# Learning to Satisfy

**Mark Coates**
McGill University
Montreal, Quebec
coates@ece.mcgill.ca

**Brian Erikkson**
University of Wisconsin
Madison, WI 53706
bceriksson@wisc.edu

**Robert Nowak**
University of Wisconsin
Madison, WI 53706
nowak@engr.wisc.edu

**Clayton Scott**
Rice University
Houston, TX
cscott@rice.edu

## Abstract

This paper investigates a class of learning problems called *learning satisfiability* (LSAT) problems, where the goal is to learn a set in the input (feature) space that satisfies a number of desired output (label/response) properties. LSAT problems are motivated, in part, by applications in computational finance, and an experimental investigation of LSAT in the context of portfolio selection is reported. A distinctive aspect of LSAT problems is that the output behavior is assessed only on the solution set, whereas in most statistical learning problems output behavior is evaluated over the entire input space. Consequently, certain learning criteria arising naturally in LSAT problems require a novel large deviation bounding technique.

## 1 Learning and Satisfaction

In most statistical learning problems, one is interested in minimizing a risk function such as expected squared error or probability of error. However, in many applications, one is interested in a solution to the learning problem that satisfies several criteria simultaneously, rather than simply optimizing one. In this paper, we introduce and study *learning satisfiability* (LSAT) problems, a class of learning problems where the goal is to learn a set in the input (feature) space that satisfies a number of desired properties expressed in terms of expectations and/or event probabilities.

Our interest in LSAT problems is motivated in part by applications in computational finance. In the portfolio selection problem we examine later in the paper, one is interested in identifying a set of stocks based on historical data such that not only is the expected return above some threshold but large losses are rare. This is a variant of the classical portfolio selection problem [1, 2], where the goal is to maximize expected return subject to a constraint on the allowable variance. Mathematically, suppose that we have training data $\{X_i, Y_i\}_{i=1}^n \overset{iid}{\sim} P$, where each input (feature vector) $X_i$ has an associated output (label) $Y_i$ and $P$ is the unknown joint distribution of the pair $(X_i, Y_i)$. We seek the largest set $G$ in the input space such that: (i) the expected output value $Y$ at every point in $G$ is non-negative; and (ii) the probability that the output $Y$ stays above a lower limit on $G$ is guaranteed to be large. Both criteria are in the form of constraints and express different measures of *confidence* in a favorable output.

LSAT problems also arise in a number of other important applications in classification and statistics. One example from this class is constrained multi-class classification, where in addition to minimizing the probability of misclassification over all classes, we ensure that the solution meets pre-specified constraints on the class-conditional error rates for some subset of classes. Another example is an extension of the false discovery rate approach for controlling the number of false positives in multiple hypothesis tests [3, 4]. In this setting, an LSAT problem might require that the false nondiscovery rate be minimized subject to the constraint that the false discovery rate does not exceed a specified threshold.

### 1.1 Related Work

An example of learning with multiple criteria is the Neyman-Pearson (NP) learning problem, in which one seeks a classifier that minimizes the false negative rate subject to a constraint on the false positive rate [5, 6]. An important distinction between NP learning and LSAT problems is that in LSAT problems output behavior is assessed on the solution set, whereas in NP learning (as well as most other standard learning problems) one is concerned with output behavior over the entire input space. Thus, LSAT criteria generally involve conditional probabilities/expectations that are functions of the target set, i.e. conditioning is on membership in the *output* set. In contrast, the conditioning in the constraints used in Neyman-Pearson learning (and in the performance metrics used in many standard classification approaches) is on the *input* class label. This difference leads to requirements for new theory and learning methods.

LSAT problems are also related to classical satisfiability (SAT) problems, most closely perhaps to stochastic SAT (SSAT) problems [7, 8]. SSAT problems involve criteria that depend on a mixture of controllable decision variables and stochastic variables, and the main objective is to determine whether there exist values for the decision variables such that the probability that the criteria are satisfied exceeds a certain threshold. A major difference between SSAT and LSAT problems is that the randomness in SSAT problems is typically known and therefore learning from data is not involved. Also, LSAT does not involve decision variables, but focuses on identification of the (possibly empty) set of inputs that satisfy stochastic criteria. Finally, since LSAT involves the identification of maximum volume sets, there are relationships with one-class neighbor (and support vector) machines and methods for learning minimum volume sets [9, 10, 11].

With regard to the computational finance application that we use to illustrate the framework, there has been related work in applying classification techniques (notably various flavours of CART) to partition stocks into outperforming and under-performing assets [12, 13]. Portfolios are then constructed by selecting a subset of the outperforming assets. These approaches make no attempt to satisfy any form of constraint on some form of risk metric associated with the portfolio. Learning theoretic approaches have been applied to many variations of the portfolio selection problem, in particular addressing the more involved task of devising sequential investment strategies, wherein multiple trading periods are considered and the aim is to devise a strategy for adapting the portfolio (potentially at some cost) after each trading period to maximize the growth [14, 15, 16, 17]. Although the application addressed is similar, our approach is fundamentally different, in that we learn a set that satisfies specified constraints over a high-dimensional feature space. The learning-theoretic approaches to sequential investment strategies primarily treat the successive returns as a multi-variate time-series; there is not the notion of informative stock descriptors that we investigate here. The extension of our approach to sequential investments is an attractive avenue of future research.

## 2 LSAT Problem Formulation

To formally define our problem, let us first introduce the following notation. Features $X$ are elements in the input space $\mathcal{X}$. An output $Y \in \mathcal{Y}$ is associated with each input. Let $\mathcal{P}$ denote a collection of probability measures on $\mathcal{X} \times \mathcal{Y}$. Each pair $(X, Y)$ is distributed independently and identically according to an unknown probability measure $P \in \mathcal{P}$ on $\mathcal{X} \times \mathcal{Y}$. We are interested in identifying a set in the input space where certain output constraints are met. Let $\mathcal{G}$ denote a collection of candidate sets and let $C : \mathcal{G} \times \mathcal{P} \rightarrow R^{k+1}$ be a constraint function mapping each set and probability measure to a $(k+1)$-dimensional vector of real numbers. For a given probability measure $P$, we are interested in the largest set $G \in \mathcal{G}$ that satisfies the constraint $C(G, P) \geq 0$, where the inequality is applied element-by-element. Specifically, we are interested in the solution to the following optimization:

$$\max_{G \in \mathcal{G}} \lambda(G) \text{ subject to } C(G, P) \geq 0,$$

where $\lambda(G)$ denotes the volume of a set $G$. It is possible for $\lambda(G)$ to be a more general objective function, but for concreteness we focus on volume, and assume that there is an associated measure on $\mathcal{G}$. A solution may not exist, depending on the nature of the constraints and $P$ (in such cases, we consider the empty set to be a default solution).

An alternate expression of the LSAT problem, which also lends itself naturally to the identification of the largest feasible set, is to express one of the constraint criteria as a risk function to be minimized subject to the other constraints. Write

$$C(G, P) = [\ C_0(G, P)\ \cdots\ C_k(G, P)\ ]^T .$$

Then we can pose our problem as

$$\min_{G \in \mathcal{G}} R(G, P) \text{ subject to } C_j(G, P) \geq 0,\ j = 1, \ldots, k$$

where the risk function $R(G, P)$ is chosen such that it is minimized by the largest set satisfying $C_0(G, P) \geq 0$. This optimization produces the *largest* feasible set, as desired. We wish to stress that any such risk function must satisfy two important properties with respect to the other constraints. First, if there exists a non-empty solution to the standard LSAT formulation, the (constrained) risk minimizer must coincide with this solution. Second, if there is no solution, the empty set must have smaller risk than any set failing to satisfy the constraint $C_0$.

### 2.1 Two Types of Constraints

One of the more innovative aspects of our work is the treatment of set-based constraints. In LSAT problems, we assess output behavior only on the solution set, whereas in most statistical learning problems output behavior is evaluated over the entire input space. We consider two types of set-based output constraints.

**1. Point-wise Constraint:** $C(G, P) = C(x, G, P)$ is a function of the input variable $x$, and the constraint takes the form

$$C(x, G, P) \geq 0,\ \forall\, x \in G$$

**2. Set-average Constraint:** $C(G, P)$ is only a function of the set $G$, and the constraint $C(G, P) \geq 0$ is only satisfied "on-average" over the set $G$.

Examples of the point-wise type of constraint include $E[Y|X = x] \geq 0$ and $P(Y \geq L|X = x) - p \geq 0, \forall\, x \in G$. Corresponding examples for the set-average constraint type are $E[Y|X \in G] \geq 0$ and $P(Y \geq L|X \in G) - p \geq 0$. Set-average constraints lead to statistical learning problems involving "self-normalizing" random sums, which require novel large deviation bounds, such as that provided in the Appendix.

### 2.2 An Illustrative LSAT Example: Portfolio Selection

We illustrate the class of LSAT problems using an example motivated by financial data analysis. We address a modification of the classical portfolio selection problem as posed by Markowitz [1, 2]. In the Markowitz model, the *return* is the expected value of the random portfolio, which has an associated *risk* as quantified by the variance of the return. The classical problem is to assign an available capital to a set of available stocks in order to maximize the return when there is an upper bound on the acceptable risk or to minimize the risk when there is a lower bound on the acceptable return. This can be posed as a convex quadratic programming problem [1], but the resultant solutions are extremely susceptible to perturbations in the model parameters, which being estimates of market behaviour exhibit substantial statistical error. Approaches have been proposed to derive more robust solutions [2], but these are parametric in nature, and good performance requires that the adopted model provides an adequate explanation of what can be high-dimensional data with complex structure.

Here we adopt a non-parametric approach and address a relaxed problem that provides more robust solutions. Instead of seeking to identify a portfolio that maximizes return or minimizes risk, we attempt to identify a set of stocks that satisfies two constraints: the expected return of any member stock must be greater than a threshold $U$ and the probability of large loss (return less than $L$) over the entire set must be smaller than a specified threshold $p$. We pose this as an LSAT problem based on historical training data and solve using a natural risk minimization formulation. Note that the LSAT framework could incorporate a constraint on risk as measured by variance over the set, providing a more natural parallel to the risk in the classical problem; we choose the alternative measure of risk because it better illustrates the self-normalizing set-average constraints that are natural in other LSAT problems and it also provides a meaningful, less-studied measure of risk.

We are interested in the largest set $G \in \mathcal{X}$ such that $E[Y|X = x] \geq U$, for all $x \in G$, and $P(Y > L|X \in G) \geq p$. The parameters $U$, $L < U$, and $p > 0$ are specified by the user. To cast this in the notation above, let

$$C(G, P) \ = \ \begin{bmatrix} \min_{x \in G} E[Y|X = x] - U \\ P(Y > L|X \in G) - p \end{bmatrix}$$

In an investment application, $U > 0$ expresses the desire for expected positive returns, and $L$ might typically be a negative value with $p > 0$ being a probability close to one in order to avoid very large losses. The probability measure governing stock features and returns is unknown, but we do have access to historical records of stock characteristics and performance. These records provide training data, possibly with the inclusion of features that reflect trends in the market. In practice, one can obtain empirical constraints by employing (nonparametric) estimators of $f(x) = E[Y|X = x]$ and $P(Y > L|X \in G)$ in place of their ensemble-average counterparts. We characterize the performance of methods based on empirical constraints in the next section.

Before concluding this section, we demonstrate the constrained risk minimization formulation of the LSAT problem. Define the risk function

$$R(G, P) \ = \ E[(U - Y_i)(1_{X_i \in G} - 1_{X_i \in \bar{G}})]$$

where $1_{X_i \in G}$ is the indicator function (outputs 1 if $X_i \in G$ and 0 otherwise) and $\bar{G}$ denotes the complement of $G$. It is easy to see that this risk is minimized (under the constraint) by the $U$-level set of the regression function $E[Y|X = x]$, i.e., the largest set satisfying $\min_{x \in G} E[Y|X \in x] - U \geq 0$ [18]. As required in the risk minimization formulation, the empty set has smaller risk than any set that fails to satisfy the constraint $C_0$.

## 3 Learning to Satisfy

In the sequel we focus our attention on the alternate formulation of the LSAT problem in terms of $R(G, P)$ because of its applicability to our finance application. The theoretical claims below have analogous statements for the original formulation in terms of $\lambda(G)$.

We are interested in identifying the set $G \in \mathcal{G}$ that satisfies the constraints $C(G, P) \geq 0$ and has minimum risk $R(G, P)$. However, since the probability measure $P$ is unknown, we aim to learn this set from a training sample $\{X_i, Y_i\}_{i=1}^n$. Suppose that we form empirical versions of the constraint functions $C_i(G, \widehat{P})$ and risk $R(G, \widehat{P})$, based on the empirical distribution $\widehat{P}$ of the training sample. For the remainder of the paper we will no longer explicitly indicate the dependence of the constraints on the underlying probability measure $P$, simply writing $C(G) = C(G, P)$, $\widehat{C}(G) = C(G, \widehat{P})$, $R(G) = R(G, P)$, and $\widehat{R}(G) = R(G, \widehat{P})$.

Define the optimal set

$$G^* \ = \ \arg\min_{G \in \mathcal{G}} R(G) \text{ subject to } C_j(G) \geq 0, \ j = 1, \ldots, k.$$

Let $\epsilon_0, \ldots, \epsilon_k > 0$ be fixed and define

$$\widehat{G} \ = \ \arg\min_{G \in \mathcal{G}} \widehat{R}(G) \ \text{ subject to } \ \widehat{C}_j(G) \geq -\epsilon_1, \ j = 1, \ldots, k.$$

By allowing the constraints to be violated by the small tolerances $\epsilon_i$, we are able to relate the performance of $\widehat{G}$ to that of $G^*$.

**Lemma 1.** *If* $\sup_{G \in \mathcal{G}} |R(G) - \widehat{R}(G)| \leq \epsilon_0$ *and* $\sup_{G \in \mathcal{G}} |C_j(G) - \widehat{C}_j(G)| \leq \epsilon_j$ *for* $j = 1, \ldots, k$ *then*

$$R(\widehat{G}) \leq R(G^*) + 2\epsilon_0 \quad \textit{and} \quad C_j(\widehat{G}) \geq -2\epsilon_j, \quad j = 1, \ldots, k$$

*Proof.* Under the assumed deviation bounds $\widehat{C}_j(G^*) \geq C_j(G^*) - \epsilon_j \geq -\epsilon_j$, which implies that $G^*$ is in the empirical constraint set. Thus $\widehat{G}$ minimizes $\widehat{R}$ subject to the empirical constraints: $R(\widehat{G}) \leq \widehat{R}(\widehat{G}) + \epsilon_0 \leq \widehat{R}(G^*) + \epsilon_0$. Applying the assumed deviation bound again to $\widehat{R}(G^*)$ produces the result. □

Note that if the deviation bounds, rather than holding deterministically as stated in the lemmas, instead hold with large probability with respect to a random training sample, then the conclusions of the lemma also hold with the same large probability.

## 4 Finance Example — Revisited

Recall the problem of identifying the largest set $G \in \mathcal{X}$ such that $\min_{x \in G} E[Y|X = x] \geq U$ and $P(Y > L|X \in G) \geq p$. This problem is equivalent to solving the constrained risk minimization

$$\min_{G \in \mathcal{G}} E[(U - Y)(1_{X \in G} - 1_{X \in \bar{G}})] \text{ subject to } P(Y > L|X \in G) \geq p$$

We do not know the probability law $P$ governing $(X, Y)$, so we use natural empirical counterparts of the risk and constraint based on available training data:

$$\widehat{R}(G) = \frac{1}{n}\sum_{i=1}^{n}(U - Y_i)(1_{X_i \in G} - 1_{X_i \in \bar{G}}) \text{ and } \widehat{C}(G) = \frac{\sum_{i=1}^{n} 1_{X_i \in G, Y_i < L}}{\sum_{i=1}^{n} 1_{X_i \in G}}$$

with the convention that $0/0 = 0$. Note that $E[\widehat{C}(G)] \approx P(Y_i < L|X_i \in G)$.

Now for deviation bounds, we make the following observations. Assume that $-B/2 \leq Y_i \leq B/2$, for $i = 1, \ldots, n$ and $B > 0$. Since $\widehat{R}(S)$ is a sum of independent and bounded random variables, Hoeffding's inequality implies that for every $\delta > 0$ with probability at least $1 - \delta$

$$|\widehat{R}(G) - R(G)| \leq B\sqrt{\frac{\log(2/\delta)}{2n}}$$

for a given set $G$. The deviations of the empirical constraint function (which is a "self-normalizing" random sum) are bounded by the following lemma.

**Lemma 2.** *Define*

$$\epsilon(k, \delta) = \begin{cases} 1 & k = 0, \\ \sqrt{\frac{\log(2/\delta)}{2k}} & k > 0. \end{cases}$$

*Then for a set G and for all $\delta$, $P(|\widehat{C}(G) - C(G)| > \epsilon(\widehat{k}, \delta)) \leq \delta$, where $\widehat{k} = \sum_i 1_{X_i \in G}$.*

The proof of Lemma 2 is given in the appendix. Now suppose that $\mathcal{G}$ is a finite collection and denote its cardinality by $|\mathcal{G}|$. Then for every $G \in \mathcal{G}$ with probability at least $1 - 2\delta$,

$$\begin{aligned} |R(G) - \widehat{R}(G)| &\leq B\sqrt{\frac{\log|\mathcal{G}| + \log(2/\delta)}{2n}} \\ |C(G) - \widehat{C}(G)| &\leq \epsilon(\widehat{k}, \delta/|\mathcal{G}|) \end{aligned}$$

From here, we are almost in a position to apply these deviation bounds to Lemma 1, with the obvious identification of $\epsilon_0$ and $\epsilon_1$. However, Lemma 1 assumed fixed tolerances where as the tolerances above are data and set dependent. Fortunately, the self-normalizing sum bound can be stated in an essentially equivalent form where the tolerances are fixed. In addition, Lemma 1 has a counterpart that allows for variable tolerances. We opted to present the current versions because of space limitations.

Similar deviation bounds can be obtained if $\mathcal{G}$ is countable or has finite VC dimension, but since the development of such bounds is fairly standard we do not provide details here. However, in anticipation of our experimental investigation, we will discuss deviation bounds specialized to dyadic decision trees (DDTs). Let $\mathcal{G}_L$ denote the collection of all dyadic rectangular partitions (dyadic decision trees) consisting of leafs/cells with side-lengths no smaller than $2^{-L}$. Let $\mathcal{G}_L^m$ denote the subset of such trees with no more than $m$ leafs. Then it is easy to show that for $m$ sufficiently large, the cardinality of $\mathcal{G}_L^m$ is no more than $(8d)^m$, where $d$ is the dimension of the input space [6]. It follows that for every $\delta > 0$ and $G \in \mathcal{G}_L^m$ with probability at least $1 - 2\delta$

$$\begin{aligned} |R(G) - \widehat{R}(G)| &\leq B\sqrt{\frac{m\log(8d) + \log(2/\delta)}{2n}} \\ |C(G) - \widehat{C}(G)| &\leq \epsilon(\widehat{k}, \delta/(8d)^m) \end{aligned} \tag{1}$$

where $m$ is the number of leafs/cells in $G$. Moreover, the same bounds hold simultaneously for all $m$ sufficiently large.

## 5 Algorithmic Considerations

Although an exact algorithm exists for direct computation of the DDT optimizing our empirical constrained optimization problem, it is computationally demanding. Efficient optimization over DDTs requires minimizing an unconstrained objective that is additive, meaning that it is a sum of terms over the leaves of its argument. To obtain an unconstrained objective we minimize the Lagrangian

$$\min_{G\in\mathcal{G}} \widehat{R}(G) + g_1\widehat{C_1}(G) + g_2\widehat{C_2}(G) + \cdots + g_k\widehat{C_k}(G). \tag{2}$$

The constraints are enforced by minimizing the Lagrangian in an iterative fashion, updating the Lagrange multipliers $g_i$ in a sequential fashion by means of a bisection search strategy.

Unfortunately constraints based on self-normalizing sums are not additive. In such cases we replace these terms by corresponding additive terms that possess similar qualitative properties, such as assigning smaller penalties to larger-volume trees. Note that verifying the actual constraints is still possible at the end of each iteration. Thus, the final solution is still guaranteed to satisfy the desired constraints to within the tolerance guaranteed by Lemma 1. On the other hand, our heuristic relies on the unnormalized additive constraints having a similar behaviour to the self-normalizing constraints. Thus, for some weight settings, the final risk may deviate from the optimal one by more than the tolerance in Lemma 1.

## 6 Experiments

We examined the performance of the LSAT methodology when applied to the computational finance problem outlined in Section 2.2 using a dataset derived from 90 securities belonging to the current NASDAQ-100 index that have been actively traded since January 1999. The dataset consists of monthly financial data over the period from January 1999 to December 2004, and comprises seven features for each security: the change in the price-to-book ratio, the price-to-book ratio, percent change in the volume traded, percent change in the price, change in the price-to-earnings ratio, price-to-earnings ratio, and the price variation (differential between month-high and month-low). The return is evaluated as the return on one dollar that is invested in the security at the start of the month and capitalized at the end of the month (effectively the ratio between the closing price and the opening price for the month).

We conducted experiments to determine the features that provided separation of the data. For illustrative purposes, we report on the results of determining LSAT sets in two dimensions corresponding to the change-in-volume and price variation features. The dataset consists of 6212 monthly returns (after elimination of data points due to missing or meaningless features). Training was performed on the three-year period of 1999-2001 and testing performed on the three-year period of 2002-2004. Figure 1(a) shows how the LSAT method can be used to reduce the risk (empirical probability of large-loss) at the cost of reducing the average return, when compared to a portfolio comprising all securities with equally-weighted investment or a portfolio derived by learning an unconstrained level set. The LSAT sets meet the imposed constraints for the training data and come close to meeting them for the test data. As shown in Figure 1(b), if the threshold U in the expected return constraint is increased, then the portfolio generates a higher average return (up to three times the return of a portfolio comprised of all securites), but the return is highly variable and the probability of large loss is high. Figure 1(c) shows that the LSAT sets with large U values have small volume and identify few investment opportunities. As the constraints are loosened, the volumes of the identified sets increase. The unconstrained set has smaller volume than some of the constrained sets because it minimizes the empirical risk $\widehat{R}$ but comes nowhere near satisfying the set-average constraint; larger sets have greater normalizing factors and can satisfy the constraint. Figure 1(d) depicts the variation in performance over time. Although the average return of the entire test-set is approximately constant for each year of the three-year test-period, the average returns of the LSAT sets decrease, indicating that the learned behaviour has greater relevance in the near-term. This suggests using a learning window and re-training prior to each investment period.

## 7 Conclusions

This paper introduced a new learning framework for handling LSAT problems. An application in portfolio selection was given to motivate and explore the potential of this framework. We anticipate many other applications for LSAT, including multi-class classification problems and false-discovery rate analysis. An algorithm for solving LSAT problems using decision trees was presented, and future work will be directed at the development of optimal tree-based methods as well as approaches based on SVMs and boosting. Substantial further investigation is required to address the many practical issues of portfolio selection, but the results presented in the paper illustrate the potential of the LSAT framework to provide solutions that effectively trade-off between the competing constraints of risk and return.

## 8 Appendix

To establish Lemma 2, we prove a slightly more general result that is also of independent interest.

**Theorem (Self-Normalizing Sum Bound).** *Let the pairs* $\{(Z_i, W_i)\}_{i=1}^n$ *be independent and identically distributed with* $a \le Z_i \le b$*, for suitable constants* $a$ *and* $b$*, and with* $W_i \in \{0, 1\}$*. Let* $\mu = E[Z|W = 1]$ *and define the estimator* $\widehat{\mu} = \frac{\sum_{i=1}^n Z_i W_i}{\sum_{i=1}^n W_i}$ *, with the convention that* $0/0 = 0$*. Define*

$$\epsilon(k,\delta) = \begin{cases} b-a & k = 0, \\ \sqrt{\frac{(b-a)^2 \log(2/\delta)}{2k}} & k > 0. \end{cases}$$

*Then for every* $\delta > 0$*,* $P(|\mu - \widehat{\mu}| < \epsilon(\widehat{k}, \delta)) \le \delta$*, where* $\widehat{k} = \sum_{i=1}^n W_i$*.*

*Proof.* For $k > 0$, $E[\widehat{\mu}|\widehat{k} = k] = \mu$. Hoeffding's inequality and the definition of $\epsilon(k, \delta)$ imply

$$P(|\mu - \widehat{\mu}| > \epsilon(\widehat{k}, \delta) \,|\, \widehat{k} = k) \quad \le \quad 2e^{-2k\epsilon^2(k)/(b-a)^2} \;=\; \delta.$$

Therefore we have

$$\begin{aligned} P(|\mu - \widehat{\mu}| > \epsilon(\widehat{k}, \delta)) &= \sum_{k \ge 0} P(|\mu - \widehat{\mu}| > \epsilon(\widehat{k}, \delta) \,|\, \widehat{k} = k) P(\widehat{k} = k) \\ &\le \sum_{k \ge 0} \delta P(\widehat{k} = k) = \delta \end{aligned}$$

This completes the proof. □

## References


[1] H.M. Markowitz. Portfolio selection. *J. Finance*, 7:77–91, 1952.

[2] D. Goldfarb and G. Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38, Feb. 2003.

[3] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J.R. Statist. Soc. (Series B)*, 57(1):289–300, 1995.

[4] J.D. Storey. A direct approach to false discovery rates. *J.R. Statist. Soc. (Series B)*, 64(3):479–498, 2002.

[5] A. Cannon, J. Howse, D. Hush, and C. Scovel. Learning with the Neyman-Pearson and min-max criteria. Technical Report LA-UR 02-2951, Los Alamos National Laboratory, 2002.

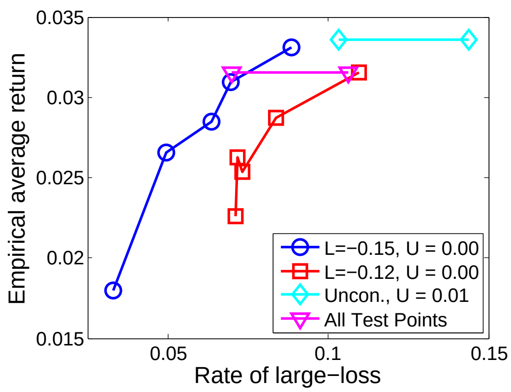


(a) Average return versus risk for U=0.

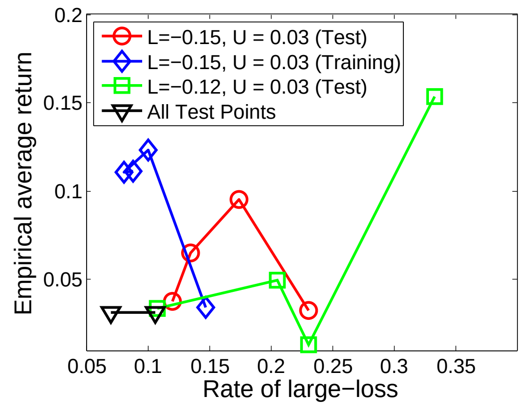


(b) Average return versus risk for U=0.03.

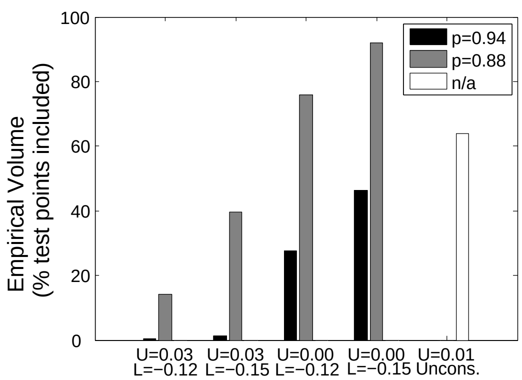


(c) Volume of learned sets

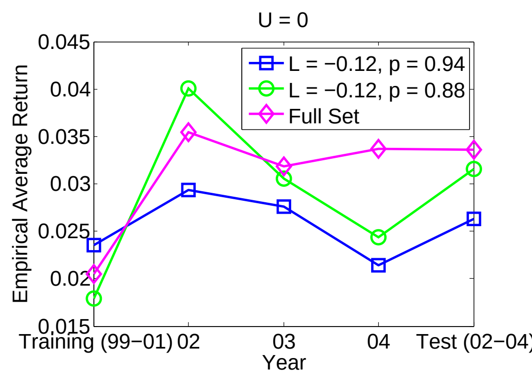


(d) Variation of performance over time

Figure 1: Performance of LSAT sets as portfolio selectors. (a) The return of sets when risk (probability of loss) is tightly controlled, showing two LSAT solutions for different settings of L with U=0. Each line plots 5 values of p (0.96, 0.94, 0.92, 0.9, and 0.88), which appear left-to-right. For comparison, we include the average return of (i) all securities and (ii) those lying in the unconstrained 0.01-level set. Two points are shown corresponding to the different probabilities of large loss depending on the choice of L. (b) The LSAT solutions for two cases when the U parameter in the first constraint is increased to 0.03. The resultant sets generate highly variable returns, in some cases outperforming the index by a factor of 3, and in other cases under-performing by a factor of 3. The sets meet constraints for the training data, but for the test data the empirical percentages of loss are large and not affected by the specified constraints. (c) The volumes of learned sets (measured as the percentage of the test data lying inside the set). The high-risk, high-return U=0.03 sets have very small volume, identifying few investment opportunities. (d) The variation of performance over time. Based on sets learned from data over the period 1999-2001, we evaluate return separately for each year 2002, 2003, and 2004. The results indicate that the learned behaviour is more relevant soon after the training period.

[6] C. Scott and R. Nowak. A Neyman-Pearson approach to statistical learning. *IEEE Trans. Information Theory*, 51(11):3806–3819, Nov. 2005.

[7] M. Majercik and M. L. Littman. Contingent planning under uncertainty via stochastic satisfiability. *Artificial Intell.*, pages 119–162, 2003.

[8] S. Manandhar, A. Tarim, and T. Walsh. Scenario-based stochastic constraint programming. In *Proc. Int. J. Conf. Artificial Intelligence (IJCAI)*, Acapulco, Mexico, Aug. 2003.

[9] A. Muñoz and J.M. Moguerza. Estimation of high-density regions using one-class neighbor machines. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(3):476–480, Mar. 2006.

[10] C. Scott and R. Nowak. Learning minimum volume sets. *J. Machine Learn. Res.*, 7:665–704, Apr. 2006.

[11] W. Polonik. Minimum volume sets and generalized quantile processes. *Stochastic Processes and their Applications*, 69:1–24, 1997.

[12] E.H. Sorensen, K.L. Miller, and C.K. Ooi. The decision tree approach to stock selection. *The Journal of Portfolio Management*, 27:42–52, 2000.

[13] G.T. Albanis and R.A. Batchelor. *Advances in Quantitative Asset Management*, chapter 11: Five classification algorithms to predict high performance stocks. Kluwer Academic Publishers, 2000.

[14] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, UK, 2006.

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

[16] L. Györfi, G. Lugosi, and F. Udina. Nonparametric kernel-based sequential investment strategies. *Mathematical Finance*, 16:337–358, 2006.

[17] E. Ordentlich and T.M. Cover. Universal portfolios with side information. *IEEE Trans. Information Theory*, 42(2):348–363, Feb. 1996.

[18] R. Willett and R. Nowak. Level set estimation via trees. In *Proc. IEEE ICASSP*, Phil., PA, Mar. 2005.