_{\tau\in[0,T]} \left\{ \mathbb{E}\Big[g(X_\tau) | X_0 = x \Big] \right\} = \mathbb{E}\Big[g(X_{\tau^*}) | X_0 = x \Big] \tag{19}$$

- $g : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ - continuous function

Deep Optimal Stopping

Object Dynamics

But this dynamics are given by a time-homogenous Itô process

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t \tag{20}$$

- $X : [0, T] \times \Omega \to \mathbb{R}^d$ - d-dimensional process
- $W : [0, T] \times \Omega \to \mathbb{R}^d$ - d-dimensional Wiener process
- $\mu : \mathbb{R}^d \to \mathbb{R}^d$ - drift function
- $\sigma : \mathbb{R}^d \to \mathbb{R}^d$ - volatility function

Deep Optimal Stopping

Discretization Process

- To discretize the process (20) we apply the Euler-Maruyama as in Maruyama [1955]:
- For a given $N \in \mathbb{N}$, we consider the mesh $t_0, t_1, ..., t_N \in [0, T]$

$$0 = t_0 < t_1 < ... < t_n = T \tag{21}$$

- Considering an approximation of the SDE (20):

$$\mathcal{X}_{t_{n+1}} = \mathcal{X}_{t_n} + \mu(\mathcal{X}_{t_n})(t_{n+1} - t_n) + \sigma(\mathcal{X}_{t_n})(W_{t_{n+1}} - W_{t_n}) \tag{22}$$

- $\mathcal{X} : [0, 1, ..., N] \times \Omega \to \mathbb{R}^d$

Deep Optimal Stopping

Discrete Optimization Problem

- Then one can state the following:

$$\sup_{\tau \in [0,T]} \left\{ \mathbb{E}\left[g(X_\tau) | X_0 = x \right] \right\} \approx \sup_{\tau \in \{0,1,\ldots,N\}} \left\{ \mathbb{E}\left[g(\mathcal{X}_\tau) | \mathcal{X}_0 = x \right] \right\} \tag{23}$$

- The approximation error depends on the mesh size.

Deep Optimal Stopping

From classic to a ML problem

- The whole procedure is based on Lemma 2.2 from Becker et al. [2019b]

Lemma 2 (Factorized stopping time)

Let $d, N \in \mathbb{N}$, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, let $\mathcal{X} : 0, 1..., N \times \Omega \to \mathbb{R}^d$ be a stochastic process, and let $\mathbb{F} = (\mathbb{F}_n)_{n \in \{0,1,\dots,N\}}$ be a filtration (i.e an increasing σ-algebra) generated by the process $\mathcal{X}$, then the following holds:

1. *for all Borel measurable functions $\mathbb{U}_n : (\mathbb{R}^d)^{n+1} \to \{0, 1\}, n \in 0, 1, \dots, N$, with $\forall x_0, x_1, \dots, x_N \in \mathbb{R}^d : \sum_{n=0}^{N} \mathbb{U}_n(x_0, x_1, \dots, x_N) = 1$ it holds that the following function is an $\mathbb{F}$-stopping time:*

$$\Omega \ni \omega \longrightarrow \sum_{n=0}^{N} n \times \mathbb{U}_n(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), \dots, \mathcal{X}_n(\omega)) \in (x_0, x_1, \dots, x_N) \tag{24}$$

2. *for every $\mathbb{F}$-stopping time $\tau : \Omega \to \{0, 1, \dots, N\}$ there exist Borel measurable functions $\mathbb{U}_n : (\mathbb{R}^d)^{n+1} \to \{0, 1\}$ which satisfy $\forall (x_0, x_1, \dots, x_N) \in \mathbb{R}^d : \sum_{n=0}^{N} \mathbb{U}_n(x_0, x_1, \dots, x_N) = 1$*

$$\tau = \sum_{n=0}^{N} n \times \mathbb{U}_n(\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_n) \tag{25}$$

Deep Optimal Stopping

From classic to a ML problem

- And so for every realization of the discrete process $\mathcal{X}_{t_n}(\omega)$:

$$\tau(\omega) = \sum_{n=0}^{N} \underbrace{n}_{\text{mesh point}} \times \underbrace{\mathbb{U}_{n,\tau}(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), ..., \mathcal{X}_n(\omega))}_{\text{binary function}} \tag{26}$$

Deep Optimal Stopping

From classic to a ML problem

- And so for every realization of the discrete process $\mathcal{X}_{t_n}(\omega)$:

$$\tau(\omega) = \sum_{n=0}^{N} \underbrace{n}_{\text{mesh point}} \times \underbrace{\mathbb{U}_{n,\tau}(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), ..., \mathcal{X}_n(\omega))}_{\text{binary function}} \tag{26}$$

- Additionally,

$$\sum_{n=0}^{N} \mathbb{U}_{n,\tau}(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), ..., \mathcal{X}_n(\omega)) = 1 \tag{27}$$

Deep Optimal Stopping

From classic to a ML problem

- And so for every realization of the discrete process $\mathcal{X}_{t_n}(\omega)$:

$$\tau(\omega) = \sum_{n=0}^{N} \underbrace{n}_{\text{mesh point}} \times \underbrace{\mathbb{U}_{n,\tau}(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), ..., \mathcal{X}_n(\omega))}_{\text{binary function}} \tag{26}$$

- Additionally,

$$\sum_{n=0}^{N} \mathbb{U}_{n,\tau}(\mathcal{X}_0(\omega), \mathcal{X}_1(\omega), ..., \mathcal{X}_n(\omega)) = 1 \tag{27}$$

- So obtaining a stopping time τ comes down to obtaining $\mathbb{U}_{n,\tau}$ for each mesh point n.

Deep Optimal Stopping

From classic to a ML problem

- Obtaining an optimal stopping time τ comes down to obtaining $\mathbb{U}_{n,\tau}$ by minimizing or maximizing a performance criteria for each mesh point n.

Deep Optimal Stopping

From classic to a ML problem

- Obtaining an optimal stopping time τ comes down to obtaining $\mathbb{U}_{n,\tau}$ by minimizing or maximizing a performance criteria for each mesh point n.
- As previously state the main goal rests in replacing the classic problem with a ML problem hence:

Deep Optimal Stopping

From classic to a ML problem

- Obtaining an optimal stopping time τ comes down to obtaining $\mathbb{U}_{n,\tau}$ by minimizing or maximizing a performance criteria for each mesh point n.
- As previously state the main goal rests in replacing the classic problem with a ML problem hence:
- In line with Becker et al. [2019a] to achieve this we will make use of a feed-forward neural network.

Deep Optimal Stopping

From classic to a ML problem

- Obtaining an optimal stopping time τ comes down to obtaining $\mathbb{U}_{n,\tau}$ by minimizing or maximizing a performance criteria for each mesh point n.
- As previously state the main goal rests in replacing the classic problem with a ML problem hence:
- In line with Becker et al. [2019a] to achieve this we will make use of a feed-forward neural network.
- We will replace the search over all $\mathbb{F}$-stopping times over a set of parameters $\theta \in \mathbb{R}^\nu \;\; \nu \in \mathbb{N}$ in order to minimize or maximize a performance criterion.

Deep Optimal Stopping

Deep Neural Network

- For every $n \in \{0, 1, ..., N\}$, $\theta \in \mathbb{R}^{\nu}$, let $u_{n,\theta} : \mathbb{R}^d \to (0, 1)$ and $U_{n,\theta} : (\mathbb{R}^d)^{n+1} \to (0, 1)$ be Borel measurable functions such that:

$$U_{n,\theta}(x_0, x_1, ..., N) = \underbrace{\max\{u_{n,\theta}, n + 1 - N\}}_{\text{Current Period}} \underbrace{\left[1 - \sum_{k=0}^{n-1} \mathbb{U}_{k,\theta}(x_0, x_1, ..., x_k)\right]}_{\text{Previous Periods}} \tag{28}$$

Deep Optimal Stopping

Deep Neural Network

- For every $n \in \{0, 1, ..., N\}$, $\theta \in \mathbb{R}^\nu$, let $u_{n,\theta} : \mathbb{R}^d \to (0,1)$ and $U_{n,\theta} : (\mathbb{R}^d)^{n+1} \to (0,1)$ be Borel measurable functions such that:

$$U_{n,\theta}(x_0, x_1, ..., N) = \underbrace{\max\{u_{n,\theta}, n+1-N\}}_{\text{Current Period}} \underbrace{\left[1 - \sum_{k=0}^{n-1} \mathbb{U}_{k,\theta}(x_0, x_1, ..., x_k)\right]}_{\text{Previous Periods}} \quad (28)$$

- $u_{n,\theta}$ is approximated using a Feed-Forward Neural Network:

$$u_{n,\theta} = \mathcal{L}_1 \circ A_{d,d}^{\theta,(2nd+n)(d+1)} \circ \mathcal{L}_d \circ A_{d,d}^{\theta,(2nd+n+1)(d+1)} \circ \mathcal{L}_d \circ A_{d,d}^{\theta,((2n+1)d+n+1)(d+1)} \quad (29)$$

Deep Optimal Stopping

Deep Neural Network

- Neuron activation functions - Multidimensional Logistic - $\mathcal{L}_k : \mathbb{R}^k \to \mathbb{R}^k$

$$\mathcal{L}_k(x) = \left(\frac{\exp(x_1)}{\exp(x_1)+1}, \frac{\exp(x_2)}{\exp(x_2)+1}, ..., \frac{\exp(x_k)}{\exp(x_k)+1} \right) \tag{30}$$

Deep Optimal Stopping

Deep Neural Network

- Neuron activation functions - Multidimensional Logistic - $\mathcal{L}_k : \mathbb{R}^k \to \mathbb{R}^k$

$$\mathcal{L}_k(x) = \left(\frac{\exp(x_1)}{\exp(x_1)+1}, \frac{\exp(x_2)}{\exp(x_2)+1}, \dots, \frac{\exp(x_k)}{\exp(x_k)+1} \right) \tag{30}$$

- Input aggregation with affine function - $A_{k,l}^{\theta,v}(x) : \mathbb{R}^l \to \mathbb{R}^k$ - $v \in \mathbb{N}_0, k, l \in \mathbb{N}$

$$A_{k,l}^{\theta,v}(x) = \begin{bmatrix} \theta_{v+1} & \theta_{v+2} & \dots & \theta_{v+l} \\ \theta_{v+l+1} & \theta_{v+l+2} & \dots & \theta_{v+2l} \\ \theta_{v+2l+1} & \theta_{v+2l+2} & \dots & \theta_{v+3l} \\ \vdots & \vdots & \vdots & \vdots \\ \theta_{v+(k-1)l+1} & \theta_{v+(k-1)l+2} & \dots & \theta_{v+kl} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_l \end{bmatrix} + \begin{bmatrix} \theta_{v+kl+1} \\ \theta_{v+kl+2} \\ \theta_{v+kl+3} \\ \vdots \\ \theta_{v+kl+k} \end{bmatrix} \tag{31}$$

Deep Optimal Stopping

Deep Neural Network Structure

Drawing II

Deep Optimal Stopping

Modified Performance Measure

- We still want to approximate:

$$\sup_{\tau \in [0,T]} \left\{ \mathbb{E}\left[g(\mathcal{X}_\tau) | \mathcal{X}_0 = x \right] \right\} \tag{32}$$

Deep Optimal Stopping

Modified Performance Measure

- We still want to approximate:

$$\sup_{\tau \in [0,T]} \left\{ \mathbb{E}\left[g(\mathcal{X}_\tau) | \mathcal{X}_0 = x \right] \right\} \tag{32}$$

- According to Becker et al. [2019b] the gain function $g(X_\tau)$ can be perceived has:

$$g(\mathcal{X}_\tau) = \left[\sum_{n=0}^{N} \mathbb{1}_{\{\tau=n\}} \right] g(\mathcal{X}_\tau) = \sum_{n=0}^{N} \mathbb{1}_{\{\tau=n\}} g(\mathcal{X}_n) = \sum_{n=0}^{N} \mathbb{U}_{n,\theta}(x_0, x_1, ..., x_n) g(\mathcal{X}_n) \tag{33}$$

Deep Optimal Stopping

Modified Performance Measure

- We still want to approximate:

$$\sup_{\tau \in [0,T]} \left\{ \mathbb{E}\left[g(\mathcal{X}_\tau) | \mathcal{X}_0 = x \right] \right\} \tag{32}$$

- According to Becker et al. [2019b] the gain function $g(X_\tau)$ can be perceived has:

$$g(\mathcal{X}_\tau) = \left[\sum_{n=0}^{N} \mathbb{1}_{\{\tau=n\}} \right] g(\mathcal{X}_\tau) = \sum_{n=0}^{N} \mathbb{1}_{\{\tau=n\}} g(\mathcal{X}_n) = \sum_{n=0}^{N} \mathbb{U}_{n,\theta}(x_0, x_1, ..., x_n) g(\mathcal{X}_n) \tag{33}$$

- We are then able to modify the original problem to:

$$\sup_{\theta \in \mathbb{R}^\nu} \left\{ \sum_{n=0}^{N} \mathbb{U}_{n,\theta}(x_0, x_1, ..., x_n) g(\mathcal{X}_n) \right\} \tag{34}$$

Deep Optimal Stopping

Stochastic Gradient Ascent

- To maximize (34) we make use of a stochastic gradient algorithm.
- Which yields a sequence of random parameter vectors defined as such:

$$\Theta_m = (\Theta_m^{(1)}, ..., \Theta_m^{(\nu)}) : \Omega \to \mathbb{R}^\nu \tag{35}$$

Deep Optimal Stopping

Stochastic Gradient Ascent

- To maximize (34) we make use of a stochastic gradient algorithm.
- Which yields a sequence of random parameter vectors defined as such:

$$\Theta_m = (\Theta_m^{(1)}, ..., \Theta_m^{(\nu)}) : \Omega \to \mathbb{R}^\nu \tag{35}$$

- Each mth element of the sequence Θ_m corresponds to the mth step of the algorithm.
- Each mth step of the algorithm results in random approximations of the local/global maximum points of (34).

This is the End...

Thank You

References

Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *Journal of Machine Learning Research*, 20, 2019a.

Sebastian Becker, Patrick Cheridito, Arnulf Jentzen, and Timo Welti. Solving high-dimensional optimal stopping problems using deep learning. *arXiv preprint arXiv:1908.01602*, 2019b.

Wendell H Fleming and Halil Mete Soner. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Hilbert J Kappen. An introduction to stochastic control theory, path integrals and reinforcement learning. In *AIP conference proceedings*, volume 887, pages 149–181. American Institute of Physics, 2007.

References

Gisiro Maruyama. Continuous Markov processes and stochastic equations. 4(1):48, 1955.

Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.

Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *journal of machine learning research*, 11(Nov):3137–3181, 2010.