

Online Portfolio Selection with Group Sparsity

Puja Das and Nicholas Johnson and Arindam Banerjee

Department of Computer Science and Engineering

University of Minnesota

Minneapolis, MN 55455

{pdas, njohnson, banerjee}@cs.umn.edu

Abstract

In portfolio selection, it often might be preferable to focus on a few top performing industries/sectors to beat the market. These top performing sectors however might change over time. In this paper, we propose an online portfolio selection algorithm that can take advantage of sector information through the use of a group sparsity inducing regularizer while making lazy updates to the portfolio. The lazy updates prevent changing ones portfolio too often which otherwise might incur huge transaction costs. The proposed formulation leads to a non-smooth constrained optimization problem at every step, with the constraint that the solution has to lie in a probability simplex. We propose an efficient primal-dual based alternating direction method of multipliers algorithm and demonstrate its effectiveness for the problem of online portfolio selection with sector information. We show that our algorithm OLU-GS has sub-linear regret *w.r.t.* the best *fixed* and best *shifting* solution in hindsight. We successfully establish the robustness and scalability of OLU-GS by performing extensive experiments on two real-world datasets.

1 Introduction

Investors often follow a top down approach which usually involves group selection followed by identifying the most profitable stocks within a group. One of the ways investors group stocks is by the type of business. The idea is to put companies in similar sectors together. However, not all sectors can yield profit and not all stocks in a particular sector can be profitable. Moreover, sectors might react differently during different economic conditions (Li, Vassalou, and Xing 2006; Arouri and Nguyen 2010). For example, defensive sectors like utilities and consumer staples are robust to economic downturns whereas cyclical sectors which include technology, financials, health care, etc., tend to react quickly to fluctuations in the market. We are particularly interested in taking advantage of and exploiting any underlying structure amongst the stocks for the problem of online portfolio selection.

Online portfolio selection has largely been a success story (Cover 1991; Helmbold et al. 1998; Cesa-Bianchi and Lugosi 2006; Agarwal et al. 2006; Borodin, El-Yaniv, and Gogan 2004; Das and Banerjee 2011; Li et al. 2011;

Das, Johnson, and Banerjee 2013; Li and Hoi 2014) over the last two decades. These new methods for portfolio selection have been designed to not make any statistical assumptions regarding the movement of stocks (Cover 1991; Cover and Ordentlich 1996; Helmbold et al. 1998). Moreover, in a well-defined technical sense, such methods are guaranteed to perform competitively with certain families of adaptive portfolios even in an adversarial market. However, the existing work has not attempted to take advantage of the group structure that could exist amongst the stocks.

We specifically focus on using a group sparsity inducing regularizer in an online learning framework where the updates to the solutions are sparse. Such lazy updates are motivated by our desire to handle proportional transaction costs in the portfolio selection problem. An investor could incur substantial transaction costs if his portfolio changes aggressively everyday (Das, Johnson, and Banerjee 2013; Blum and Kalai 1997; Davis and Norman 1990).

In this paper, we first propose our general online lazy updates with group sparsity framework and go on to show that the online portfolio selection with sector information is a special case of this framework. Next, we introduce our OLU-GS algorithm which induces group sparsity and ensures that the updates are lazy. This results in solving a constrained non-smooth convex optimization problem at every iteration. We propose a novel alternating direction method of multipliers (ADMM) algorithm to solve this problem efficiently. In our analysis, which applies to any convex composite function with lazy updates, we show that our algorithm has $O(\sqrt{T})$ regret for general convex functions and $O(\log T)$ regret for strongly convex functions. Additionally, we prove regret bounds with respect to a shifting solution which has the benefit of hindsight.

We conduct extensive experiments on two real-world datasets and use the Global Industry Classification Standard to group the stocks into sectors for 22 years of the benchmark NYSE dataset with 30 stocks and 8 sectors and 22 years of a S&P500 dataset with 243 stocks and 9 sectors. Our experiments show that our sparse group lazy portfolios can take advantage of the sector information to beat the market and are scalable with transaction costs. It shows an interesting group switching behavior and could be especially beneficial for individual investors who have expertise in select market sectors and are averse to changing their portfolio.

2 Online Portfolio Selection

We consider a stock market consisting of n stocks $\{s_1, \dots, s_n\}$ over a span of T periods. For ease of exposition, we will consider a period to be a day, but the analysis presented holds for any valid definition of a ‘period’ such as an hour or a month. Let $x_t(i)$ denote the *price relative* of stock s_i in day t , i.e., the multiplicative factor by which the price of s_i changes in day t . Hence, $x_t(i) > 1$ implies a gain, $x_t(i) < 1$ implies a loss, and $x_t(i) = 1$ implies the price remained unchanged. We assume, $x_t(i) > 0 \forall i, t$. Let $\mathbf{x}_t = \langle x_t(1), \dots, x_t(n) \rangle$ denote the vector of price relatives for day t , and let $\mathbf{x}_{1:t}$ denote the collection of such price relative vectors up to and including day t . A portfolio $\mathbf{p}_t = \langle p_t(1), \dots, p_t(n) \rangle$ on day t can be viewed as a probability distribution over the stocks that prescribes investing $p_t(i)$ fraction of the current wealth in stock s_i . Note that the portfolio \mathbf{p}_t has to be decided before knowing \mathbf{x}_t which will be revealed only at the end of the day. The multiplicative gain in wealth at the end of day t is then simply $\mathbf{p}_t^T \mathbf{x}_t = \sum_{i=1}^n p_t(i) x_t(i)$. For a sequence of price relatives $\mathbf{x}_{1:t-1} = \{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}\}$ up to day $(t-1)$, the sequential portfolio selection problem in day t is to determine a portfolio \mathbf{p}_t based on past performance of the stocks. At the end of day t , \mathbf{x}_t is revealed and the actual performance of \mathbf{p}_t gets determined by $\mathbf{p}_t^T \mathbf{x}_t$. Over T periods, for a sequence of portfolios $\mathbf{p}_{1:T} = \{\mathbf{p}_1, \dots, \mathbf{p}_T\}$, the multiplicative gain in wealth is $S_T(\mathbf{p}_{1:T}, \mathbf{x}_{1:T}) = \prod_{t=1}^T (\mathbf{p}_t^T \mathbf{x}_t)$ and the logarithmic gain in wealth is, $LS_T(\mathbf{p}_{1:T}, \mathbf{x}_{1:T}) = \sum_{t=1}^T \log(\mathbf{p}_t^T \mathbf{x}_t)$. Ideally, for a costless environment (no transaction costs) we would like to maximize $LS_T(\mathbf{p}_{1:T}, \mathbf{x}_{1:T})$ over $\mathbf{p}_{1:T}$. However, online portfolio selection cannot be posed as an optimization problem due to the temporal nature of the choices: \mathbf{x}_t is not available when one has to decide on \mathbf{p}_t . Further, (statistical) assumptions regarding \mathbf{x}_t can be difficult to make.

3 Portfolio Selection with Group Sparsity

We focus on the problem of online portfolio selection with group sparsity where the groups are the pre-specified market sectors. The goal is to adaptively identify and invest in a few top performing sectors at any given period. In order to make our approach practical, we do not want the portfolios to change drastically everyday as an investor will have to pay transaction costs. So we encourage lazy updates to our portfolios along with group sparsity.

3.1 Related Work

Previous work (Cover 1991; Helmbold et al. 1998; Agarwal et al. 2006; Das and Banerjee 2011) have shown that their algorithms are guaranteed to perform competitively with certain families of adaptive portfolios even in an adversarial market in a costless setting without making any statistical assumptions regarding the movement of the stocks. (Das, Johnson, and Banerjee 2013) considers transaction costs in their formulation, and can be posed as a special case of our framework which we are about to present. Heuristics have been shown to have empirical advantages in certain settings (Borodin, El-Yaniv, and Gogan 2004; Li et al. 2011;

Li and Hoi 2014). However, none of the existing work in online portfolio selection has attempted to investigate or take advantage of the group structure (pre-specified or modeled) within the stocks in their algorithm setting.

3.2 Problem Formulation

We present a general formulation for our online lazy algorithm with group sparsity and go on to show how the portfolio selection problem is a special case of our setting. In an online lazy setting, the optimization proceeds in rounds where in round t the algorithm has to pick a solution from a feasible set, $\mathbf{p}_t \in \mathcal{P}$, such that it is sparse in the number of groups picked and close to the previous solution \mathbf{p}_{t-1} . Nature then reveals a convex loss function, f_t , and we observe its value $f_t(\mathbf{p}_t)$. Ideally, over T rounds we would like to minimize the quantity,

$$\sum_{t=1}^T \{f_t(\mathbf{p}_t) + \lambda_1 \Omega(\mathbf{p}_t)\} + \lambda_2 \sum_{t=1}^{T-1} \|\mathbf{p}_{t+1} - \mathbf{p}_t\|_1. \quad (1)$$

In (1), the $\Omega(\cdot)$ penalty function can be any group norm which will ensure group sparsity. We adopt the ‘groupwise’ ℓ_2 -norm used in group lasso (Yuan and Lin 2006; Friedman, Hastie, and Tibshirani 2010) as our regularizer, i.e.,

$$\lambda_1 \Omega(\mathbf{p}) = \lambda_1 \sum_{g=1}^{\mathcal{G}} w_g \|\mathbf{p}_{|g}\|. \quad (2)$$

We call \mathcal{G} our set of groups and $\forall g \in \mathcal{G}, g \subseteq \{1, \dots, n\}$. $\mathbf{p}_{|g}$ is the vector whose coordinates are equal to those of \mathbf{p} for indices in the set g . $(w_g)_{g \in \mathcal{G}}$ denotes positive weights and $\|\cdot\|$ is the euclidean norm. To introduce group sparsity, it is also possible to impose other joint regularization on the weight, e.g. the $\ell_{1,\infty}$ -norm (Quattoni et al. 2009). We consider the case where the groups are disjoint, i.e. \mathcal{G} is separable over $\{1, \dots, n\}$, however our framework and algorithm can be extended to the overlapping group lasso case (Jacob, Obozinski, and Vert 2009). The ℓ_1 penalty term in (1) ensures that the updates to the solution \mathbf{p}_t are lazy.

Absolute minimization of (1) is not reasonable because we do not know the sequence of f_t *a priori*. If the f_t s are known, (1) reduces to a batch optimization problem: a special case is the fused group lasso when f_t is quadratic (Friedman, Hastie, and Tibshirani 2010; Tibshirani et al. 2005) or TV regularization (Rudin, Osher, and Fatemi 1992). Alternatively, over T iterations we intend to select a sequence of \mathbf{p}_t such that the following *regret bound* is sub-linear in T ,

$$R_T = \sum_{t=1}^T \psi_t(\mathbf{p}_t) - \min_{\mathbf{p}^* \in \mathcal{P}} \sum_{t=1}^T \psi_t(\mathbf{p}^*) \leq o(T) \quad (3)$$

where $\psi_t(\mathbf{p}) = f_t(\mathbf{p}) + \lambda_1 \Omega(\mathbf{p}) + \lambda_2 \|\mathbf{p} - \mathbf{p}_{t-1}\|_1$ is non-smooth and \mathbf{p}^* is the minimizer of $\sum_{t=1}^T \psi_t$ in hindsight. Note that while the \mathbf{p}_t s can change over time, \mathbf{p}^* is fixed. That is, the minimizer, $\mathbf{p}^* = \operatorname{argmin}_{\mathbf{p}} \sum_{t=1}^T \psi_t(\mathbf{p}) = \operatorname{argmin}_{\mathbf{p}} \sum_{t=1}^T f_t(\mathbf{p}) + \lambda_1 \Omega(\mathbf{p})$, since it incurs zero ℓ_1 penalty in every iteration.

Additionally, we examine the case where the comparator class can also change over time. In particular, we consider

the sequence $\{\mathbf{p}_1^*, \dots, \mathbf{p}_T^*\}$ which has the power of hindsight. Then, over T iterations we ensure that the following *shifting* regret bound is sub-linear in T :

$$\sum_{t=1}^T \psi_t(\mathbf{p}_t) - \min_{\mathbf{p}_1^*, \dots, \mathbf{p}_T^*} \sum_{t=1}^T \psi_t(\mathbf{p}_t^*) \leq o(T) + c \text{size}(\langle \mathbf{p}_1^*, \dots, \mathbf{p}_T^* \rangle), \quad (4)$$

where $\text{size}(\langle \mathbf{p}_1^*, \dots, \mathbf{p}_T^* \rangle)$ intuitively measures the amount of shifting that occurs in the best sequence of solutions in hindsight and c is a constant. Online portfolio selection with group sparsity can now be viewed as a special case of the above setting where $f_t(\mathbf{p}) = -\log(\mathbf{p}^T \mathbf{x}_t)$ and the ℓ_1 penalty term on the difference of two consecutive portfolios measures the fraction of wealth traded. The parameters λ_1 controls how many groups are selected (setting $\lambda_1 = 0$ reduces to (Das, Johnson, and Banerjee 2013)) and λ_2 controls the amount that can be traded every day.

4 Online Lazy Updates with Group Sparsity

We now present our Online Lazy Updates with Group Sparsity (OLU-GS) algorithm. In the sequel, we show that using the solutions generated by OLU-GS, we can achieve sub-linear regret for the non-shifting (3) and shifting case (4). At the beginning of day $t + 1$, we find a new solution \mathbf{p}_{t+1} by minimizing the following:

$$\mathbf{p}_{t+1} = \underset{\mathbf{p} \in \Delta_n}{\text{argmin}} \langle \nabla f_t(\mathbf{p}_t), \mathbf{p} \rangle + \lambda_1 \Omega(\mathbf{p}) + \lambda_2 \|\mathbf{p} - \mathbf{p}_t\|_1 + \frac{1}{2\eta} \|\mathbf{p} - \mathbf{p}_t\|_2^2, \quad (5)$$

where we have linearized f_t around \mathbf{p}_t . Our objective function in (5) is composite with smooth and non-smooth terms with the probability simplex as a constraint set. Although there is literature on solving composite functions (J. Duchi et al. 2010; Yu 2013), composite functions with linear constraints have not been adequately investigated. We propose an Alternating Direction Method of Multipliers (ADMM) (Boyd et al. 2011) based efficient primal-dual algorithm to solve (5). ADMM has been applied in many large scale statistics and machine learning problems because of its computational benefits and fast convergence in practice (Boyd et al. 2011). We rewrite (5) in ADMM form by introducing auxiliary variables \mathbf{y} and \mathbf{z}

$$\underset{\mathbf{p} \in \Delta_n, \mathbf{p}=\mathbf{y}, \mathbf{p}-\mathbf{p}_t=\mathbf{z}}{\text{argmin}} \langle \nabla f_t(\mathbf{p}_t), \mathbf{p} \rangle + \lambda_1 \Omega(\mathbf{y}) + \lambda_2 \|\mathbf{z}\|_1 + \frac{1}{2\eta} \|\mathbf{p} - \mathbf{p}_t\|_2^2. \quad (6)$$

Using variable splitting, we write the *augmented lagrangian* as, $L(\mathbf{p}, \mathbf{y}, \mathbf{z}, \mathbf{w}, \mathbf{v}) = \langle \nabla f_t(\mathbf{p}_t), \mathbf{p} \rangle + \lambda_1 \Omega(\mathbf{y}) + \lambda_2 \|\mathbf{z}\|_1 + \frac{1}{2\eta} \|\mathbf{p} - \mathbf{p}_t\|_2^2 + \frac{\beta}{2} \|\mathbf{p} - \mathbf{y} + \mathbf{w}\|_2^2 + \frac{\beta}{2} \|\mathbf{p} - \mathbf{p}_t - \mathbf{z} + \mathbf{v}\|_2^2$, where \mathbf{w} and \mathbf{v} are the scaled dual variables and $\mathbf{p} \in \Delta_n$. Splitting the variables as we do in (6) has two advantages. Firstly, we will show there is a closed form solution for each update. Secondly, the updates for \mathbf{y} and \mathbf{z} can be done in parallel and the same is true for the scaled dual variables \mathbf{w} and \mathbf{v} .

Algorithm 1 OLU-GS Algorithm with ADMM

- 1: Input $\mathbf{p}_t, \mathbf{x}_t, \nabla f_t(\mathbf{p}_t), \mathcal{G}, \lambda_1, \lambda_2, \eta, \beta$
- 2: Initialize $\mathbf{p}, \mathbf{y}, \mathbf{z}, \mathbf{w}, \mathbf{v} \in \mathbf{0}^n, k = 0$
- 3: Set $\hat{a} = \frac{1+\eta\beta}{1+2\eta\beta}$ and $\hat{b} = \frac{\eta\beta}{1+2\eta\beta}$ and $\hat{c} = \frac{\eta}{1+2\eta\beta}$
- 4: ADMM iterations

$$\mathbf{p}_{t+1}^{k+1} = \prod_{\mathbf{p} \in \Delta_n} \left\{ \hat{a} \mathbf{p}_t - \hat{c} \nabla f_t(\mathbf{p}_t) + \hat{b} (\mathbf{y}^{(k)} + \mathbf{z}^{(k)} - \mathbf{w}^{(k)} - \mathbf{v}^{(k)}) \right\} \quad (12)$$

$$\mathbf{y}_{|g}^{(k+1)} = S_{\lambda_1/\beta}(\mathbf{p}_{|g}^{(k+1)} - \mathbf{p}_{t|g} + \mathbf{w}_{|g}^{(k)}), \forall g \in \mathcal{G} \quad (13)$$

$$\mathbf{z}^{(k+1)} = S_{\lambda_2/\beta}(\mathbf{p}_{t+1}^{k+1} - \mathbf{p}_t + \mathbf{v}^{(k)}) \quad (14)$$

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} + (\mathbf{p}^{(k+1)} - \mathbf{y}^{(k+1)} + \mathbf{w}^{(k)}) \quad (15)$$

$$\mathbf{v}^{(k+1)} = \mathbf{v}^{(k)} + (\mathbf{p}^{(k+1)} - \mathbf{p}_t - \mathbf{z}^{(k+1)} + \mathbf{v}^{(k)}) \quad (16)$$

where \prod_{Δ_n} is the projection to the simplex and S_ρ is the shrinkage operator.

- 5: Continue until **Stopping Criteria** is satisfied
-

ADMM consists of the following iterations for solving \mathbf{p}_{t+1} ,

$$\mathbf{p}_{t+1}^{(k+1)} = \underset{\mathbf{p} \in \Delta_n}{\text{argmin}} \langle \nabla f_t(\mathbf{p}_t), \mathbf{p} \rangle + \frac{1}{2\eta} \|\mathbf{p} - \mathbf{p}_t\|_2^2 \quad (7)$$

$$+ \frac{\beta}{2} \|\mathbf{p} - \mathbf{y}^{(k)} + \mathbf{w}^{(k)}\|_2^2 + \frac{\beta}{2} \|\mathbf{p} - \mathbf{p}_t - \mathbf{z}^{(k)} + \mathbf{v}^{(k)}\|_2^2$$

$$\mathbf{y}^{(k+1)} = \underset{\mathbf{y}}{\text{argmin}} \lambda_1 \Omega(\mathbf{y}) + \frac{\beta}{2} \|\mathbf{p}_{t+1}^{(k+1)} - \mathbf{y} + \mathbf{w}^{(k)}\|_2^2 \quad (8)$$

$$\mathbf{z}^{(k+1)} = \underset{\mathbf{z}}{\text{argmin}} \lambda_2 \|\mathbf{z}\|_1 + \frac{\beta}{2} \|\mathbf{p}_{t+1}^{(k+1)} - \mathbf{p}_t - \mathbf{z} + \mathbf{v}^{(k)}\|_2^2 \quad (9)$$

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} + (\mathbf{p}_{t+1}^{(k+1)} - \mathbf{y}^{(k+1)}) \quad (10)$$

$$\mathbf{v}^{(k+1)} = \mathbf{v}^{(k)} + (\mathbf{p}_{t+1}^{(k+1)} - \mathbf{p}_t - \mathbf{z}^{(k+1)}). \quad (11)$$

p-update: We take the derivative of (7) *w.r.t.* \mathbf{p} and set it to zero to get a closed form update of \mathbf{p} as $\nabla f_t(\mathbf{p}_t) + \frac{1}{\eta}(\mathbf{p} - \mathbf{p}_t) + \beta(\mathbf{p} - \mathbf{y}^{(k)} + \mathbf{w}^{(k)}) + \beta(\mathbf{p} - \mathbf{p}_t - \mathbf{z}^{(k)} + \mathbf{v}^{(k)}) = 0$. Rearranging this and setting $\hat{a} = \frac{1+\eta\beta}{1+2\eta\beta}$, $\hat{b} = \frac{\eta\beta}{1+2\eta\beta}$, and $\hat{c} = \frac{\eta}{1+2\eta\beta}$, we get (12). $\prod_{\mathbf{p} \in \Delta_n}$ is the projection operator which is carried out as in (Duchi, Singer, and Chandra 2008).

y-update: We can rewrite (8) as

$$\mathbf{y}^{(k+1)} = \underset{\mathbf{y}}{\text{argmin}} \frac{1}{2} \|\mathbf{p}^{(k+1)} + \mathbf{w}^{(k)} - \mathbf{y}\|_2^2 + \frac{\lambda_1}{\beta} \Omega(\mathbf{y}). \quad (17)$$

When $\Omega(\cdot)$ is a group lasso penalty with l_2 -norm, with \mathcal{G} being a partition of $\{1, \dots, n\}$, (17) is *separable* in every group, and the solution is a generalization of the soft thresholding operator to groups of variables (Jenatton et al. 2010):

$$\forall g \in \mathcal{G}, \mathbf{y}_{|g} = \begin{cases} 0 & \text{if } \|\mathbf{q}_{|g}\|_2 \leq \tilde{\lambda} \\ \frac{\|\mathbf{q}_{|g}\|_2 - \tilde{\lambda}}{\|\mathbf{q}_{|g}\|_2} \mathbf{q}_{|g} & \text{otherwise} \end{cases} \quad (18)$$

where $\mathbf{q} = \mathbf{p}^{(k+1)} + \mathbf{w}^{(k)}$, $\tilde{\lambda} = \frac{\lambda_1}{\beta}$ and $\mathbf{y}_{|g}$ is a vector of size n whose coordinates are equal to those of \mathbf{y} for indices in the set g . We obtain a closed form solution for \mathbf{z}^{k+1} by using the soft-thresholding operator $S_\rho(a)$ (Boyd et al. 2011). The

Algorithm 2 Portfolio Selection with Group Sparsity

- 1: Input $\mathcal{G}, \lambda_1, \lambda_2, \eta, \beta$; Transaction cost γ
- 2: Initialize $p_{1,g} = \frac{1}{|\mathcal{G}|}, g = 1, \dots, |\mathcal{G}|; p_0 = p_1; S_0^\gamma = 1$
- 3: For $t = 1, \dots, T$
- 4: Receive \mathbf{x}_t , the vector of price relatives
- 5: Compute cumulative wealth: $S_t^\gamma = S_{t-1}^\gamma \times (\mathbf{p}_t^T \mathbf{x}_t) - \gamma \times S_{t-1}^\gamma \times \|\mathbf{p}_t - \mathbf{p}_{t-1}\|_1$
- 6: Update portfolio:
- 7: $\mathbf{p}_{t+1} = \text{OLU-GS}(\mathbf{p}_t, \mathbf{x}_t, -\frac{\mathbf{x}_t}{\mathbf{p}_t^T \mathbf{x}_t}, \mathcal{G}, \lambda_1, \lambda_2, \eta, \beta)$
- 8: end for

updates for $\mathbf{w}^{(k+1)}$ and $\mathbf{v}^{(k+1)}$ are already in closed form. We iterate over the updates until convergence according to the stopping criteria in (Boyd et al. 2011). Algorithm 1 summarizes the ADMM updates for OLU-GS.

Algorithm 2 outlines our algorithm for computing the transaction cost-adjusted wealth S_T^γ , where γ is a proportional transaction cost (Das, Johnson, and Banerjee 2013).

5 Analysis

In this paper, we consider updates of the following form:

$$\mathbf{p}_{t+1} = \underset{\mathbf{p} \in \mathcal{P}}{\operatorname{argmin}} \left\{ \eta \langle \nabla f_t(\mathbf{p}_t), \mathbf{p} \rangle + \eta r(\mathbf{p}) + \eta \lambda_2 \|\mathbf{p} - \mathbf{p}_t\|_1 + d_\phi(\mathbf{p}, \mathbf{p}_t) \right\}, \quad (19)$$

where r is any non-smooth regularizer and d_ϕ is a Bregman divergence. Note, that our analysis is different from (J. Duchi et al. 2010), because of the presence of the $\|\mathbf{p} - \mathbf{p}_t\|_1$ term in our updates. We call such class of updates: Composite Objective with Lazy Updates. Our OLU-GS updates in Section 4 are a special case of (19), $r = \lambda_1 \Omega$ and $d_\phi(\mathbf{p}, \mathbf{p}_t) = \frac{1}{2} \|\mathbf{p} - \mathbf{p}_t\|_2^2$. In this section we prove regret bounds for two cases: when (1) f_t s are general convex functions and (2) f_t s are strongly convex. Moreover, we prove shifting bounds, where the comparator class can itself change over time. The proofs of the theorems will be made available in a longer version of the paper.

(a) General Convex Functions: We assume that f_t are general convex functions with bounded (sub)gradients, i.e., for any $\hat{g} \in \partial f_t(\mathbf{p})$ we have $\|\hat{g}\| \leq G$.

Theorem 1 *Let the sequence of $\{\mathbf{p}_t\}$ be defined by (19). Let f_t be a Lipschitz continuous function for which $\|\nabla f_t(\mathbf{p}_t)\|_2^2 \leq G$. Then, by choosing $\eta \propto \frac{1}{\sqrt{T}}$ and $\lambda_2 \propto \frac{1}{\sqrt{T}}$, we have*

$$\sum_{t=1}^T [f_t(\mathbf{p}_t) + r(\mathbf{p}_t) + \lambda_2 \|\mathbf{p}_{t+1} - \mathbf{p}_t\|_1 - f(\mathbf{p}^*) - r(\mathbf{p}^*)] \leq O(\sqrt{T}). \quad (20)$$

(b) Strongly Convex Functions: We assume that f_t are all β -strongly convex functions. For any $(\mathbf{p}, \mathbf{p}_t)$, $f_t(\mathbf{p}) \geq f_t(\mathbf{p}_t) + \langle \mathbf{p} - \mathbf{p}_t, \nabla f_t(\mathbf{p}_t) \rangle + \frac{\beta}{2} \|\mathbf{p} - \mathbf{p}_t\|_2^2$.

Sector	Example Companies
Consumer Discretionary	Nike Inc., Target Corp.
Consumer Staples	Costco Co., Beam Inc.
Energy	Chevron Corp., Noble Corp.
Financials	Equifax Inc., AFLAC Inc.
Health Care	Cerner, Pfizer Inc.
Industrials	Raytheon Co., 3M Co.
Information Tech	Apple Inc., Dell Inc.
Materials	Alcoa Inc., Ecolab Inc.
Utilities	AGL Resources, AES Corp.

Table 1: Overview of GICS sectors used in our dataset.

Theorem 2 *Let the sequence of $\{\mathbf{p}_t\}$ be defined by (19). Let f_t be all β -strongly convex. Then, for any $\lambda_2 < \beta/4$, choosing $\eta_t = \frac{2}{\gamma t}$, where $\gamma \in (0, \frac{\beta}{4} - \lambda_2]$, we have*

$$\sum_{t=1}^T [f_t(\mathbf{p}_t) + r(\mathbf{p}_t) + \lambda_2 \|\mathbf{p}_{t+1} - \mathbf{p}_t\|_1 - f(\mathbf{p}^*) - r(\mathbf{p}^*)] \leq O(\log(T)). \quad (21)$$

(c) Shifting Bounds: For general convex functions as defined in Section 5, we can show the following *shifting regret* bound.

Theorem 3 *Let $\{\mathbf{p}_1^*, \dots, \mathbf{p}_T^*\}$ be the best sequence obtained by minimizing (1). For, $\|\nabla \phi(\mathbf{p}_{t+1})\|_\infty \leq \zeta$ we have*

$$\sum_{t=1}^T [f_t(\mathbf{p}_t) + r(\mathbf{p}_t) - f_t(\mathbf{p}_t^*) - r(\mathbf{p}_t^*)] + \lambda_2 \sum_{t=1}^{T-1} [\|\mathbf{p}_{t+1} - \mathbf{p}_t\|_1 - \|\mathbf{p}_{t+1}^* - \mathbf{p}_t^*\|_1] \leq O(\sqrt{T}) + \zeta \sum_{t=1}^{T-1} \|\mathbf{p}_t^* - \mathbf{p}_{t+1}^*\|_1. \quad (22)$$

Here, $\sum_{t=1}^{T-1} \|\mathbf{p}_t^* - \mathbf{p}_{t+1}^*\|_1$ measures the amount of shifting that occurs for the best sequence of \mathbf{p}_t^* s.

6 Experiments and Results

Dataset: The experiments were conducted on data taken from the New York Stock Exchange (NYSE) and Standard & Poor's 500 (S&P 500) stock market index. The NYSE dataset (Helmbold et al. 1998; Agarwal et al. 2006; Borodin, El-Yaniv, and Gogan 2004; Cover 1991) consists of 36 stocks with data accumulated over a period of 22 years from July 3, 1962 to December 31, 1984. The dataset captures the bear market that lasted between January 1973 and December 1974. The S&P500 dataset consists of 258 stocks with data accumulated over a period of 22 years from 1991 to 2012. The dataset captures the bull and bear markets of recent times such as the dot-com bubble which occurred between 1997-2000, the following bubble burst starting in March 2000 and continuing through 2002, and the recent financial and housing bubble burst between 2007-2009.

We used the Global Industry Classification Standard to group the stocks in the datasets into their designated sectors.




Figure 1: As λ_1 increases the (a) total group lasso value and (b) number of active group changes decrease.

This resulted in 8 sectors and 30 stocks in the NYSE dataset and 9 sectors and 243 stocks in the S&P500 dataset. Table 1 shows the sectors represented in the two datasets, and a couple representative companies from each sector.

Methodology and Parameter Setting: In all experiments we started with \$1 as our initial investment and an initial portfolio uniformly distributed over the groups to avoid group bias. We use OLU-GS to obtain our portfolios sequentially and compute the transaction cost-adjusted wealth for each day. The parameters consist of λ_1 : weight on group sparsity norm, λ_2 : lazy updates weight, η : weight on the ℓ_2 norm, and β : the parameter for the augmentation term. For all our experiments, we set $\beta = 2$ which we found to give reasonable accuracy and use group lasso for group sparsity.

Since the two datasets are very different in nature (stock composition and duration), we experimented extensively with a large range of λ_1 , λ_2 , and η values from $1e-9$ to 1 to observe their effect on group sparsity and lazy updates to our portfolio. Moreover, we chose a reasonable range of γ values between 0% and 2% to compute the proportional transaction costs incurred due to the portfolio update every day. We have illustrated some of our results with representative plots from either the NYSE or S&P500 dataset.

We use the wealth obtained (without transaction costs) from the EG algorithm with experimentally tuned parameters, a Buy-and-Hold strategy, and the best single stock as benchmarks for our experiments with initial investments of \$1. EG has been shown to outperform a uniform constantly rebalanced portfolio (Helmbold et al. 1998; Das and Banerjee 2011). For the Buy-and-Hold case we start with a uniformly distributed portfolio and do a hold on the positions thereafter (i.e. no trades). For the best single stock case we observe how the market performs and select the stock that has accumulated the most wealth at the end of the period. Note, in a real world situation, this strategy is infeasible since it is not possible to know the best stock *a priori*.

6.1 Effect of λ_1 for Group Sparsity ($\Omega(\mathbf{p})$)

The regularization parameter λ_1 for the group lasso term ($\Omega(\mathbf{p})$) is varied from $[1e-9, 1]$ to obtain different levels of group sparsity. The value of λ_1 has a strong effect on (a) the total group lasso penalty value, (b) the number of active groups, and (c) which groups are active.

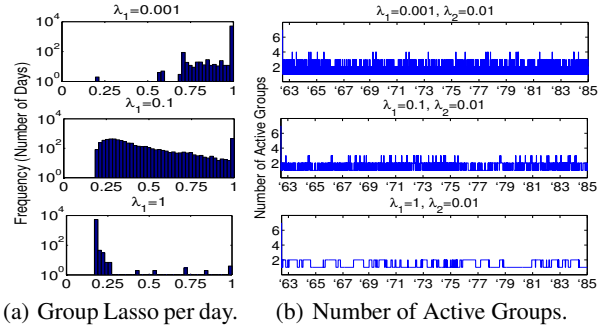


Figure 2: As λ_1 increases the number of days with high group lasso value and the number of active groups decrease.

(a) Total Group Lasso penalty: Figure 1(a) plots the value of the total group lasso penalty ($\sum_{t=1}^T \Omega(\mathbf{p}_t)$) as we increase λ_1 , keeping λ_2 and η fixed. For both the NYSE and S&P500 datasets, we observe that $\sum_{t=1}^T \Omega(\mathbf{p}_t)$ decreases as we increase λ_1 , which is in conformance with our objective. Since the two datasets are different in terms of the total number of stocks and the number of stocks composing each sector, Figure 1(a) specifically illustrates how to choose λ_1 to attain a desired level of sparsity for each of the datasets. Figure 2(a) plots a histogram of the total per day group lasso penalty with increasing λ_1 values for the S&P500. It is fairly evident that there is a decrease in the number of days with high group lasso penalty as λ_1 increases.

(b) Active Groups: We compute the *active groups* each day by selecting the groups in which the majority (80%) of the wealth is invested. Figure 2(b) plots the number of active groups per day for the NYSE dataset. With $\lambda_1 = 1e-3$, OLU-GS picks up to 4 groups on a particular day. For a higher value of $\lambda_1 = 1$ a maximum of 2 groups are selected to invest in. In particular, the two sectors picked are Basic Materials and Consumer Discretionary.

(c) Active Groups Changes: Figure 1(b) plots the total number of times the *active groups* change for the NYSE dataset (over 22 years) as λ_1 increases. We consider an active group change as anytime the group composition changes. The individual line plots indicate different values of λ_2 and η . For λ_1 between $1e-6$ to $1e-2$: with low values for λ_2 , the total number of changes in the active groups are quite high but for a higher value of $\lambda_2 = 1e-2$ we see a decrease in the number of active group changes illustrating the portfolio laziness. With larger values of $\lambda_1 \geq 1e-2$ we see a dramatic drop in the number of active group changes and high λ_2 values only reemphasize this behavior.

6.2 Wealth and Group Sparsity

To evaluate the practical application of our proposed algorithm, we now analyze its performance when calculating the transaction cost-adjusted cumulative wealth. Figure 4 shows how the choice of different λ_1 values affect the transaction cost-adjusted cumulative wealth for the NYSE dataset (for a fixed λ_2 and η value). Figure 5 demonstrates that there exists a combination of λ_1 and λ_2 values which make an optimal





Figure 3: Picking non-cyclic sectors, sectors that tend to perform well during economic downturns, during the 1970s and dot-com bear markets and cyclic sectors, sectors that perform well during economic booms, during the dot-com bull market.

choice between group sparsity and lazy updates.

EG, Buy-and-Hold, Best Single Stock: We compared the total wealth without transaction costs of OLU-GS with that of EG, a Buy-and-Hold strategy, and the best performing single stock. These strategies are plotted as horizontal lines and we can see that for the NYSE dataset EG returns \$20.89, Buy-and-Hold returns \$20.88, and the best single stock, Phillip Morris (MO), returns \$54.14. In comparison, OLU-GS returns \$71.18 without transaction costs. OLU-GS returns over 3x as much wealth for the NYSE dataset as EG or Buy-and-Hold do and about \$15 more than the best stock. Figure 4 also shows that OLU-GS is able to return more wealth than EG and Buy-and-Hold with reasonable transaction costs (0.001%, 0.005%, and 0.01%).

6.3 Switching Sectors

We desire that OLU-GS is able to identify the best sectors automatically. We illustrate the strength of OLU-GS in selecting the best sectors with two examples. A recurring trend that we observe from our experiments with both the NYSE and S&P500 datasets is that OLU-GS selects stocks in Consumer Staples during the bear markets. Figure 3(a) clearly shows that OLU-GS selects and invests in this defensive sector during the historical bear markets of 1969-1971 and 1975-1977. Another example of a defensive or non-cyclic sector is Utilities. Figure 3(b) shows that the weight on the Utilities sector sees a considerable increase during the dot-com crash. This is interesting because unlike other areas of the economy, even during bear markets, the demand for Consumer Staples and Utilities do not slow down. These sectors




Figure 4: Transaction cost-adjusted wealth. OLU-GS returns more than competing algorithms even with transaction costs.

consist of stocks which are defensive in nature and usually outperform the S&P500 Index during bearish markets and under-perform during bullish markets. Sectors like Information Technology and Financials comprise of cyclical stocks which are sensitive to market movements and can take advantage of the bullish markets. In Figure 3(b), we see that Information Technology sector is picked up during the bullish markets which preceded the dot-com bubble.

7 Conclusion

In this paper, we have developed a general lazy online learning with group sparsity framework and an online learning algorithm (OLU-GS) and show how it can be applied to the problem of online portfolio selection with sector information and transaction costs. Our analysis shows that OLU-GS is competitive with reasonable fixed and shifting strategies which have the power of hindsight. Our experimental results illustrate the behavior of group sparsity and lazy updates and show that OLU-GS is able to outperform baseline strategies with reasonable transaction costs. Finally, we demonstrate that OLU-GS is able to select the best performing sectors during different economic conditions. In the future we wish to explore the possibility of learning the group structure from the data itself.

Acknowledgements: The research was supported by NSF CAREER award IIS-0953274, NSF grants IIS-0916750, and IIS-1029711, and NASA grant NNX12AQ39A.




Figure 5: Wealth as a function of λ_1 and λ_2 . There exists a parameter range that gives good wealth performance.

References

- Agarwal, A.; Hazan, E.; Kale, S.; and Schapire, R. 2006. Algorithms for portfolio management based on the newton method. In *Proceedings of the 23rd International Conference on Machine Learning*, 9–16.
- Arouri, M. E. H., and Nguyen, D. K. 2010. Oil prices, stock markets and portfolio investment: Evidence from sector analysis in europe over the last decade. *Energy Policy* 38(8):4528–4539.
- Blum, A., and Kalai, A. 1997. Universal portfolios with and without transaction costs. In *Proceedings of the 10th Annual Conference on Learning Theory*.
- Borodin, A.; El-Yaniv, R.; and Gogan, V. 2004. Can we learn to beat the best stock. *Journal of Artificial Intelligence Research* 21:579–594.
- Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; and Eckstein, J. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3(1):1–122.
- Cesa-Bianchi, N., and Lugosi, G. 2006. *Prediction, Learning, and Games*. Cambridge University Press.
- Cover, T. M., and Ordentlich, E. 1996. Universal portfolios with side information. *IEEE Transactions on Information Theory* 42:348–363.
- Cover, T. 1991. Universal portfolios. *Mathematical Finance* 1:1–29.
- Das, P., and Banerjee, A. 2011. Meta optimization and its application to portfolio selection. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Das, P.; Johnson, N.; and Banerjee, A. 2013. Online lazy updates for portfolio selection with transaction costs. In *Proceedings of the 27th Association for the Advancement of Artificial Intelligence Conference*.
- Davis, M., and Norman, A. 1990. Portfolio selection with transaction costs. *Mathematics of Operations Research* 15(4):676–713.
- Duchi, J.; Singer, Y.; and Chandra, T. 2008. Efficient projections onto the l_1 -ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning*.
- Friedman, J.; Hastie, T.; and Tibshirani, R. 2010. A note on the group lasso and a sparse group lasso. *Arxiv preprint arXiv:1001.0736*.
- Helmbold, D.; Scahpire, E.; Singer, Y.; and Warmuth, M. 1998. Online portfolio selection using multiplicative weights. *Mathematical Finance* 8(4):325–347.
- J. Duchi, J.; Shalev-Shwartz, S.; Singer, Y.; and Tewari, A. 2010. Composite objective mirror descent. In *Proceedings of the 23rd Annual Conference on Learning Theory*, 14–26.
- Jacob, L.; Obozinski, G.; and Vert, J. 2009. Group lasso with overlap and graph lasso. In *Proceedings of the 26th International Conference on Machine Learning*.
- Jenatton, R.; Mairal, J.; Obozinski, G.; and Bach, F. 2010. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research* 2297–2334.
- Li, B., and Hoi, S. C. 2014. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)* 46(3).
- Li, B.; Hoi, S.; Zhao, P.; and Gopalkrishnan, V. 2011. Confidence weighted mean reversion strategy for on-line portfolio selection. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*.
- Li, Q.; Vassalou, M.; and Xing, Y. 2006. Sector investment growth rates and the cross section of equity returns*. *The Journal of Business* 79(3):1637–1665.
- Quattoni, A.; Carreras, X.; Collins, M.; and Darrell, T. 2009. An efficient projection for $\ell_{1,\infty}$ regularization. In *Proceedings of the 26th International Conference on Machine Learning*.
- Rudin, L. I.; Osher, S.; and Fatemi, E. 1992. Nonlinear Total Variation Based Noise Removal Algorithms. *Physica D* 60:259–268.
- Tibshirani, R.; Saunders, M.; Rosset, S.; Zhu, J.; and Knight, K. 2005. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society Series B* 91–108.
- Yu, Y. 2013. Better approximation and faster algorithm using the proximal average. In *Advances in Neural Information Processing Systems*, 458–466.
- Yuan, M., and Lin, Y. 2006. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B* 68(1):49–67.