Online Gradient Descent with Momentum

Dealing with Transaction Costs in Portfolio Optimization

E. Vittori, M. Bernasconi de Luca, F. Trovó, M. Restelli

IMI CORPORATE & INVESTMENT BANKING

INTESA SANPAOLO

POLITECNICO MILANO 1863

Table of contents

Introduction

Online Portfolio Optimization

Given M assets, decide what proportion to invest in each asset in a sequential manner

- $\mathbf{x}_t$ s.t. $\sum_{j=1}^{M} x_t^j = 1, x_t^j > 0 \Rightarrow \mathbf{x}_t \in \Delta_{M-1}$ is the portfolio allocation
- $\mathbf{r}_t = \left(\frac{p_{t,1}}{p_{t-1,1}}, \ldots, \frac{p_{t,M}}{p_{t-1,M}}\right)$ are the price relatives
- $t \leq 1$ day

Algorithm 1 Online Portfolio Optimization

1: Input M assets, set $\mathbf{x}_1 \leftarrow \frac{1}{M}\mathbf{1}$
2: **for** $t \in \{1, \ldots, T\}$ **do**
3: Select $\mathbf{x}_{t+1} \leftarrow$ selection policy
4: Observe $\mathbf{r}_{t+1}$ from the market
5: Calculate wealth $W_{T+1} = \prod_{t=1}^{T+1} \langle \mathbf{x}_t, \mathbf{r}_t \rangle$
6: **end for**

Regret guarantees

Modern Portfolio Optimization

Algorithm 2 Efficient Frontier (Markowitz, 1952)

1: Input M assets and a batch of historical data, set $\lambda > 0$
2: Calculate μ, Σ from the data
3: Select $\mathbf{x} \leftarrow \mathrm{argmax}_{\mathbf{x} \in \Delta_{M-1}} \mathbf{x}'\mu - \frac{\lambda}{2}\mathbf{x}'\Sigma\mathbf{x}$
4: Hold the allocation

- Assumption of stochastic stock dynamics
- Single-period portfolio
- No theoretical guarantees

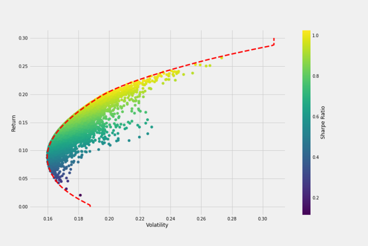

Online Learning with Experts Advice

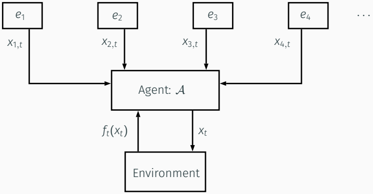

- Sequential decision problem
- Adversarial Environment
- Regret: $R_T = \sum_{t=1}^{T} f_t(x_t) - \inf_{e \in \mathcal{E}} \sum_{t=1}^{T} f_t(x_{e,t})$
  - No regret: $R_T = o(T)$ for any sequence $f_1, f_2, \ldots$
- Per-round Computational Complexity

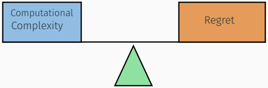

Halving Algorithm

Predict a sequence of heads or tails $y_t = \{H, T\}$, assume there are N experts e_i and one of them predicts perfectly the sequence.

Algorithm 2 Halving Algorithm

Initialize a weight for each expert $w_{i,t} = 1$

for $t \in \{1, \dots, T\}$ **do**

  Select $\mathbf{x}_{t+1} \leftarrow \begin{cases} H & \text{if majority of experts with } w_{i,t} = 1 \text{ predicts so} \\ T & \text{otherwise} \end{cases}$

  Observe y_{t+1}

  If $x_{t+1} \neq y_{t+1}$, set $w_{i,t} = 0$ for each $e_{i,t+1} \neq y_{t+1}$

end for

Halving Algorithm Regret

Predict a sequence of heads or tails $y_t = \{H, T\}$, assume there are N experts e_i and one of them predicts perfectly the sequence.

$R_T = \sum_{t=1}^{T} \mathbb{1}(x_t, y_t) - \mathbb{1}(e_t^*, y_t) = m$ where m is the number of mistakes

Proof:

Define $W_t = \sum_i w_{i,t}$
At $t = 0$, $m = 0$, $W_0 = N$
At each **mistake** $W_m \leq \frac{W_{m-1}}{2}$
Recursively $W_m \leq \frac{W_0}{2^m} = \frac{N}{2^m}$
Since at least one expert is perfect: $1 \geq \frac{N}{2^m} \Rightarrow m \leq \lfloor \log_2 N \rfloor$

Online Portfolio Optimization

From Expert Learning to Online Portfolio Optimization

- The experts are Constant Rebalancing Portfolios (CRP)
- $\mathbf{x}^* = \arg\inf_{x \in \Delta_{M-1}} \sum_{t=1}^{T} f_t(\mathbf{x})$ is the Best CRP
- $f_t(\mathbf{x}) = -\log(\langle \mathbf{x}, \mathbf{r}_t \rangle)$ is the loss

Regret in Online Portfolio Optimization

$$R_T = \log(W_T(\mathbf{x}^*, \ldots, \mathbf{x}^*)/W_T(\mathbf{x}_1, \ldots, \mathbf{x}_T)),$$

where $W_T(\mathbf{x}_1, \ldots, \mathbf{x}_T) = \prod_{t=1}^{T} \langle \mathbf{x}_t, \mathbf{r}_t \rangle$ is the wealth

Limitations: No transaction costs

Universal Portfolios (UP)

Algorithm 3 Universal Portfolios (Cover and Ordentlich, 1996)

1: Input M assets, set $\mathbf{x}_1 \leftarrow \frac{1}{M}\mathbf{1}$, initialize $\mathbf{W}_1$
2: **for** $t \in \{1, \ldots, T\}$ **do**
3: Select $\mathbf{x}_{t+1} \leftarrow \frac{\int_{\mathbf{b} \in \Delta_{M-1}} \mathbf{b} W_t(\mathbf{b}) \mathrm{d}\mu(\mathbf{b})}{\int_{\mathbf{b} \in \Delta_{M-1}} W_t(\mathbf{b}) \mathrm{d}\mu(\mathbf{b})}$
4: Observe $\mathbf{r}_{t+1}$ from the market
5: $W_{t+1} = \prod_{t=1}^{t+1} \langle \mathbf{r}_{t+1}, \mathbf{x}_{t+1} \rangle$
6: **end for**

Regret

$$R_T \leq (M-1)\log(T+1)$$

Online Newton Step (ONS)

Algorithm 4 Online Newton Step (Agarwal et al., 2006)

Require: β, δ

1: Input M assets, set $\mathbf{x}_1 \leftarrow \frac{1}{M}\mathbf{1}_M$

2: **for** $t \in \{1, \ldots, T\}$ **do**

3: Select $\mathbf{x}_{t+1} \leftarrow \Pi^{\mathbf{A}_t}_{\Delta_{M-1}}(\mathbf{x}_t - \frac{1}{\beta}\mathbf{A}_t^{-1}\mathbf{b}_t)$, where:

$\mathbf{b}_t = \sum_{\tau=1}^{t} \nabla[\log_\tau(\mathbf{x}_\tau \cdot \mathbf{r}_\tau)])$

$\mathbf{A}_t = \sum_{\tau=1}^{t} \nabla^2[\log(\mathbf{x}_\tau \cdot \mathbf{r}_\tau)] + \mathbf{1}_M$

$\Pi^{\mathbf{A}_t}_{\Delta_{M-1}}$ is the projection in the norm induced by $\mathbf{A}_t$

4: Observe $\mathbf{r}_{t+1}$ from the market

5: $W_{t+1} = \prod_{t=1}^{t+1} \langle \mathbf{r}_{t+1}, \mathbf{x}_{t+1} \rangle$

6: **end for**

Regret

$$R_T \leq \frac{10M}{8\beta} \log\left[\frac{T}{64\beta^2}\right]$$

Online Gradient Descent (OGD)

Algorithm 5 Online Gradient Descent (Zinkevich, 2003)

Require: learning rate sequence $\{\eta_1, \ldots, \eta_T\}$

1: Input M assets, set $\mathbf{x}_1 \leftarrow \frac{1}{M}\mathbf{1}$
2: **for** $t \in \{1, \ldots, T\}$ **do**
3: Select $\mathbf{x}_{t+1} \leftarrow \Pi_{\Delta_{M-1}}\left(\mathbf{x}_t + \eta_t \frac{\mathbf{r}_t}{\langle \mathbf{r}_t, \mathbf{x}_t \rangle}\right)$
4: Observe $\mathbf{r}_{t+1}$ from the market
5: Get wealth increase $\langle \mathbf{r}_{t+1}, \mathbf{x}_{t+1} \rangle$
6: **end for**

Total Regret

$$R_T \leq \left(\frac{D^2}{2K} + G^2 K\right)\sqrt{T}$$

where $D = \sup_{\mathbf{x},\mathbf{y} \in X} ||\mathbf{x} - \mathbf{y}||_2$, $G = \sup_{\mathbf{x} \in X} ||\nabla f_t(\mathbf{x})||_2$, $\eta_t = \frac{K}{\sqrt{t}}$

Online Gradient Descent with Momentum

Total Regret: Adding Transaction Costs

In order to deal with transaction costs, we extend the definition of regret to include portfolio turnover.

Total Regret

$$R_T^C = \underbrace{\log(W_T(\mathbf{x}^*, \ldots, \mathbf{x}^*)/W_T(\mathbf{x}_1, \ldots, \mathbf{x}_T))}_{R_T\text{: standard regret}} + \underbrace{\gamma \sum_{t=1}^{T} ||\mathbf{x}_t - \mathbf{x}_{t-1}||_1}_{C_T\text{: transaction costs}},$$

γ is the proportional transaction rate for buying and selling stocks

Online Gradient Descent with Momentum (OGDM)

Algorithm 6 OGDM (Vittori et al., 2020)

Require: learning rate sequence $\{\eta_1, \dots, \eta_T\}$, momentum parameter sequence $\{\lambda_1, \dots, \lambda_T\}$

1: Input M assets, set $\mathbf{x}_1 \leftarrow \frac{1}{M}\mathbf{1}$
2: **for** $t \in \{1, \dots, T\}$ **do**
3: Select $\mathbf{x}_{t+1} \leftarrow \Pi_{\Delta_{M-1}}\left(\mathbf{x}_t + \eta_t \frac{\mathbf{r}_t}{\langle \mathbf{r}_t, \mathbf{x}_t \rangle} - \frac{\lambda_t}{2}(\mathbf{x}_t - \mathbf{x}_{t-1})\right)$
4: Observe $\mathbf{r}_{t+1}$ from the market
5: Get wealth $\log(\langle \mathbf{r}_{t+1}, \mathbf{x}_{t+1} \rangle) - \gamma ||\mathbf{x}_{t+1} - \mathbf{x}_t||_1$
6: **end for**

Total Regret

$$R_T^C \leq \left[\frac{D^2}{K_\eta}\left(\frac{1}{2} + K_\lambda\right) + K_\eta \tilde{G}\left(2\gamma\sqrt{M} + \tilde{G}\right)\right]\sqrt{T}$$

where $D = \sup_{\mathbf{x},\mathbf{y} \in X} ||\mathbf{x} - \mathbf{y}||_2$, $\tilde{G} = \sup_{\mathbf{x} \in X} ||\nabla f_t(\mathbf{x})||_2 + \frac{DK_\lambda}{2K_\eta}$, $\eta_t = \frac{K_\eta}{\sqrt{t}}$, $\lambda_t = \frac{K_\lambda}{t}$

Comparison with State of the Art

- Online Portfolio Optimization:
  - Universal Portfolios (UP) [Cover and Ordentlich (1996)]
  - Online Newton Step (ONS) [Agarwal et al. (2006)]
- Online Portfolio Optimization with Transaction Costs:
  - Online Lazy Updates (OLU) [Das et al. (2013)]

	UP	ONS	OLU	OGDM
R_T	$\mathcal{O}(\log T)$	$\mathcal{O}(\log T)$	$\mathcal{O}(\sqrt{T})$	$\mathcal{O}(\sqrt{T})$
R_T^C	-	-	$\mathcal{O}(T)$	$\mathcal{O}(\sqrt{T})$
Complexity	$\Theta(T^M)$	$\Theta(M^2)$	$\Theta(M)$	$\Theta(M)$

Experiments

Experimental Setting

Datasets				
Name	Market	Year Span	Days	Assets
NYSE(O)	New York Stock Exchange	1962 - 1984	5651	36
SP500	Standard Poor's 500	1998 - 2003	1276	25
Corona	Global	2019 - 2020	280	4

Corona Dataset (03/29/2019 - 05/08/2020)		
Ticker	Description	Market Category
SPY	SPDR S&P 500 ETF Trust	Equity
BNDX	Vanguard Bond Index Fund ETF	Fixed Income
DAX	Global X DAX Germany ETF	Equity
VIX	CBOE Volatility Index	Derivatives

Experiments: Wealth $W_T^C(\mathfrak{U})$

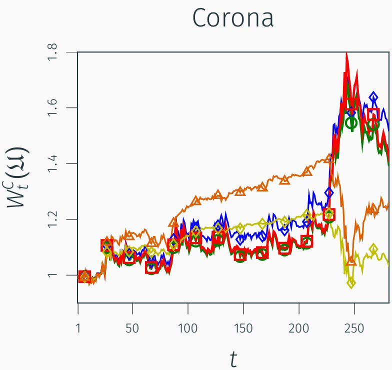

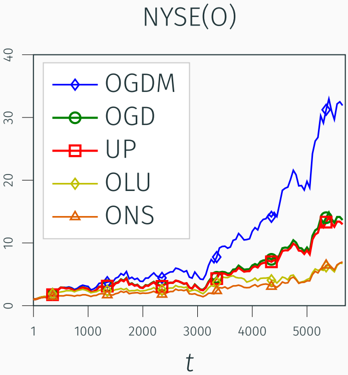

Specific run, on the Corona dataset for $\gamma = 0$

Specific run on 5 stocks of the NYSE(O) for $\gamma = 0.01$

Experiments: Average APY

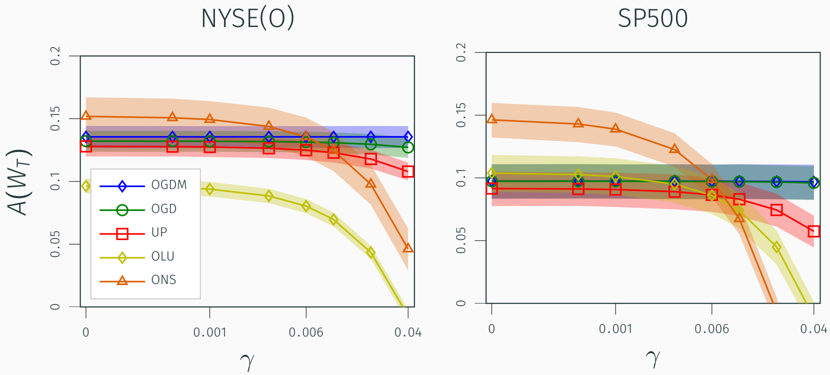

Average Average Annual Percentage Yield $A(W_T)$ computed on the wealth $W_T^C(\mathbf{x}_{1:T}, \mathbf{r}_{1:T})$:

$$A(W_T) = W_T^{250/T} - 1$$

Experiments: Average variation of the portfolio

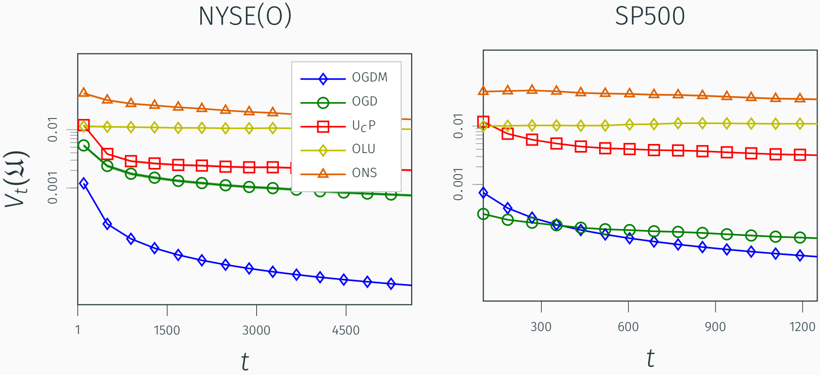

Average variation of the portfolio incurred on a varying time horizon:

$$V_t(\mathfrak{U}) = \frac{C_t(\mathfrak{U})}{t\gamma}$$

Conclusions

Conclusions and Future Work

Contributions:

- OGDM in online portfolio optimization
- Experimental campaign on real data

Future Works:

- Model stochastic non stationary markets
- Generalize the total regret analysis to other online learning algorithms

Remark:

- In order to rebalance a portfolio with hundreds of assets an entire day if not more is necessary. Expert learning algorithms are made to work with small timesteps thanks to the low computational complexity. A solution could be to work with few very liquid assets.

Q&A

Q&A

For any further questions please contact me on
edoardo.vittori@polimi.it

References

Agarwal, A., Hazan, E., Kale, S., and Schapire, R. (2006). Algorithms for portfolio management based on the newton method. In *ICML*.

Cover, T. and Ordentlich, E. (1996). Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363.

Das, P., Johnson, N., and Banerjee, A. (2013). Online lazy updates for portfolio selection with transaction costs. In *AAAI*, pages 202–208.

Li, B., Hoi, S., Sahoo, D., and Liu, Z. (2015). Moving average reversion strategy for on-line portfolio selection. *ARTIF INTELL*, 222:104–123.

Li, B., Zhao, P., Hoi, S., and Gopalkrishnan, V. (2012). Pamr: Passive aggressive mean reversion strategy for portfolio selection. *MACH LEARN*, 87(2):221–258.

Markowitz, H. (1952). Portfolio selection. *The journal of finance*, 7(1):77–91.

Vittori, E., de Luca, M. B., Trovo, F., and Restelli, M. (2020). Dealing with transaction costs in portfolio optimization: Online gradient descent with momentum. In *ACM International Conference on AI in Finance*, pages 1–8.

Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936.

Appendix

State of the Art Algorithms

- Algorithms for Online Portfolio Optimization:
  - Universal Portfolios (UP) [Cover and Ordentlich (1996)]
  - Online Newton Step (ONS) [Agarwal et al. (2006)]
- Algorithms for Online Portfolio Optimization with transaction costs:
  - Online Lazy Updates [Das et al. (2013)]

State of the Art Algorithms

- Online Portfolio Optimization:
  - Universal Portfolios (UP) [Cover and Ordentlich (1996)]
  - Online Newton Step (ONS) [Agarwal et al. (2006)]
- Online Portfolio Optimization with Transaction Costs:
  - Online Lazy Updates (OLU) [Das et al. (2013)]
  - Universal portfolios with costs
- Heuristics:
  - Passive Aggressive Mean Reversion (PAMR) [Li et al. (2012)]
  - Online Moving Average Reversion (OLMAR) [Li et al. (2015)]