Universal Data Compression and Portfolio Selection

Thomas M. Cover
Stanford University
Information Systems Laboratory
Durand 121
Stanford, CA 94305-9510
cover@isl.stanford.edu

Abstract

We consider universal data compression, universal portfolio selection (online portfolio algorithms) and the relationship of both to information theory. Apparently the fundamental minimax redundancy game in data compression and the minimax regret game for the growth rate of wealth in investment have the same answer. There is also a duality between entropy rate and the growth rate of wealth.

1 Introduction

We will discuss recent developments in computational finance, paying special attention to the developing theory in online portfolio selection.

Kelly [17] first noted a relationship between gambling and information theory in 1956. Subsequent work by Breiman [8] and others [1, 2, 3, 4, 5, 10, 11, 12, 18, 20, 22, 21] developed this connection into a theory of investment. Meanwhile an active investigation of robust methods in data compression was underway [19, 20, 23, 26, 27], where we have mentioned only some of those papers relevant to universal portfolio selection. The development of online algorithms in computer science in computational finance has sought out a much broader set of problems, as seen in [6, 9, 14, 15, 16, 29] and nicely summarized in the survey by El-Yaniv [14]. (See also the work on universal prediction by Haussler, Warmuth, Freund, Schapire, Littlestone, and others.) Some outstanding successes of computational finance, like the Black Scholes option pricing theory, can be addressed from these viewpoints.

We now turn attention to universal data compression and investment.

It is well known how to achieve data compression for stochastic processes at the entropy limit. Either arithmetic coding or Lempel-Ziv coding achieves this. If the source is unknown, then universal data compression algorithms exist which achieve the corresponding entropy limit for the unknown source within a minimax redundancy given by the information radius of the set of unknown distributions. The cost of universality is roughly $\frac{1}{2}(\log n)/n$ bits per symbol per degree of freedom for sequences of length n. The minimax redundancy also can be shown to be the channel capacity associated with this set of distributions.

We exhibit counterpart algorithms for universal investment in the stock market. It is known that the growth rate optimal portfolio choice for an independent identically distributed stock market is a constant rebalanced portfolio. One's wealth then grows exponentially.

We now ask if there is an online investment algorithm that does uniformly well if the underlying distribution of the market is unknown. It will be shown that the individual sequence minimax regret for data compression and for stock market investment are identical.

In universal data compression the number of degrees of freedom turns out to be the dimension of the family of distributions, while in universal investment, the number of degrees of freedom is the number of dimensions in the action (portfolio) space.

We show, in fact, that the value V_n of the two-person investment game of duration n, in which one player chooses a causal investment portfolio sequence and the other chooses the sequence of market price-relatives and the associated best constant rebalanced portfolio, is the same as the value of the minimax redundancy in the data compression game. The portfolio selection game contains data compression as a special case.

We begin by describing work in statistics and information theory which forms the intellectual background for current work on universal data compression and investment.

0272-5428/96 $05.00

2 Background

Perhaps the first time it was found to be advantageous to combine independent problems and solve them jointly was in Shannon's [25] 1948 work on coding with a fidelity criterion, later known as rate distortion theory. Here, for example, one wishes to quantize independent identically distributed random variables to minimize the sum of the mean squared errors.

Since the problems are independent, it makes sense that one should quantize each random variable separately to minimize the total mean squared error of the result. However, this results in a rectangular grid quantization scheme for the n-dimensional problem and can easily be bettered by quantizing the entire vector of observations. This is done optimally by the use of rate distortion theory. It can be proved that the distortion, for a given description rate R in bits per sample, is the minimum expected value of the distortion $d(x, \hat{x})$ over all distributions $p(\hat{x}|x)$ that satisfy the rate constraint $I(X; \hat{X}) \leq R$. Thus, not only is a joint treatment of the problem called for, but it yields an elegant theoretical answer.

However, rate distortion theory wasn't followed up for several years, so it had little influence as a strategy for combining problems. The major introduction of such a point of view into statistics is due to Robbins [24] in 1951 with the advent of compound Bayes and compound sequential Bayes approaches.

Robbins advocated combining problems in the following way: Suppose one has a collection of conditional distributions $f(x|\theta)$, where x is the observation and $\theta \in \{0, 1\}$ is the unknown classification of the observation. It is assumed that $f(x|\theta)$ is known. Nature chooses $(\theta_1, \ldots, \theta_n)$, and one observes conditionally independent observations $(X_1, \ldots, X_n) \sim \prod_{i=1}^{n} f_{\theta_i}(x_i)$. How should one decide on the classifications $\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_n$ in order to minimize the expected number of errors? Robbins argued that combining the problems, finding the empirical distribution $f(x)$, inferring the empirical prior (the proportion of 1's and 0's in $(\theta_1, \ldots, \theta_n)$) and then applying Bayes decision rule using this prior will yield a better proportion of errors than if one solved each problem singly.

Robbins also showed how to do this on the fly (compound sequential Bayes decision theory), that is, he developed an online approach to the problem and showed there was no asymptotic loss by so doing.

Taking another point of view, Blackwell [7] (1956) introduced the ideas of approachability and excludability and proved that one could do almost as well in a sequence of two-person zero sum games as if one had known ahead of time the empirical distribution of the plays of the second player.

3 Universal data compression

Suppose one has a random variable X with probability mass function $p(x)$. It is well known that the ideal code work length for binary descriptions of x is

$$l(x) = \log \frac{1}{p(x)}.$$

This yields an expected codeword length of

$$H(p) = \sum_x p(x) \log \frac{1}{p(x)},$$

and no uniquely decodable code (in which no codeword is the prefix of another) achieves a lower expected description length. (Throughout, we will ignore the integer constraint on $l(x)$.)

The problem of universal coding arises when one has a family $p_\theta(x)$ of probability mass functions on X indexed by θ. If θ is known, the ideal codeword length is $l_\theta(x) = \log(\frac{1}{p_\theta(x)})$. If θ is unknown, one wishes to find a (universal) codeword length $l(x)$ (which can equally well be written as $\log \frac{1}{p(x)}$ for some probability mass function p) which is close simultaneously to all of the ideal codeword lengths indexed by θ.

We consider the individual sequence regret of

$$R(x) = \max_\theta (\log \frac{1}{p(x)} - \log \frac{1}{p_\theta(x)})$$

and an expected regret

$$\begin{aligned} R(\theta) &= \sum_x p_\theta(x)(\log \tfrac{1}{p(x)} - \log \tfrac{1}{p_\theta(x)}) \\ &= \sum_x p_\theta(x) \log \tfrac{1}{p(x)} - H(p_\theta) \end{aligned}.$$

Let

$$R = \min_p \max_\theta R(\theta),$$

denote the minimax expected regret. The latter is achieved by choosing that distribution p which is the center of the minimal information ball containing all $p_\theta(x)$. Note that $R(\theta) = D(p_\theta \parallel p)$, where $D(p \parallel q) = \sum p \log \frac{p}{q}$ is the relative entropy distance of p from q. Thus

$$R = \min_{p(x)} \max_\theta R(\theta) = \min_p \max_\theta D(p_\theta \parallel p).$$

In fact, using the minimax theorem, a correspondence between minimax regret data compression and channel capacity can be made. It has been shown [Gallager (1968) unpublished] that the minimax regret is equal to the channel capacity

$$\begin{aligned} C &= \max_{p(\theta)} \sum_{\theta, x} p(\theta) p_\theta(x) \log \tfrac{p_\theta(x)}{p(x)} \\ &= \max_{p(\theta)} I(\theta; X) = R \end{aligned}$$

of a channel with inputs θ and outputs x.

Thus, a correspondence is made between the redundancy of descriptions of sources p_θ and the channel capacity. The excess description length R needed to describe random variables X uniformly well with respect to the family of pmf's $\{p_\theta(x)\}$ is precisely C, the log of the number of distinguishable parameters θ as viewed by the observation X.

Finally, the individual sequence redundancy

$$R_n^* = \min_p \max_{\theta, x^n} \left(\log \frac{1}{p(x^n)} - \log \frac{1}{p_\theta(x^n)} \right)$$

has been shown by Shtarkov [26] to be equal to V_n (defined in section 4) for n-sequences x^n drawn from an m-symbol alphabet and where the maximization is over the set of all product distributions $p_\theta(x^n) = \prod_{i=1}^n \theta(x_i)$.

4 Universal portfolios

We now ask for the counterpart to this theory in investment. Consider a sequence of portfolio vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n \geq \mathbf{0}$, $\mathbf{x}_i \in R^m$, and a nonanticipating sequence of portfolio choices $\mathbf{b}_1, \mathbf{b}_2, \ldots$, with the interpretation that x_{ij} is the price relative (ratio of closing to opening price) of stock j on day i, and b_{ij} is the proportion of current wealth invested in stock j on day i. Let $\mathbf{x}^n = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ denote the market sequence. We assume $\mathbf{b}_i$ depends only on $\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}$ and that $b_{ij} \geq 0$, $\sum_{j=1}^m b_{ij} = 1$. The wealth factor associated with this sequence is

$$S_n = \prod_{i=1}^n \mathbf{b}_i^t \mathbf{x}_i.$$

A constant rebalanced portfolio strategy uses the same portfolio $\mathbf{b}$ for each trading day. This results in a wealth increase of

$$S_n(\mathbf{x}^n, \mathbf{b}) = \prod_{i=1}^n \mathbf{b}^t \mathbf{x}_i.$$

For a sequence of price relatives $\mathbf{x}^n$ it is possible with hindsight to compute the best constant rebalanced portfolio $\mathbf{b}^*$ as the portfolio achieving

$$\max_{\mathbf{b}} S_n(\mathbf{x}^n, \mathbf{b}).$$

This achieves a wealth factor

$$S_n^*(\mathbf{x}^n) = \max_{\mathbf{b}} S_n(\mathbf{x}^n, \mathbf{b}),$$

and an exponential growth rate of wealth

$$W_n^*(\mathbf{x}^n) = \frac{1}{n} \log S_n^*(\mathbf{x}^n).$$

The best constant rebalanced portfolio $\mathbf{b}^*$ can only be computed with knowledge of market performance through n; it clearly cannot be implemented online.

Definition *An online portfolio selection strategy is a sequence of maps*

$$\hat{\mathbf{b}}_i : \mathbb{R}_+^{m(i-1)} \to \mathcal{B}, \quad i = 1, 2, \ldots$$

where

$$\hat{\mathbf{b}}_i = \hat{\mathbf{b}}_i(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1})$$

is the portfolio used on day i given past market outcomes $\mathbf{x}^{i-1} = \mathbf{x}_1, \ldots, \mathbf{x}_{i-1}$.

For any online portfolio selection strategy $\hat{\mathbf{b}}(\cdot)$ achieving a wealth factor

$$\hat{S}_n(\mathbf{x}^n) = \prod_{i=1}^n \hat{\mathbf{b}}_i^t \mathbf{x}_i$$

and exponential growth rate

$$\hat{W}_n(\mathbf{x}^n) = \frac{1}{n} \log \hat{S}_n(\mathbf{x}^n),$$

the regret is defined as

$$R_n(\mathbf{x}^n) = W_n^*(\mathbf{x}^n) - \hat{W}_n(\mathbf{x}^n).$$

Accordingly, we have the following definition of minimax regret.

Definition *The min-max regret for horizon n is defined as*

$$R_n^* = \min_{\hat{\mathbf{b}}(\cdot)} \max_{\mathbf{x}^n} R_n(\mathbf{x}^n),$$

where the minimization is over online portfolio selection strategies and the maximization is over individual sequences of asset price changes.

The min-max optimum portfolio selection strategy is the one achieving the minimum above.

The min-max regret is determined in the following theorem[21].

Theorem *The min-max regret R_n^* for n investment periods and m assets is given by*

$$R_n^* = \frac{1}{n} \log V_n,$$

where

$$V_n = \sum \binom{n}{n_1, n_2, \ldots, n_m} 2^{-nH(\frac{n_1}{n}, \ldots, \frac{n_m}{n})}.$$

The proof is given in [21].

Further, it can be shown that R_n^* is the value of the two-person zero sum game with payoff to the investor of

$$E \log \prod_{i=1}^n \hat{\mathbf{b}}_i^t \mathbf{x}_i / \prod_{i=1}^n \mathbf{b}^t \mathbf{x}_i,$$

where player 1, the investor, chooses an online portfolio strategy $\hat{\mathbf{b}}_i(\cdot)$, and the market chooses a market sequence $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ and a constant rebalanced portfolio $\mathbf{b}$. The investor has an optimal pure strategy, but the optimal market strategy is mixed.

It can be shown that

$$V_n \sim \frac{\Gamma(\frac{m}{2})}{\sqrt{\pi}(\frac{2}{n})^{(m-1)/2}}.$$

In particular, for $m = 2$ assets,

$$V_n \sim \sqrt{\frac{2}{\pi n}}.$$

Also, for all n,

$$\frac{1}{2\sqrt{n+1}} \leq V_n \leq \frac{2}{\sqrt{n+1}}.$$

Consequently, there exists an online portfolio strategy $\hat{\mathbf{b}}_i(\mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_{i-1})$ achieving wealth $\hat{S}_n(\mathbf{x}^n)$ such that

$$\frac{\hat{S}_n(\mathbf{x}^n)}{S_n^*(\mathbf{x}^n)} \geq V_n \geq \frac{1}{2\sqrt{n+1}},$$

for all market sequences $\mathbf{x}^n$.

This polynomial factor is asymptotically negligible to first order in the exponential growth rate of S_n^*.

This leads naturally to the definition of a new "hindsight" option [21], a derivative security which pays $S_n^*(\mathbf{x}^n)$, the result of investing one dollar according to the best constant rebalanced portfolio computed in hindsight for the observed market behavior $\mathbf{x}^n$. Thus $\frac{1}{V_n}$ is an upper bound on the price of this option. Standard Black Scholes assumptions (e.g., binomial lattice or geometric Brownian motion market behavior) yield an option price between 1 and $\frac{1}{V_n}$, but still behaving like $\sqrt{n}$ for $m = 2$ stocks.

5 Horizon free portfolios

We now show how to do uniformly well for all n. The infinite horizon Dirichlet-weighted universal portfolio, defined in [10, 12], uses, at time i, the portfolio

$$\hat{\mathbf{b}}_i = \hat{\mathbf{b}}_i(\mathbf{x}^{i-1}) = \frac{\int_B \mathbf{b} S_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}{\int_B S_{i-1}(\mathbf{b}, \mathbf{x}^{i-1}) d\mu(\mathbf{b})}, \quad i = 1, 2, \ldots$$

where

$$S_i(\mathbf{b}, \mathbf{x}^i) = S_i(\mathbf{b}) = \prod_{j=1}^{i} \mathbf{b}^t \mathbf{x}_j, \text{ and } S_0(\mathbf{b}, \mathbf{x}^0) = 1.$$

The measure μ on the portfolio simplex $\mathcal{B}$ is the Dirichlet $(\frac{1}{2}, \ldots \frac{1}{2})$ prior with density

$$d\mu(\mathbf{b}) = \frac{\Gamma(\frac{m}{2})}{[\Gamma(\frac{1}{2})]^m} \prod_{j=1}^{m} b_j^{-\frac{1}{2}} d\mathbf{b},$$

where $\Gamma(\cdot)$ denotes the Gamma function.

Note that this portfolio $\hat{b}_i$ does not depend on an arbitrary horizon n.

Let R_n denote the worst case regret of this portfolio, which from [12, 21] is given by

$$R_n = \frac{1}{n} \log \frac{\Gamma(\frac{1}{2})\Gamma(n + \frac{m}{2})}{\Gamma(\frac{m}{2})\Gamma(n + \frac{1}{2})}.$$

Thus it follows for $m = 2$ that

$$\hat{S}_n / S_n^* \geq \frac{1}{\sqrt{2\pi}} V_n,$$

for all n and all market sequences $\mathbf{x}_1, \mathbf{x}_2, \ldots$. Since the game theoretic solution [21] yields

$$\hat{S}_n / S_n^* \geq V_n,$$

for all n, and all $\mathbf{x}^n$, we see that demanding good minimax performance for all n costs at most a factor of $\sqrt{2\pi}$.

6 Conclusions

We find the following is true for the market. If the market $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$ is independent identically distributed, then the wealth S_n grows exponentially at rate

$$W = \max_{\mathbf{b}} E \log \mathbf{b}^t \mathbf{X},$$

in the sense that

$$\frac{1}{n} \log S_n \to W, \text{ a.e.}$$

Universal investment schemes exist which achieve this growth rate of wealth within $\frac{1}{2}(\log n)/n$ per degree of freedom. Thus, for example, if there are m stocks, there are $m - 1$ degrees of freedom in the choice of portfolio at each stage. If one has side information $Y_i \in \{1, 2, \ldots, k\}$ available at each time i, then there are $k(m-1)$ degrees of freedom.

The reason the family of constant rebalanced portfolios is a natural family is if, by chance, the stock sequence $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ is independent identically distributed, then constant rebalanced portfolios are optimal. In fact, it can be shown that if $\mathbf{b}$ maximizes $E \log \mathbf{b}^t \mathbf{x}$ over all portfolios, then $\mathbf{b}$ maximizes the growth rate of wealth. Thus, no loss in generality is incurred by the restriction of portfolios to this class if the market is i.i.d.

Although we have pointed out the similarities between data compression and investment, there are some differences. In data compression, one observes the entire sequence before encoding it into a description. In investment, one must do it on the fly (i.e., online or causally).

Also, there is an extra ingredient in market investment. One does not map market sequences into finite alphabet sequences (of descriptions), but rather into the simplex of actions B. So there is an extra maximization (the maximization over $\mathbf{b}$) involved in investment that does not occur in data compression.

We see two correspondences between data compression theory and information theory. First, minimax redundancy in data compression is equal to the channel capacity for the same set of distributions. The second duality is between minimax regret individual sequence data compression and minimax regret investment. Here it is found that the competitive ratio for the investment game is precisely equal to the exponentiated competitive difference for individual sequence minimax regret data compression.

Acknowledgement

I would like to thank Erik Ordentlich for his many contributions.

References

[1] P. Algoet. Universal schemes for prediction, gambling and portfolio selection. *Ann. Probab.*, pages 901–941, April 1992.

[2] P. Algoet and T. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *Ann. Probab.*, 16(2):876–898, 1988.

[3] A. Barron and T. Cover. A bound on the financial value of information. *IEEE Transactions on Information Theory*, 34(5):1097–1100, September 1988.

[4] R. Bell and T. Cover. Competitive optimality of logarithmic investment. *Mathematics of Operations Research*, 5(2):161–166, May 1980.

[5] R. Bell and T. Cover. Game-theoretic optimal portfolios. *Management Science*, 34(6):724–733, June 1988.

[6] S. Ben-David, A. Borodin, R. M. Karp, G. Tardos, and A. Wigderson. On the power of randomization in on-line algorithms. *Algorithmica*, 11(1):2–14, January 1994.

[7] D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific J. Mat.*, 6:1–8, 1956.

[8] L. Breiman. Optimal gambling systems for favorable games. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1:65–78, 1961.

[9] A. Chou, J. Cooperstock, R. El-Yaniv, M. Klugerman, and T. Leighton. The statistical adversary allows optimal money-making trading strategies. *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1995.

[10] T. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, January 1991.

[11] T. Cover and D. Gluss. Empirical Bayes stock market portfolios. *Advances in Applied Mathematics*, 7:170–181, 1986.

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

[13] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.

[14] R. El-Yaniv. Competitive solutions for online financial problems: A survey. To appear in Springer-Verlag Computer Science Lecture Notes series, 1996.

[15] R. El-Yaniv, A. Fiat, R. Karp, and G. Turpin. Competitive analysis of financial games. *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, pages 327–333, 1992.

[16] R. El-Yaniv and R. Karp. Nearly optimal competitive online replacement policies. *2nd Israel Symposium on Theory of Computing and Systems*, 1996.

[17] J. J. Kelly. A new interpretation of information rate. *Bell Syst. Tech. J.*, 35:917–926, 1956.

[18] F. Jamshidian. Asymptotically optimal portfolios. *Mathematical Finance*, 2(2):131–150, 1992.

[19] R. Krichevsky and V. Trofimov. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207, March 1981.

[20] N. Merhav and M. Feder. Universal schemes for sequential decision from individual data sequences. *IEEE Transactions on Information Theory*, 39(4):1280–1292, July 1993.

[21] E. Ordentlich and T. Cover. The cost of achieving the best portfolio in hindsight. *Stanford Statistics Technical Report #90*, under review by *Mathematics of Operations Research* 1996.

[22] E. Ordentlich and T. Cover. Online portfolio selection. *Proceedings of the Ninth Conference on Computational Learning Theory*, Desenzano del Garda, Italy 1996.

[23] J. Rissanen. A universal data compression system. *IEEE Transactions on Information Theory*, 29:656–664, 1983.

[24] H. Robbins. Asymptotically subminimax solutions of compound statistical decision problems. *Proceedings of the 2nd Berkeley Symposium Math. Statistics Probab.*, pages 131–148, 1951.

[25] C. Shannon. A mathematical theory of communication. *BSTJ*, 27:379–423, 623–656, July, October 1948.

[26] Y. Shtarkov. Universal sequential coding of single messages. *Problems of Information Transmission*, 23(3):3–17, July-September 1987.

[27] Y. Shtarkov, T. Tjalkens, and F. Willems. Multi-alphabet universal coding of memoryless sources. *Problems of Information Transmission*, 31:114–127, 1995.

[28] W. Szpankowski. On asymptotics of certain sums arising in coding theory. *IEEE Transactions on Information Theory*, 41(6), November 1995.

[29] M. Taksar, M. Klass, and D. Assaf. A diffusion model for optimal portfolio selection in the presence of brokerage fees. *Mathematics of Operations Research*, 13(2):277–294, May 1988.