COS 511: Foundations of Machine Learning

Rob Schapire — Lecture #22
Scribe: Siun-Chuon Mau — May 2, 2006

Portfolio Selection

In this lecture, we consider the problem of maximizing the wealth of an investment portfolio in the framework of on-line learning model. Since the price sequence in this framework is considered non-statistical, our results can be regarded as for the worst-case, in the sense that our results apply even for the "worst" price sequence that could have been designed by an adversary.

In the first part of the lecture, we will use Bayes' algorithm and its associated bound to guarantee that the performance (the resulting log wealth) of Bayes' algorithm is not much worse than that of the best stock picked in hindsight. In the second part, an improved strategy, dubbed universal portfolios ([1] and references therein; however, see [2] for an opposing view), that applies Bayes' algorithm to an ensemble of constantly rebalanced portfolios is discussed.

Note that the active discussions occurred during the lecture contained a few criticisms of the approaches presented here. Some of these criticisms will be discussed in this note.

1 Assumptions, Model, and Notations

Here is the list:

- There are N (fixed) investments (say stocks) available to fill the portfolio with.
- The wealth of the portfolio is reinvested each time period (say time period = day) labeled by $t = 1, 2, \ldots, T$.
- The price relative for stock i on day t is defined as:

$$p_t(i) \triangleq \frac{\text{Price of stock } i \text{ at the end of day } t}{\text{Price of stock } i \text{ at the beginning of day } t} \ . \tag{1}$$

- S_t denotes total wealth at the start of day t.
- And, $w_t(i)$ labels the fraction of wealth invested in stock i at the start of day t.

Therefore:

- The amount of wealth in stock i at the beginning of day t is $S_t w_t(i)$,
- The amount of wealth in stock i at the end of day t is $S_t w_t(i) p_t(i)$.

– The wealth at the beginning of day $t+1$ is:

$$S_{t+1} = \sum_{i=1}^{N} S_t w_t(i) p_t(i) = S_t \, \mathbf{w}_t \cdot \mathbf{p}_t \ . \tag{2}$$

– And, the final wealth is:

$$S_{T+1} = S_1 \prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t \ . \tag{3}$$

Without loss of generality, we set

$$S_1 = 1, \quad \text{or equivalently,} \quad S_{T+1} = \prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t \ . \tag{4}$$

1.1 Log-Wealth Maximization

Our goal is to choose $\mathbf{w}_t$ to maximize final wealth:

$$\max_{\mathbf{w}_t} \ S_{T+1} \quad \equiv \quad \max_{\mathbf{w}_t} \prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t \equiv \min_{\mathbf{w}_t} \sum_{t=1}^{T} \underbrace{- \ln \left(\mathbf{w}_t \cdot \mathbf{p}_t \right)}_{\text{log loss}} \ . \tag{5}$$

Hence, our portfolio selection problem is equivalent to the online learning problem of choosing the weights $\mathbf{w}_t$, for all $t = 1, 2, \cdots, T$, to minimize the total log-loss function, where $-\ln(\mathbf{w}_t \cdot \mathbf{p}_t)$ is regarded as the log-loss at time t. In particular, we would like an portfolio-selection algorithm whose performance is not much worse than the best performing stock in hindsight.

Again, the price-relative sequence is modeled to be deterministic. Although it is possible to design algorithms with better performance if a good stochastic model is used for the time-sequence of price relatives, the deterministic-sequence approach has the advantage of being simpler and providing a worst-case performance bound even for highly volatile price sequences whose stochastic model could be poorly understood.

2 Bayes' Algorithm Approach

To map our problem to an online algorithm we start by normalizing all price relatives to their maximum over both t and i. Let

$$C \triangleq \max_{t,i} \ p_t(i) \ , \tag{6}$$

and $p_t(i) \in [0, C]$, $\forall t = 1, 2, \cdots, T$, and $\forall i = 1, 2, \cdots, N$. Next, we define the outcome set for each letter of the associated time sequence as $\mathbf{X} = \{0, 1\}$, and the prediction of expert i at time t as:

$$p_{t,i}(x_t) \triangleq \begin{cases} p_t(i), & \text{if } x_t = 1 \\ C - p_t(i), & \text{if } x_t = 0 \end{cases} \tag{7}$$

In order for the expert $p_{t,i}\left(x_t\right)$ to predict $p_t(i)$ always, i.e,

$$p_{t,i}\left(x_t\right) = p_t(i), \ \forall t = 1, 2, \cdots, T, \tag{8}$$

we set

$$x_t \quad = \quad 1, \ \forall t = 1, 2, \cdots, T \tag{9}$$

and identify the (normalized) weights over experts with those over stocks:

$$w_{t,i} = w_t(i), \quad \forall t, \ \forall i, \quad \text{with} \ \sum_{i=1}^{N} w_{t,i} = 1. \tag{10}$$

The prediction of the master algorithm is therefore

$$q_t(x_t) = q_t(1) \quad = \quad \sum_{i=1}^{N} w_{t,i}\, p_{t,i} = \mathbf{w}_t \cdot \mathbf{p}_t, \quad \forall t. \tag{11}$$

A lower bound for the final wealth can be obtained by the upper bound for the total-loss function in Bayes' algorithm (assuming that initial weights $w_t(i)$ are all identically $1/N$):

$$\sum_{t=1}^{T} -\ln q_t\left(x_t\right) \quad \leq \quad \min_{i=1,2,\cdots,N} \sum_{i=t}^{T} -\ln p_{t,i}\left(x_t\right) + \ln N \ . \tag{12}$$

Since the total-loss function is related to the final wealth $\prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t$ via

$$\sum_{t=1}^{T} -\ln q_t(x_t) \quad = \quad -\ln\left(\prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t\right), \tag{13}$$

the Bayes' algorithm bound in Eq. (12) implies

$$\underbrace{\prod_{t=1}^{T} \mathbf{w}_t \cdot \mathbf{p}_t}_{\substack{\text{algorithm wealth} \\ \text{at } T}} \quad \geq \quad \frac{1}{N} \max_i \underbrace{\prod_{t=1}^{T} p_t(i)}_{\substack{\text{wealth at } T \\ \text{if invested} \\ \text{solely in stock } i}} \quad . \tag{14}$$

That is,

$$\text{(wealth of Bayes' algorithm at } T) \geq \frac{1}{N} \times \text{(wealth of the best stock at } T) \ , \tag{15}$$

which indicates that the portfolio growth rate managed by the Bayes' algorithm is as good as that of the best stock in hindsight.

It is known that the weight update rule of Bayes' algorithm is:

$$w_{t+1}(i) = \frac{w_t(i)\, p_t(i)}{\text{normalization}},\ \forall t \text{ and } \forall i\ . \tag{16}$$

As it implies no trading over the course of the algorithm (except for the initial buying to populate the portfolio), the Bayes' algorithm is effectively prescribing the buy-and-hold strategy. This fact can be used to derive the above bound directly:

$$\begin{aligned} (\text{wealth of algorithm at } T) &= \sum_{i=1}^{N} \frac{1}{N} \times (\text{wealth of stock } i) && (17)\\ &\geq \frac{1}{N} \times (\text{wealth of the best stock at } T) && (18) \end{aligned}$$

It is well known that improvements over the buy-and-hold strategy exists. One option is to use switching experts as described in the last lecture. However, we will consider universal portfolios ([1] and references therein) instead.

3 Universal Portfolios

3.1 Why Rebalance a Portfolio?

The advantage of portfolio rebalancing can be illustrated by a toy example. The portfolio consists of only two stocks with the following properties. The price of stock 1 never changes, and the price of stock 2 doubles on odd days and halves on even days. That is,

$$\begin{aligned} p_t(1) &= 1,\ \ \forall t && (19)\\ p_t(2) &= \begin{cases} 2, & t = \text{odd} \\ 1/2, & t = \text{even} \end{cases} && (20) \end{aligned}$$

Suppose the portfolio starts with unit wealth equally distributed among the two stocks. Clearly, buy-and-hold does not grow portfolio wealth over time. However, constantly rebalancing does, as illustrated below:

$$\begin{aligned} &S_1 \\ S_2 &= S_1\left(\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot\frac{1}{2}\right) = \frac{3}{4}\, S_1 \\ S_3 &= S_2\left(\frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 2\right) = \frac{3}{2}\, S_2 = \frac{9}{8}\, S_1 \\ &\vdots \\ S_5 &= \frac{9}{8} S_3 \\ &\vdots \end{aligned}$$

We can see that wealth increases exponentially by a factor of 9/8 every two days.

3.2 Constant Rebalanced Portfolios and The Universal Portfolio Algorithm

The central idea in universal portfolios is to apply Bayes' algorithm to find the optimum weighted average of an ensemble of constant rebalanced portfolios (CRP's) so that the portfolio does almost as well as the best CRP in hindsight. Initial wealth is divided equally among all (infinite number of) CRP's, then the buy-and-hold strategy is used on the ensemble of CRP's (no trading of one CRP for another).

Since the *normalized* weight vector $\mathbf{b} = (b_1, b_2, \ldots, b_N)$ of a CRP is a constant by definition, each CRP can be labeled by its own $\mathbf{b}$. We will use both $\mathrm{CRP}_{\mathbf{b}}$ and $\mathbf{b}$ to denote the CRP indexed by $\mathbf{b}$. In order to apply the Bayes' algorithm and its weight update rule in Eq. (16), we need to express the fraction, $w_t(i)$, of wealth invested in stock i at day t in terms of $\mathbf{b}$ and price relatives $\mathbf{p}$.

The differential wealth of $\mathrm{CRP}_{\mathbf{b}}$ at the end of day t is given by:

$$dS_t(\mathbf{b}) \;=\; d\mu(\mathbf{b}) \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \;, \tag{21}$$

where $d\mu(\mathbf{b})$ is the initial differential wealth of $\mathrm{CRP}_{\mathbf{b}}$. Summing over all CRP's yields the total wealth at the end of day t:

$$S_t \;=\; \int_{\text{all CRP}_{\mathbf{b}}} dS_t(\mathbf{b}) = \int_{\text{all CRP}_{\mathbf{b}}} \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \, d\mu(\mathbf{b}) \;. \tag{22}$$

Note that this means S_t is the expectation, over all $\mathbf{b}$, of the wealth of $\mathrm{CRP}_{\mathbf{b}}$:

$$S_t = \mathsf{E}_{\text{over all }\mathbf{b}}\left[\prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s\right] \;. \tag{23}$$

The part of $dS_t(\mathbf{b})$ invested in stock i is:

$$b_i \, dS_t(\mathbf{b}) \;=\; b_i \, d\mu(\mathbf{b}) \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \;. \tag{24}$$

(Remember $\sum_{i=1}^{N} b_i = 1$.) The total amount invested in stock i at the end of day t is obtained by summing over all CRP's.

$$\int_{\text{all CRP}_{\mathbf{b}}} b_i \, dS_t(\mathbf{b}) \;=\; \int_{\text{all CRP}_{\mathbf{b}}} b_i \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \, d\mu(\mathbf{b}) \;. \tag{25}$$

Therefore, the Bayes' algorithm weight, the fraction of wealth invested in stock i at the end of day t, is given by:

$$w_t(i) \;=\; \frac{\int_{\text{all CRP}_{\mathbf{b}}} b_i \, dS_t(\mathbf{b})}{\int_{\text{all CRP}_{\mathbf{b}}} dS_t(\mathbf{b})} \tag{26}$$

$$\;=\; \frac{\int_{\text{all CRP}_{\mathbf{b}}} b_i \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \, d\mu(\mathbf{b})}{\int_{\text{all CRP}_{\mathbf{b}}} \prod_{s=1}^{t} \mathbf{b}\cdot\mathbf{p}_s \, d\mu(\mathbf{b})} \;. \tag{27}$$

The algorithm defined by Eqs. (27) and (16) is called the "universal portfolio algorithm" (UPA). These integrals are generally difficult to evaluate numerically. Computationally simpler estimates can be obtained by evaluating only on a random sample of all CRP's, in effect discretizing the simplex of all CRP's.

A standard performance bound exists for this algorithm:

$$\begin{pmatrix} \text{expected wealth of} \\ \text{UPA at time } T,\ S_t \end{pmatrix} \geq \frac{1}{(T+1)^{N-1}} \times \begin{pmatrix} \text{wealth of the best CRP at } T \\ \text{(if it was the sole investment)} \end{pmatrix} \tag{28}$$

This bound is the most moderate in T dependence. Other algorithms with more moderate N dependences exist, but they are worse in other aspects. Instead of proving the above bound, we will prove a slightly weaker one, which admits a much simpler proof [1] .

The idea is to show that CRP's that are near perform similarly and that there are enough CRP's that are close enough to optimal. Let $\mathbf{b}^*$ be the best performing portfolio in hindsight, and

$$\Delta \triangleq \left\{ \mathbf{b} = (b_1, b_2, \ldots, b_N) : \sum_{i=1}^{N} b_i = 1 \right\} \tag{29}$$

be the simplex of all possible portfolio weights $\mathbf{b}$. Consider the α simplex-neighborhood of $\mathbf{b}^*$ given by

$$\delta^*(\alpha) \triangleq \{\mathbf{b} = (1-\alpha)\,\mathbf{b}^* + \alpha\,\mathbf{z} : \mathbf{z} \in \Delta\} , \tag{30}$$

where $\alpha \in [0,1]$ is fixed. Over the period of the s^{th} day,

$$\begin{aligned} (\text{one day gain of a } \mathbf{b} \in \delta^*(\alpha)) &= \mathbf{b} \cdot \mathbf{p}_s & (31)\\ &= (1-\alpha)\,\mathbf{b}^* \cdot \mathbf{p}_s + \alpha\,\mathbf{z} \cdot \mathbf{p}_s & (32)\\ &\geq (1-\alpha)\,\mathbf{b}^* \cdot \mathbf{p}_s & (33)\\ &= (1-\alpha) \times (\text{one day gain of } \mathbf{b}^*) & (34) \end{aligned}$$

Hence,

$$\begin{aligned} (\text{wealth gain of } \mathbf{b} \in \delta^*(\alpha) \text{ at day } T) &= \prod_{s=1}^{T} \mathbf{b} \cdot \mathbf{p}_s & (35)\\ &\geq (1-\alpha)^T \prod_{s=1}^{T} \mathbf{b}^* \cdot \mathbf{p}_s & (36)\\ &= (1-\alpha)^T\, (\text{wealth gain of } \mathbf{b}^* \text{ at day } T) & (37) \end{aligned}$$

The number of portfolios in $\delta^*(\alpha)$ is given by its volume, which is shrunk from Δ by a factor of α in all its $N-1$ dimensions:

$$\begin{aligned} \mathrm{Vol}\,(\delta^*(\alpha)) &= \mathrm{Vol}\,(\{\mathbf{b} = (1-\alpha)\,\mathbf{b}^* + \alpha\,\mathbf{z} : \mathbf{z} \in \Delta\}) & (38)\\ &= \mathrm{Vol}\,(\{\alpha\,\mathbf{z} : \mathbf{z} \in \Delta\}) & (39)\\ &= \mathrm{Vol}\,(\alpha\,\Delta) & (40)\\ &= \alpha^{N-1}\mathrm{Vol}\,(\Delta) & (41) \end{aligned}$$

As the distribution of initial portfolio wealth is uniform over Δ,

$$\frac{\text{initial wealth invested in } \delta^*(\alpha)}{\text{initial wealth invested in } \Delta} = \frac{\text{Vol}(\delta^*(\alpha))}{\text{Vol}(\Delta)} = \alpha^{N-1} \tag{42}$$

Putting these together:

$$\begin{aligned}
&\begin{pmatrix} \text{expected wealth of} \\ \text{UPA at time } T \end{pmatrix} && (43)\\
\geq\ &\begin{pmatrix} \text{expected wealth of} \\ \delta^*(\alpha) \text{ at time } T \end{pmatrix} && \text{(since } \delta^*(\alpha) \subseteq \Delta\text{)} \quad (44)\\
\geq\ &(1-\alpha)^T \times \begin{pmatrix} \text{wealth gain} \\ \text{of } \mathbf{b}^* \text{ at } T \end{pmatrix} \times \begin{pmatrix} \text{initial wealth} \\ \text{invested in } \delta^*(\alpha) \end{pmatrix} && \text{(using Eq. (37))} \quad (45)\\
=\ &\alpha^{N-1}(1-\alpha)^T \times \begin{pmatrix} \text{wealth gain} \\ \text{of } \mathbf{b}^* \text{ at } T \end{pmatrix} \times \begin{pmatrix} \text{initial wealth} \\ \text{invested in } \Delta \end{pmatrix} && \text{(using Eq. (42))} \quad (46)\\
=\ &\alpha^{N-1}(1-\alpha)^T \times \begin{pmatrix} \text{wealth of the best CRP at } T \\ \text{(if it was the sole investment)} \end{pmatrix} && (47)
\end{aligned}$$

This bound holds for all $\alpha \in [0,1]$, and in particular, for $\alpha = 1/(T+1)$. Together with the inequality $(1-\frac{1}{T+1})^T > 1/e$, we have

$$\begin{pmatrix} \text{expected wealth of} \\ \text{UPA at time } T \end{pmatrix} \geq \frac{1}{e(T+1)^{N-1}} \times \begin{pmatrix} \text{wealth of the best CRP at } T \\ \text{(if it was the sole investment)} \end{pmatrix} \tag{48}$$

3.3 Numerical Results

Some numerical results showing the performance comparisons of UPA and a few other algorithms when they are applied to historical stock price data were shown in class. The experimental portfolios consists of a few stocks and the algorithms were run on data lasting five to ten years. The general trend appeared to suggest that UPA performs best when stock prices are uncorrelated and more volatile, and not better otherwise. It was speculated that uncorrelatedness of stock prices played a more important role than price volatilities.

One valid criticism of the numerical experiment is that since the stocks used were listed on NYSE for many years, they were not the worst to invest in and could have been forgiving to the performance of UPA.

References

[1] Avrim Blum and Adam Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35:193–205, 1999.

[2] Paul A. Samuelson. Why we should not make mean log of wealth big though years to act are long. *Journal of Banking and Finance*, 3:305–307, 1979.