Learning Methods for Online Prediction Problems

Peter Bartlett
Statistics and EECS
UC Berkeley

Course Synopsis

- A finite comparison class: $\mathcal{A} = \{1, \ldots, m\}$.
- Converting online to batch.
- Online convex optimization.
- Log loss.
  - Three views of log loss.
  - Normalized maximum likelihood.
  - Sequential investment.
  - Constantly rebalanced portfolios.

Log Loss

A family of decision problems with several equivalent interpretations:

- Maximizing long term rate of growth in portfolio optimization.
- Minimizing redundancy in data compression.
- Minimizing likelihood ratio in sequential probability assignment.

See Nicolò Cesa-Bianchi and Gàbor Lugosi, *Prediction, Learning and Games*, Chapters 9, 10.

Log Loss

- Consider a finite outcome space $\mathcal{Y} = \{1, \ldots, m\}$.
- The comparison class $\mathcal{A}$ is a set of sequences $f_1, f_2, \ldots$ of maps $f_t : \mathcal{Y}^t \to \Delta^{\mathcal{Y}}$.
- We write $f_t(y_t|y_1, \ldots, y_{t-1})$, notation that is suggestive of a conditional probability distribution.
- The adversary chooses, at round t, a value $y_t \in \mathcal{Y}$, and the loss function for a particular sequence f is

$$\ell_t(f) = -\ln(f_t(y_t|y_1, \ldots, y_{t-1})).$$

Log Loss: Notation

$$
\begin{aligned}
y^n = y_1^n &= (y_1, \ldots, y_n), \\
f_n(y^n) &= \prod_{t=1}^{n} f_t(y_t|y^{t-1}), \\
a_n(y^n) &= \prod_{t=1}^{n} a_t(y_t|y^{t-1}).
\end{aligned}
$$

Again, this notation is suggestive of probability distributions.
Check:

$$f_n(y^n) \geq 0 \qquad\qquad \sum_{y^n \in \mathcal{Y}^n} f_n(y^n) = 1.$$

Log Loss: Three applications

- Sequential probability assignment.
- Gambling/investment.
- Data compression.

Log Loss: Sequential Probability Assignment

Think of y_t as the indicator for the event that it rains on day t. Minimizing log loss is forecasting $\Pr(y_t|y^{t-1})$ sequentially:

$$L_n^* = \inf_{f \in F} \sum_{t=1}^{n} \ln \frac{1}{f_t(y_t|y^{t-1})}$$

$$\hat{L}_n = \sum_{t=1}^{n} \ln \frac{1}{a_t(y_t|y^{t-1})}$$

$$L_n^* - \hat{L}_n = \sup_{f \in F} \ln \frac{f_n(y^n)}{a_n(y^n)},$$

which is the worst ratio of log likelihoods.

Log Loss: Gambling

Suppose we are investing our initial capital C in proportions

$$a_t(1), \ldots, a_t(m)$$

across m horses. If horse i wins, it pays odds $o_t(i) \geq 0$. In that case, our capital becomes $Ca_t(i)o_t(i)$.
Let $y_t \in \{1, \ldots, m\}$ denote the winner of race t.
Suppose that $a_t(y|y_1, \ldots, y_{t-1})$ depends on the previous winners. Then our capital goes from C to

$$C \prod_{t=1}^{n} a_t(y_t|y^{t-1})o_t(y_t).$$

Log Loss: Gambling

Compared to a set F of experts (who also start with capital C), the ratio of the best expert's final capital to ours is

$$
\begin{aligned}
&\sup_{f \in F} \frac{C \prod_{t=1}^{n} f_t(y_t | y^{t-1}) o_t(y_t)}{C \prod_{t=1}^{n} a_t(y_t | y^{t-1}) o_t(y_t)} \\
&= \sup_{f \in F} \frac{f_n(y^n)}{a_n(y^n)} \\
&= \exp\left(\sup_{f \in F} \ln \frac{f_n(y^n)}{a_n(y^n)} \right).
\end{aligned}
$$

Log Loss: Data Compression

We can identify probability distributions with codes, and view $\ln p(y^n)$ as the length (in *nats*) of an optimal sequentially constructed codeword encoding the sequence y^n, under the assumption that y^n is generated by p.

Then

$$-\ln p_n(y^n) - \inf_{f \in F} \left(-\ln f_n(y^n) \right) = \hat{L} - L^*$$

is the *redundancy* (excess length) of the code with respect to a family F of codes.

Log Loss: Optimal Prediction

The minimax regret for a class F is

$$V_n(F) = \inf_a \sup_{y^n \in \mathcal{Y}^n} \ln \frac{\sup_{f \in F} f_n(y^n)}{a_n(y^n)}.$$

For a class F and $n > 0$, define the normalized maximum likelihood strategy a^* by

$$a_n^*(y^n) = \frac{\sup_{f \in F} f_n(y^n)}{\sum_{x^n \in \mathcal{Y}^n} \sup_{f \in F} f_n(x^n)}.$$

Log Loss: Optimal Prediction

Theorem

1. a^* is the unique strategy that satisfies

$$\sup_{y^n \in \mathcal{Y}^n} \ln \frac{\sup_{f \in F} f_n(y^n)}{a_n^*(y^n)} = V_n(F).$$

2. For all $y^n \in \mathcal{Y}^n$,

$$\ln \frac{\sup_{f \in F} f_n(y^n)}{a_n^*(y^n)} = \ln \sum_{x^n \in \mathcal{Y}^n} \sup_{f \in F} f_n(x^n).$$

Log Loss: Optimal Prediction

Proof.

2. By the definition of a_n^*,

$$\ln \frac{\sup_{f\in F} f_n(y^n)}{a_n^*(y^n)} = \ln \sum_{x^n \in \mathcal{Y}^n} \sup_{f\in F} f_n(x^n).$$

1. For any other a, there must be a $y^n \in \mathcal{Y}^n$ with $a_n(y^n) < a_n^*(y^n)$. Then

$$\ln \frac{\sup_{f\in F} f_n(y^n)}{a_n(y^n)} > \ln \frac{\sup_{f\in F} f_n(y^n)}{a_n^*(y^n)},$$

which implies the sup over y^n is bigger than its value for a^*. □

Log Loss: Optimal Prediction

How do we compute the normalized maximum likelihood strategy?

$$a_n^*(y^n) = \frac{\sup_{f \in F} f_n(y^n)}{\sum_{x^n \in \mathcal{Y}^n} \sup_{f \in F} f_n(x^n)}.$$

This a_n^* is a probability distribution on $\mathcal{Y}^n$. We can calculate it sequentially via

$$a_t^*(y_t | y^{t-1}) = \frac{a_t^*(y^t)}{a_{t-1}^*(y^{t-1})},$$

where

$$a_t^*(y^t) = \sum_{y_{t+1}^n \in \mathcal{Y}^{n-t}} a_n^*(y^n).$$

Log Loss: Optimal Prediction

- In general, these are big sums.
- The normalized maximum likelihood strategy does not exist if we cannot sum $\sup_{f \in F} f_n(x^n)$ over $x^n \in \mathcal{Y}^\backslash$.
- We need to know the horizon n: it is not possible to extend the strategy for $n - 1$ to the strategy for n.
- In many cases, there are efficient strategies that approximate the performance of the optimal (normalized maximum likelihood) strategy.

Log Loss: Minimax Regret

Example

Suppose $|F| = m$. Then we have

$$\begin{aligned} V_n(F) &= \ln \sum_{y^n \in \mathcal{Y}^n} \sup_{f \in F} f_n(y^n) \\ &\leq \ln \sum_{y^n \in \mathcal{Y}^n} \sum_{f \in F} f_n(y^n) \\ &= \ln \sum_{f \in F} \sum_{y^n \in \mathcal{Y}^n} f_n(y^n) \\ &= \ln N. \end{aligned}$$

Log Loss: Minimax Regret

Example

Consider the class F of all constant experts:

$$f_t(y|y^{t-1}) = f_t(y).$$

For $|\mathcal{Y}| = 2$,

$$V_n(F) = \frac{1}{2}\ln n + \frac{1}{2}\ln\frac{\pi}{2} + o(1).$$

Minimax Regret: Proof Idea

$$V_n(F) = \ln \sum_{y^n \in \mathcal{Y}^n} \sup_{f \in F} f_n(y^n).$$

Suppose that $f(1) = q$, $f(0) = 1 - q$. Clearly, $f_n(y^n)$ depends only on the number n_1 of 1s in y^n, and it's easy to check that the maximizing value of q is n_1/n, so

$$\sup_{f \in F} f_n(y^n) = \max_q (1-q)^{n-n_1} q^{n_1} = \left(\frac{n-n_1}{n}\right)^{n-n_1} \left(\frac{n_1}{n}\right)^{n_1}.$$

Thus (using Stirling's approximation),

$$V_n(F) = \ln \sum_{n_1=1}^{n-1} \binom{n}{n_1} \left(\frac{n-n_1}{n}\right)^{n-n_1} \left(\frac{n_1}{n}\right)^{n_1}$$

$$\vdots$$

$$= \ln \left((1 + o(1)) \sqrt{\frac{n\pi}{2}} \right).$$

Log Loss

- Three views of log loss.
- Normalized maximum likelihood.
- Sequential investment.
- Constantly rebalanced portfolios.

Sequential Investment

Suppose that we have n financial instruments (let's call them $1, 2, \ldots, n$), and at each period we need to choose how to spread our capital. We invest a proportion p_i in instrument i (with $p_i \geq 0$ and $\sum_i p_i = 1$). During the period, the value of instrument i increases by a factor of $x_i \geq 0$ and so our wealth increases by a factor of

$$p'x = \sum_{i=1}^{n} p_i x_i.$$

For instance, $x_1 = 1$ and $x_2 \in \{0, 2\}$ corresponds to a choice between doing nothing and placing a fair bet at even odds.

Asymptotic growth rate optimality of logarithmic utility

Logarithmic utility has the attractive property that, if the vectors of market returns $X_1, X_2, \ldots X_t, \ldots$ are random, then maximizing expected log wealth leads to the optimal asymptotic growth rate.

We'll illustrate with a simple example, and then state a general result. Suppose that we are betting on two instruments many times. Their one-period returns (that is, the ratio of the instrument's value after period t to that before period t) satisfy

$$\begin{aligned} \mathrm{Pr}(X_{t,1} = 1) &= 1, \\ \mathrm{Pr}(X_{t,2} = 0) &= p, \\ Pr(X_{t,2} = 2) &= 1 - p. \end{aligned}$$

Clearly, one is risk free, and the other has two possible outcomes: complete loss of the investment, and doubling of the investment.

Asymptotic growth rate optimality of logarithmic utility

For instance, suppose that we start with wealth at $t = 0$ of $V_0 > 0$, and $0 < p < 1$. If we bet all of our money on instrument 2 at each step, then after T rounds we end up with expected wealth of

$$\mathbf{E}\, V_T = (2(1-p))^T V_0,$$

and this is the maximum value of expected wealth over all strategies. But with probability one, we will eventually have wealth zero if we follow this strategy. What should we do?

Asymptotic growth rate optimality of logarithmic utility

Suppose that, for period t, we bet a fraction b_t of our wealth on instrument 2. Then if we define

$$W_t = 1[X_{t,2} = 2] \qquad \text{(that is, we win the bet)},$$

then we have

$$V_{t+1} = (1 + b_t)^{W_t}(1 - b)^{1-W_t} V_t.$$

Consider the asymptotic growth rate of wealth,

$$G = \lim_{T\to\infty} \frac{1}{T} \log_2 \frac{V_T}{V_0}.$$

(This extracts the exponent.)

Asymptotic growth rate optimality of logarithmic utility

By the weak law of large numbers, we have

$$G = \lim_{T\to\infty}\left(\frac{1}{T}\sum_{t=1}^{T}\left(W_t\log_2(1+b_t)+(1-W_t)\log_2(1-b_t)\right)\right)$$
$$= \lim_{T\to\infty}\left(\frac{1}{T}\sum_{t=1}^{T}\left((1-p)\log_2(1+b_t)+p\log_2(1-b_t)\right)\right).$$

For what values of b_t is this maximized? Well, the concavity of $\log_2$, together with Jensen's inequality, implies that, for all $x_i \geq 0$ with $\sum_i x_i = x$,

$$\begin{aligned} \max \quad & \sum x_i \log y_i \\ \text{s.t.} \quad & \sum y_i = y \end{aligned}$$

has the solution $y_i = x_i y/x$. Thus, we should set $b_t = 1 - 2p$.

Asymptotic growth rate optimality of logarithmic utility

That is, if we choose the proportion b_t to allocate to each instrument so as to maximize the expected log return,

$$((1-p)\log_2(1+b_t)+p\log_2(1-b_t))\,,$$

then we obtain the optimal exponent in the asymptotic growth rate, which is

$$G=(1-p)\log_2(2(1-p))+p\log_2(2p).$$

Notice that if p is strictly less than $1/2$, $G>0$. That is, we have exponential growth. Compare this with the two individual alternatives: choosing instrument 1 gives no growth, whereas choosing instrument 2 gives expected wealth that grows exponentially, but it leads to ruin, almost surely.

Asymptotic growth rate optimality of logarithmic utility

This result was first pointed out by Kelly [5]. Kelly viewed p as the probability that a one-bit message containing the future outcome X_t was transmitted through a communication channel incorrectly, and then the optimal exponent G is equal to the channel capacity,

$$G = 1 - \left((1-p)\log_2 \frac{1}{1-p} + p\log_2 \frac{1}{p} \right).$$

Asymptotic growth rate optimality of logarithmic utility

Maximizing expected log return is asymptotically optimal much more generally. To define the general result, suppose that, in period t, we need to distribute our wealth over m instruments. We allocate proportion $b_{t,i}$ to the ith, and assume the $b_t \in \Delta_m$, the m-simplex. Then, if the period t returns are $X_{t,1}, \ldots, X_{t,m} \geq 0$, the yield per dollar invested is $b_t \cdot X_t$, so that our initial capital of V_t becomes

$$V_{t+1} = V_t b_t \cdot X_t.$$

By a strategy, we mean a sequence of functions $\{b_t\}$ which, at time t, uses the allocation $b_t(X_1, \ldots, X_{t-1}) \in \Delta_m$.

Asymptotic growth rate optimality of logarithmic utility

Definition

If $X_t \in \mathbb{R}_+^m$ denotes the random returns of m instruments during period t, we say that strategy b^* is log-optimal if

$$b_t^*(X_0, \ldots, X_{t-1}) = \arg\max_{b \in \Delta_m} \mathbf{E}\left[\log(b \cdot X_t) | X_0, \ldots, X_{t-1}\right].$$

Asymptotic growth rate optimality of logarithmic utility

Breiman [3] proved the following result for i.i.d. discrete-valued returns; Algoet and Cover [1] proved the general case.

Theorem

Suppose that the log-optimal strategy b^* has capital growth $V_0, V_1^*, \ldots, V_T^*$ over T periods and some strategy b has capital growth $V_0, V_1, \ldots, V_T$. Then almost surely

$$\limsup_{T\to\infty} \frac{1}{T} \log \frac{V_T}{V_T^*} \leq 0.$$

In particular, if the returns are i.i.d., then in each period the optimal strategy (at least, optimal to first order in the exponent) allocates its capital according to some fixed mixture $b^* \in \Delta_m$. This mixture is the one that maximizes the expected logarithm of the one-period yield.

Asymptotic growth rate optimality of logarithmic utility

This is an appealing property: if we are interested in what happens asymptotically, then we should use log as a utility function, and maximize the expected log return during each period.

Constantly rebalanced portfolios

A constantly rebalanced portfolio (CRP) is an investment strategy defined by a mixture vector $b \in \Delta_m$. At every time step, it allocates proportion b_j of the total capital to instrument j. We have seen that, for i.i.d. returns, the asymptotic growth rate is maximized by a particular CRP. The Dow Jones Industrial Average measures the performance of another CRP (the one that allocates one thirtieth of its capital to each of thirty stocks). Investing in a single stock is another special case of a CRP. (As an illustration of the benefits provided by rebalancing, consider an i.i.d. market with two instruments and return vectors chosen uniformly from $\{(2, 1/2), (1/2, 2)\}$. Investing in any single instrument leads to a growth rate of 0, whereas a $(1/2, 1/2)$ CRP will have wealth that increases by a factor of $5/4$ in each period.)

Constantly rebalanced portfolios

Now that we've motivated CRPs, we'll drop all probabilistic assumptions and move back to an online setting. Suppose that the market is adversarial (a reasonable assumption), and consider the problem of competing with the best CRP *in hindsight*. That is, at each step t we must choose an allocation of our capital b_t so that, after T rounds, the logarithm of our wealth is close to that of the best CRP.

Constantly rebalanced portfolios

The following theorem is due to Cover [4] (the proof we give is due to Blum and Kalai [2]). It shows that there is a *universal portfolio strategy*, that is, one that competes with the best CRP.

Theorem

There is a strategy (call it b_U) for which

$$\log(V_T) \geq \log(V_T(b^*)) - (m-1)\log(T+1) - 1,$$

where b^* is the best CRP.

The strategy is conceptually very simple. It involves distributing capital uniformly across all CRPs at each period.

Constantly rebalanced portfolios

Consider competing with the m single instrument portfolios. We could just place our money uniformly across the m instruments at the start, and leave it there. Then we have

$$\begin{aligned}
\log(V_T) &= \log\left(\sum_{j=1}^{m}\prod_{t=1}^{T} X_{t,j}(V_0/m)\right) \\
&\geq \max_j \log\left(\prod_{t=1}^{T} X_{t,j}(V_0/m)\right) \\
&= \max_j \log\left(\prod_{t=1}^{T} X_{t,j}V_0\right) - \log(m),
\end{aligned}$$

that is, our regret with respect to the best single instrument portfolio (in hindsight) is no more than $\log m$.

Constantly rebalanced portfolios

To compete with the set of CRPs, we adopt a similar strategy: we allocate our capital uniformly over Δ_m, and then calculate the mixture b_t that corresponds at time t to this initial distribution. Consider an infinitesimal region around a point $b \in \Delta_m$. If μ is the uniform measure on Δ_m, the initial investment in CRP b is $d\mu(b)V_0$. By time $t-1$, this has grown to $V_{t-1}(b)d\mu(b)V_0$, and so this is the contribution to the overall mixture b_t. And of course we need to appropriately normalize (by the total capital at time $t-1$):

$$b_t = \frac{\int_{\Delta_m} b V_{t-1}(b) d\mu(b)}{\int_{\Delta_m} V_{t-1}(b) d\mu(b)}.$$

Constantly rebalanced portfolios

How does this strategy perform? Suppose that b^* is the best CRP in hindsight. Then the region around b^* contains very similar mixtures, and provided that there is enough volume of sufficiently similar CRPs, our strategy should be able to compete with b^*. Indeed, consider the set of mixtures of b^* with some other vector $a \in \Delta_m$,

$$S = \{(1-\epsilon)b^* + \epsilon a : a \in \Delta_m\}.$$

For every $b \in S$, we have

$$\frac{V_1(b)}{V_0} = \frac{V_1((1-\epsilon)b^* + \epsilon a)}{V_0} \geq \frac{(1-\epsilon)V_1(b^*)}{V_0}.$$

Thus, after T steps,

$$\frac{V_T(b)}{V_T(b^*)} \geq (1-\epsilon)^T.$$

Constantly rebalanced portfolios

Also, the proportion of initial wealth allocated to CRPs in S is

$$\mu(S) = \mu(\{\epsilon a : a \in \Delta_m\}) = \epsilon^{m-1}.$$

Combining these two facts, we have that

$$\log\left(\frac{V_T(b_U)}{V_T(b^*)}\right) \geq \log\left((1-\epsilon)^T \epsilon^{m-1}\right).$$

Setting $\epsilon = 1/(T+1)$ gives a regret of

$$\log\left((1-1/(T+1))^T (T+1)^{-(m-1)}\right) > -1-(m-1)\log(T+1).$$

Constantly rebalanced portfolios

There are other approaches to portfolio optimization based on the online prediction strategies that we have seen earlier in lectures. For instance, the exponential weights algorithm can be used in this setting, although it leads to $\sqrt{T}$ regret, rather than $\log T$. Also, gradient descent approaches have also been investigated. For a Newton update method, logarithmic regret bounds have been proved.

Paul H. Algoet and Thomas M. Cover.
Asymptotic optimality and asymptotic equipartition properties of log-optimum investment.
The Annals of Probability, 16(2):876–898, 1988.

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

Leo Breiman.
Optimal gambling systems for favorable games.
In *Proc. Fourth Berkeley Symp. Math. Statist. Probab.*, volume 1, pages 60–77. Univ. California Press, 1960.

Thomas M. Cover.
Universal portfolios.
Mathematical Finance, 1(1):1–29, 1991.

Jr. J. L. Kelly.
A new interpretation of information rate.
J. Oper. Res. Soc., 57:975–985, 1956.

Course Synopsis

- A finite comparison class: $\mathcal{A} = \{1, \ldots, m\}$.
- Converting online to batch.
- Online convex optimization.
- Log loss.