

---

# Efficient Second Order Online Learning by Sketching

---

**Haipeng Luo**

Princeton University, Princeton, NJ USA  
haipengl@cs.princeton.edu

**Alekh Agarwal**

Microsoft Research, New York, NY USA  
alekha@microsoft.com

**Nicolò Cesa-Bianchi**

Università degli Studi di Milano, Italy  
nicolo.cesa-bianchi@unimi.it

**John Langford**

Microsoft Research, New York, NY USA  
jcl@microsoft.com

## Abstract

We propose Sketched Online Newton (SON), an online second order learning algorithm that enjoys substantially improved regret guarantees for ill-conditioned data. SON is an enhanced version of the Online Newton Step, which, via sketching techniques enjoys a running time linear in the dimension and sketch size. We further develop sparse forms of the sketching methods (such as Oja’s rule), making the computation linear in the sparsity of features. Together, the algorithm eliminates all computational obstacles in previous second order online learning approaches.

## 1 Introduction

Online learning methods are highly successful at rapidly reducing the test error on large, high-dimensional datasets. First order methods are particularly attractive in such problems as they typically enjoy computational complexity linear in the input size. However, the convergence of these methods crucially depends on the geometry of the data; for instance, running the same algorithm on a rotated set of examples can return vastly inferior results. See Fig. 1 for an illustration.

Second order algorithms such as Online Newton Step [18] have the attractive property of being invariant to linear transformations of the data, but typically require space and update time quadratic in the number of dimensions. Furthermore, the dependence on dimension is not improved even if the examples are sparse. These issues lead to the key question in our work: *Can we develop (approximately) second order online learning algorithms with efficient updates?* We show that the answer is “yes” by developing efficient sketched second order methods with regret guarantees. Specifically, the three main contributions of this work are:

**1. Invariant learning setting and optimal algorithms (Section 2).** The typical online regret minimization setting evaluates against a benchmark that is bounded in some fixed norm (such as the  $\ell_2$ -norm), implicitly putting the problem in a nice geometry. However, if all the features are scaled down, it is desirable to compare with accordingly larger weights, which is precluded by an apriori fixed norm bound. We study an invariant learning setting similar to the paper [33] which compares the learner to a benchmark only constrained to generate bounded predictions on the sequence of examples. We show that a variant of the Online Newton Step [18], while quadratic in computation, stays regret-optimal with a nearly matching lower bound in this more general setting.

**2. Improved efficiency via sketching (Section 3).** To overcome the quadratic running time, we next develop sketched variants of the Newton update, approximating the second order information using a small number of carefully chosen directions, called a *sketch*. While the idea of data sketching is widely studied [36], as far as we know our work is the first one to apply it to a general adversarial

online learning setting and provide rigorous regret guarantees. Three different sketching methods are considered: Random Projections [1, 19], Frequent Directions [12, 23], and Oja’s algorithm [28, 29], all of which allow linear running time per round. For the first two methods, we prove regret bounds similar to the full second order update whenever the sketch-size is large enough. Our analysis makes it easy to plug in other sketching and online PCA methods (e.g. [11]).

**3. Sparse updates (Section 4).** For practical implementation, we further develop sparse versions of these updates with a running time linear in the sparsity of the examples. The main challenge here is that even if examples are sparse, the sketch matrix still quickly becomes dense. These are the first known sparse implementations of the Frequent Directions<sup>1</sup> and Oja’s algorithm, and require new sparse eigen computation routines that may be of independent interest.

Empirically, we evaluate our algorithm using the sparse Oja sketch (called Oja-SON) against first order methods such as diagonalized ADAGRAD [6, 25] on both ill-conditioned synthetic and a suite of real-world datasets. As Fig. 1 shows for a synthetic problem, we observe substantial performance gains as data conditioning worsens. On the real-world datasets, we find improvements in some instances, while observing no substantial second-order signal in the others.




Figure 1: Error rate of SON using Oja’s sketch, and ADAGRAD on a synthetic ill-conditioned problem.  $m$  is the sketch size ( $m = 0$  is Online Gradient,  $m = d$  resembles Online Newton). SON is nearly invariant to condition number for  $m = 10$ .

**Related work** Our online learning setting is closest to the one proposed in [33], which studies scale-invariant algorithms, a special case of the invariance property considered here (see also [31, Section 5]). Computational efficiency, a main concern in this work, is not a problem there since each coordinate is scaled independently. Orabona and Pál [30] study unrelated notions of invariance. Gao et al. [9] study a specific randomized sketching method for a special online learning setting.

The L-BFGS algorithm [24] has recently been studied in the stochastic setting<sup>2</sup> [3, 26, 27, 34, 35], but has strong assumptions with pessimistic rates in theory and reliance on the use of large mini-batches empirically. Recent works [7, 15, 14, 32] employ sketching in stochastic optimization, but do not provide sparse implementations or extend in an obvious manner to the online setting. The Frank-Wolfe algorithm [8, 20] is also invariant to linear transformations, but with worse regret bounds [17] without further assumptions and modifications [10].

**Notation** Vectors are represented by bold letters (e.g.,  $\mathbf{x}, \mathbf{w}, \dots$ ) and matrices by capital letters (e.g.,  $M, A, \dots$ ).  $M_{i,j}$  denotes the  $(i, j)$  entry of matrix  $M$ .  $\mathbf{I}_d$  represents the  $d \times d$  identity matrix,  $\mathbf{0}_{m \times d}$  represents the  $m \times d$  matrix of zeroes, and  $\text{diag}\{\mathbf{x}\}$  represents a diagonal matrix with  $\mathbf{x}$  on the diagonal.  $\lambda_i(A)$  denotes the  $i$ -th largest eigenvalue of  $A$ ,  $\|\mathbf{w}\|_A$  denotes  $\sqrt{\mathbf{w}^\top A \mathbf{w}}$ ,  $|A|$  is the determinant of  $A$ ,  $\text{TR}(A)$  is the trace of  $A$ ,  $\langle A, B \rangle$  denotes  $\sum_{i,j} A_{i,j} B_{i,j}$ , and  $A \preceq B$  means that  $B - A$  is positive semidefinite. The sign function  $\text{SGN}(a)$  is 1 if  $a \geq 0$  and  $-1$  otherwise.

## 2 Setup and an Optimal Algorithm

We consider the following setting. On each round  $t = 1, 2, \dots, T$ : **(1)** the adversary first presents an example  $\mathbf{x}_t \in \mathbb{R}^d$ , **(2)** the learner chooses  $\mathbf{w}_t \in \mathbb{R}^d$  and predicts  $\mathbf{w}_t^\top \mathbf{x}_t$ , **(3)** the adversary reveals a loss function  $f_t(\mathbf{w}) = \ell_t(\mathbf{w}^\top \mathbf{x}_t)$  for some convex, differentiable  $\ell_t : \mathbb{R} \rightarrow \mathbb{R}_+$ , and **(4)** the learner suffers loss  $f_t(\mathbf{w}_t)$  for this round.

The learner’s regret to a comparator  $\mathbf{w}$  is defined as  $R_T(\mathbf{w}) = \sum_{t=1}^T f_t(\mathbf{w}_t) - \sum_{t=1}^T f_t(\mathbf{w})$ . Typical results study  $R_T(\mathbf{w})$  against all  $\mathbf{w}$  with a bounded norm in some geometry. For an invariant update,

<sup>1</sup>Recent work by [13] also studies sparse updates for a more complicated variant of Frequent Directions which is randomized and incurs extra approximation error.

<sup>2</sup>Stochastic setting assumes that the examples are drawn i.i.d. from a distribution.

we relax this requirement and only put bounds on the predictions  $\mathbf{w}^\top \mathbf{x}_t$ . Specifically, for some pre-chosen constant  $C$  we define  $\mathcal{K}_t \stackrel{\text{def}}{=} \{\mathbf{w} : |\mathbf{w}^\top \mathbf{x}_t| \leq C\}$ . We seek to minimize regret to all comparators that generate bounded predictions on every data point, that is:

$$R_T = \sup_{\mathbf{w} \in \mathcal{K}} R_T(\mathbf{w}) \quad \text{where } \mathcal{K} \stackrel{\text{def}}{=} \bigcap_{t=1}^T \mathcal{K}_t = \{\mathbf{w} : \forall t = 1, 2, \dots, T, |\mathbf{w}^\top \mathbf{x}_t| \leq C\}.$$

Under this setup, if the data are transformed to  $M\mathbf{x}_t$  for all  $t$  and some invertible matrix  $M \in \mathbb{R}^{d \times d}$ , the optimal  $\mathbf{w}^*$  simply moves to  $(M^{-1})^\top \mathbf{w}^*$ , which still has bounded predictions but might have significantly larger norm. This relaxation is similar to the comparator set considered in [33].

We make two structural assumptions on the loss functions.

**Assumption 1.** (*Scalar Lipschitz*) The loss function  $\ell_t$  satisfies  $|\ell'_t(z)| \leq L$  whenever  $|z| \leq C$ .

**Assumption 2.** (*Curvature*) There exists  $\sigma_t \geq 0$  such that for all  $\mathbf{u}, \mathbf{w} \in \mathcal{K}$ ,  $f_t(\mathbf{w})$  is lower bounded by  $f_t(\mathbf{u}) + \nabla f_t(\mathbf{u})^\top (\mathbf{w} - \mathbf{u}) + \frac{\sigma_t}{2} (\nabla f_t(\mathbf{u})^\top (\mathbf{w} - \mathbf{u}))^2$ .

Note that when  $\sigma_t = 0$ , Assumption 2 merely imposes convexity. More generally, it is satisfied by squared loss  $f_t(\mathbf{w}) = (\mathbf{w}^\top \mathbf{x}_t - y_t)^2$  with  $\sigma_t = \frac{1}{8C^2}$  whenever  $|\mathbf{w}^\top \mathbf{x}_t|$  and  $|y_t|$  are bounded by  $C$ , as well as for all exp-concave functions (see [18, Lemma 3]).

Enlarging the comparator set might result in worse regret. We next show matching upper and lower bounds qualitatively similar to the standard setting, but with an extra unavoidable  $\sqrt{d}$  factor.<sup>3</sup>

**Theorem 1.** For any online algorithm generating  $\mathbf{w}_t \in \mathbb{R}^d$  and all  $T \geq d$ , there exists a sequence of  $T$  examples  $\mathbf{x}_t \in \mathbb{R}^d$  and loss functions  $\ell_t$  satisfying Assumptions 1 and 2 (with  $\sigma_t = 0$ ) such that the regret  $R_T$  is at least  $CL\sqrt{dT}/2$ .

We now give an algorithm that matches the lower bound up to logarithmic constants in the worst case but enjoys much smaller regret when  $\sigma_t \neq 0$ . At round  $t+1$  with some invertible matrix  $A_t$  specified later and gradient  $\mathbf{g}_t = \nabla f_t(\mathbf{w}_t)$ , the algorithm performs the following update *before* making the prediction on the example  $\mathbf{x}_{t+1}$ :

$$\mathbf{u}_{t+1} = \mathbf{w}_t - A_t^{-1} \mathbf{g}_t, \quad \text{and} \quad \mathbf{w}_{t+1} = \underset{\mathbf{w} \in \mathcal{K}_{t+1}}{\text{argmin}} \|\mathbf{w} - \mathbf{u}_{t+1}\|_{A_t}. \quad (1)$$

The projection onto the set  $\mathcal{K}_{t+1}$  differs from typical norm-based projections as it only enforces boundedness on  $\mathbf{x}_{t+1}$  at round  $t+1$ . Moreover, this projection step can be performed in closed form.

**Lemma 1.** For any  $\mathbf{x} \neq \mathbf{0}$ ,  $\mathbf{u} \in \mathbb{R}^d$  and positive definite matrix  $A \in \mathbb{R}^{d \times d}$ , we have

$$\underset{\mathbf{w} : |\mathbf{w}^\top \mathbf{x}| \leq C}{\text{argmin}} \|\mathbf{w} - \mathbf{u}\|_A = \mathbf{u} - \frac{\tau_C(\mathbf{u}^\top \mathbf{x})}{\mathbf{x}^\top A^{-1} \mathbf{x}} A^{-1} \mathbf{x}, \quad \text{where } \tau_C(y) = \text{SGN}(y) \max\{|y| - C, 0\}.$$

If  $A_t$  is a diagonal matrix, updates similar to those of Ross et al. [33] are recovered. We study a choice of  $A_t$  that is similar to the Online Newton Step (ONS) [18] (though with different projections):

$$A_t = \alpha \mathbf{I}_d + \sum_{s=1}^t (\sigma_s + \eta_s) \mathbf{g}_s \mathbf{g}_s^\top \quad (2)$$

for some parameters  $\alpha > 0$  and  $\eta_t \geq 0$ . The regret guarantee of this algorithm is shown below:

**Theorem 2.** Under Assumptions 1 and 2, suppose that  $\sigma_t \geq \sigma \geq 0$  for all  $t$ , and  $\eta_t$  is non-increasing. Then using the matrices (2) in the updates (1) yields for all  $\mathbf{w} \in \mathcal{K}$ ,

$$R_T(\mathbf{w}) \leq \frac{\alpha}{2} \|\mathbf{w}\|_2^2 + 2(CL)^2 \sum_{t=1}^T \eta_t + \frac{d}{2(\sigma + \eta_T)} \ln \left( 1 + \frac{(\sigma + \eta_T) \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{d\alpha} \right).$$

<sup>3</sup>In the standard setting where  $\mathbf{w}_t$  and  $\mathbf{x}_t$  are restricted such that  $\|\mathbf{w}_t\| \leq D$  and  $\|\mathbf{x}_t\| \leq X$ , the minimax regret is  $\mathcal{O}(DXL\sqrt{T})$ . This is clearly a special case of our setting with  $C = DX$ .

---

**Algorithm 1** Sketched Online Newton (SON)

---

**Input:** Parameters  $C$ ,  $\alpha$  and  $m$ .

- 1: Initialize  $\mathbf{u}_1 = \mathbf{0}_{d \times 1}$ .
  - 2: Initialize sketch  $(S, H) \leftarrow \mathbf{SketchInit}(\alpha, m)$ .
  - 3: **for**  $t = 1$  **to**  $T$  **do**
  - 4:   Receive example  $\mathbf{x}_t$ .
  - 5:   **Projection step:** compute  $\hat{\mathbf{x}} = S\mathbf{x}_t$ ,  $\gamma = \frac{\tau_C(\mathbf{u}_t^\top \mathbf{x}_t)}{\mathbf{x}_t^\top \mathbf{x}_t - \hat{\mathbf{x}}^\top H \hat{\mathbf{x}}}$  and set  $\mathbf{w}_t = \mathbf{u}_t - \gamma(\mathbf{x}_t - S^\top H \hat{\mathbf{x}})$ .
  - 6:   Predict label  $y_t = \mathbf{w}_t^\top \mathbf{x}_t$  and suffer loss  $\ell_t(y_t)$ .
  - 7:   Compute gradient  $\mathbf{g}_t = \ell'_t(y_t)\mathbf{x}_t$  and the *to-sketch vector*  $\hat{\mathbf{g}} = \sqrt{\sigma_t + \eta_t}\mathbf{g}_t$ .
  - 8:    $(S, H) \leftarrow \mathbf{SketchUpdate}(\hat{\mathbf{g}})$ .
  - 9:   **Update weight:**  $\mathbf{u}_{t+1} = \mathbf{w}_t - \frac{1}{\alpha}(\mathbf{g}_t - S^\top H S \mathbf{g}_t)$ .
  - 10: **end for**
- 

The dependence on  $\|\mathbf{w}\|_2^2$  implies that the method is not completely invariant to transformations of the data. This is due to the part  $\alpha \mathbf{I}_d$  in  $A_t$ . However, this is not critical since  $\alpha$  is fixed and small while the other part of the bound grows to eventually become the dominating term. Moreover, we can even set  $\alpha = 0$  and replace the inverse with the Moore-Penrose pseudoinverse to obtain a truly invariant algorithm, as discussed in Appendix D. We use  $\alpha > 0$  in the remainder for simplicity.

The implication of this regret bound is the following: in the worst case where  $\sigma = 0$ , we set  $\eta_t = \sqrt{d/C^2 L^2 t}$  and the bound simplifies to

$$R_T(\mathbf{w}) \leq \frac{\alpha}{2} \|\mathbf{w}\|_2^2 + \frac{CL}{2} \sqrt{Td} \ln \left( 1 + \frac{\sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{\alpha CL \sqrt{Td}} \right) + 4CL\sqrt{Td},$$

essentially only losing a logarithmic factor compared to the lower bound in Theorem 1. On the other hand, if  $\sigma_t \geq \sigma > 0$  for all  $t$ , then we set  $\eta_t = 0$  and the regret simplifies to

$$R_T(\mathbf{w}) \leq \frac{\alpha}{2} \|\mathbf{w}\|_2^2 + \frac{d}{2\sigma} \ln \left( 1 + \frac{\sigma \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{d\alpha} \right), \quad (3)$$

extending the  $\mathcal{O}(d \ln T)$  results in [18] to the weaker Assumption 2 and a larger comparator set  $\mathcal{K}$ .

### 3 Efficiency via Sketching

Our algorithm so far requires  $\Omega(d^2)$  time and space just as ONS. In this section we show how to achieve regret guarantees nearly as good as the above bounds, while keeping computation within a constant factor of first order methods.

Let  $G_t \in \mathbb{R}^{t \times d}$  be a matrix such that the  $t$ -th row is  $\hat{\mathbf{g}}_t^\top$  where we define  $\hat{\mathbf{g}}_t = \sqrt{\sigma_t + \eta_t}\mathbf{g}_t$  to be the *to-sketch vector*. Our previous choice of  $A_t$  (Eq. (2)) can be written as  $\alpha \mathbf{I}_d + G_t^\top G_t$ . The idea of sketching is to maintain an approximation of  $G_t$ , denoted by  $S_t \in \mathbb{R}^{m \times d}$  where  $m \ll d$  is a small constant called the sketch size. If  $m$  is chosen so that  $S_t^\top S_t$  approximates  $G_t^\top G_t$  well, we can redefine  $A_t$  as  $\alpha \mathbf{I}_d + S_t^\top S_t$  for the algorithm.

To see why this admits an efficient algorithm, notice that by the Woodbury formula one has  $A_t^{-1} = \frac{1}{\alpha} (\mathbf{I}_d - S_t^\top (\alpha \mathbf{I}_m + S_t S_t^\top)^{-1} S_t)$ . With the notation  $H_t = (\alpha \mathbf{I}_m + S_t S_t^\top)^{-1} \in \mathbb{R}^{m \times m}$  and  $\gamma_t = \tau_C(\mathbf{u}_{t+1}^\top \mathbf{x}_{t+1}) / (\mathbf{x}_{t+1}^\top \mathbf{x}_{t+1} - \mathbf{x}_{t+1}^\top S_t^\top H_t S_t \mathbf{x}_{t+1})$ , update (1) becomes:

$$\mathbf{u}_{t+1} = \mathbf{w}_t - \frac{1}{\alpha} (\mathbf{g}_t - S_t^\top H_t S_t \mathbf{g}_t), \quad \text{and} \quad \mathbf{w}_{t+1} = \mathbf{u}_{t+1} - \gamma_t (\mathbf{x}_{t+1} - S_t^\top H_t S_t \mathbf{x}_{t+1}).$$

The operations involving  $S_t \mathbf{g}_t$  or  $S_t \mathbf{x}_{t+1}$  require only  $\mathcal{O}(md)$  time, while matrix vector products with  $H_t$  require only  $\mathcal{O}(m^2)$ . Altogether, these updates are at most  $m$  times more expensive than first order algorithms as long as  $S_t$  and  $H_t$  can be maintained efficiently. We call this algorithm Sketched Online Newton (SON) and summarize it in Algorithm 1.

We now discuss three sketching techniques to maintain the matrices  $S_t$  and  $H_t$  efficiently, each requiring  $\mathcal{O}(md)$  storage and time linear in  $d$ .

---

**Algorithm 2** FD-Sketch for FD-SON

---

**Internal State:**  $S$  and  $H$ .**SketchInit**( $\alpha, m$ )

- 1: Set  $S = \mathbf{0}_{m \times d}$  and  $H = \frac{1}{\alpha} \mathbf{I}_m$ .
- 2: Return  $(S, H)$ .

**SketchUpdate**( $\hat{g}$ )

- 1: Insert  $\hat{g}$  into the last row of  $S$ .
  - 2: Compute eigendecomposition:  $V^\top \Sigma V = S^\top S$  and set  $S = (\Sigma - \Sigma_{m,m} \mathbf{I}_m)^{\frac{1}{2}} V$ .
  - 3: Set  $H = \text{diag} \left\{ \frac{1}{\alpha + \Sigma_{1,1} - \Sigma_{m,m}}, \dots, \frac{1}{\alpha} \right\}$ .
  - 4: Return  $(S, H)$ .
- 

---

**Algorithm 3** Oja's Sketch for Oja-SON

---

**Internal State:**  $t, \Lambda, V$  and  $H$ .**SketchInit**( $\alpha, m$ )

- 1: Set  $t = 0, \Lambda = \mathbf{0}_{m \times m}, H = \frac{1}{\alpha} \mathbf{I}_m$  and  $V$  to any  $m \times d$  matrix with orthonormal rows.
- 2: Return  $(\mathbf{0}_{m \times d}, H)$ .

**SketchUpdate**( $\hat{g}$ )

- 1: Update  $t \leftarrow t + 1, \Lambda$  and  $V$  as Eqn. 4.
  - 2: Set  $S = (t\Lambda)^{\frac{1}{2}} V$ .
  - 3: Set  $H = \text{diag} \left\{ \frac{1}{\alpha + t\Lambda_{1,1}}, \dots, \frac{1}{\alpha + t\Lambda_{m,m}} \right\}$ .
  - 4: Return  $(S, H)$ .
- 

**Random Projection (RP).** Random projections are classical methods for sketching [19, 1, 21]. Here we consider Gaussian Random Projection sketch:  $S_t = S_{t-1} + \mathbf{r}_t \hat{\mathbf{g}}_t^\top$ , where each entry of  $\mathbf{r}_t \in \mathbb{R}^m$  is an independent random Gaussian variable drawn from  $\mathcal{N}(0, 1/\sqrt{m})$ . One can verify that the update of  $H_t^{-1}$  can be realized by two rank-one updates:  $H_t^{-1} = H_{t-1}^{-1} + \mathbf{q}_t \mathbf{r}_t^\top + \mathbf{r}_t \mathbf{q}_t^\top$  where  $\mathbf{q}_t = S_t \hat{\mathbf{g}}_t - \frac{\|\hat{\mathbf{g}}_t\|_2^2}{2} \mathbf{r}_t$ . Using Woodbury formula, this results in  $\mathcal{O}(md)$  update of  $S$  and  $H$  (see Algorithm 6 in Appendix E). We call this combination of SON with RP-sketch RP-SON. When  $\alpha = 0$  this algorithm is invariant to linear transformations for each fixed realization of the randomness.

Using the existing guarantees for RP-sketch, in Appendix E we show a similar regret bound as Theorem 2 up to constants, provided  $m = \tilde{\Omega}(r)$  where  $r$  is the rank of  $G_T$ . Therefore RP-SON is near invariant, and gives substantial computational gains when  $r \ll d$  with small regret overhead.

**Frequent Directions (FD).** When  $G_T$  is near full-rank, however, RP-SON may not perform well. To address this, we consider Frequent Directions (FD) sketch [12, 23], a deterministic sketching method. FD maintains the invariant that the last row of  $S_t$  is always  $\mathbf{0}$ . On each round, the vector  $\hat{\mathbf{g}}_t^\top$  is inserted into the last row of  $S_{t-1}$ , then the covariance of the resulting matrix is eigendecomposed into  $V_t^\top \Sigma_t V_t$  and  $S_t$  is set to  $(\Sigma_t - \rho_t \mathbf{I}_m)^{\frac{1}{2}} V_t$  where  $\rho_t$  is the smallest eigenvalue. Since the rows of  $S_t$  are orthogonal to each other,  $H_t$  is a diagonal matrix and can be maintained efficiently (see Algorithm 2). The sketch update works in  $\mathcal{O}(md)$  time (see [12] and Appendix G.2) so the total running time is  $\mathcal{O}(md)$  per round. We call this combination FD-SON and prove the following regret bound with notation  $\Omega_k = \sum_{i=k+1}^d \lambda_i(G_T^\top G_T)$  for any  $k = 0, \dots, m-1$ .

**Theorem 3.** *Under Assumptions 1 and 2, suppose that  $\sigma_t \geq \sigma \geq 0$  for all  $t$  and  $\eta_t$  is non-increasing. FD-SON ensures that for any  $\mathbf{w} \in \mathcal{K}$  and  $k = 0, \dots, m-1$ , we have*

$$R_T(\mathbf{w}) \leq \frac{\alpha}{2} \|\mathbf{w}\|_2^2 + 2(CL)^2 \sum_{t=1}^T \eta_t + \frac{m}{2(\sigma + \eta_T)} \ln \left( 1 + \frac{\text{TR}(S_T^\top S_T)}{m\alpha} \right) + \frac{m\Omega_k}{2(m-k)(\sigma + \eta_T)\alpha}.$$

Instead of the rank, the bound depends on the spectral decay  $\Omega_k$ , which essentially is the only extra term compared to the bound in Theorem 2. Similarly to previous discussion, if  $\sigma_t \geq \sigma$ , we get the bound  $\frac{\alpha}{2} \|\mathbf{w}\|_2^2 + \frac{m}{2\sigma} \ln \left( 1 + \frac{\text{TR}(S_T^\top S_T)}{m\alpha} \right) + \frac{m\Omega_k}{2(m-k)\sigma\alpha}$ . With  $\alpha$  tuned well, we pay logarithmic regret for the top  $m$  eigenvectors, but a square root regret  $\mathcal{O}(\sqrt{\Omega_k})$  for remaining directions not controlled by our sketch. This is expected for deterministic sketching which focuses on the dominant part of the spectrum. When  $\alpha$  is not tuned we still get sublinear regret as long as  $\Omega_k$  is sublinear.

**Oja's Algorithm.** Oja's algorithm [28, 29] is not usually considered as a sketching algorithm but seems very natural here. This algorithm uses online gradient descent to find eigenvectors and eigenvalues of data in a streaming fashion, with the to-sketch vector  $\hat{\mathbf{g}}_t$ 's as the input. Specifically, let  $V_t \in \mathbb{R}^{m \times d}$  denote the estimated eigenvectors and the diagonal matrix  $\Lambda_t \in \mathbb{R}^{m \times m}$  contain the estimated eigenvalues at the end of round  $t$ . Oja's algorithm updates as:

$$\Lambda_t = (\mathbf{I}_m - \Gamma_t) \Lambda_{t-1} + \Gamma_t \text{diag}\{V_{t-1} \hat{\mathbf{g}}_t\}^2, \quad V_t \stackrel{\text{orth}}{\leftarrow} V_{t-1} + \Gamma_t V_{t-1} \hat{\mathbf{g}}_t \hat{\mathbf{g}}_t^\top \quad (4)$$

where  $\Gamma_t \in \mathbb{R}^{m \times m}$  is a diagonal matrix with (possibly different) learning rates of order  $\Theta(1/t)$  on the diagonal, and the “ $\sqrt{\cdot}^{\text{orth}}$ ” operator represents an orthonormalizing step.<sup>4</sup> The sketch is then  $S_t = (t\Lambda_t)^{\frac{1}{2}} V_t$ . The rows of  $S_t$  are orthogonal and thus  $H_t$  is an efficiently maintainable diagonal matrix (see Algorithm 3). We call this combination Oja-SON.

The time complexity of Oja’s algorithm is  $\mathcal{O}(m^2 d)$  per round due to the orthonormalizing step. To improve the running time to  $\mathcal{O}(md)$ , one can only update the sketch every  $m$  rounds (similar to the block power method [16, 22]). The regret guarantee of this algorithm is unclear since existing analysis for Oja’s algorithm is only for the stochastic setting (see e.g. [2, 22]). However, Oja-SON provides good performance experimentally.

## 4 Sparse Implementation

In many applications, examples (and hence gradients) are sparse in the sense that  $\|\mathbf{x}_t\|_0 \leq s$  for all  $t$  and some small constant  $s \ll d$ . Most online first order methods enjoy a per-example running time depending on  $s$  instead of  $d$  in such settings. Achieving the same for second order methods is more difficult since  $A_t^{-1} \mathbf{g}_t$  (or sketched versions) are typically dense even if  $\mathbf{g}_t$  is sparse.

We show how to implement our algorithms in sparsity-dependent time, specifically, in  $\mathcal{O}(m^2 + ms)$  for RP-SON and FD-SON and in  $\mathcal{O}(m^3 + ms)$  for Oja-SON. We emphasize that since the sketch would still quickly become a dense matrix even if the examples are sparse, achieving purely sparsity-dependent time is highly non-trivial (especially for FD-SON and Oja-SON), and may be of independent interest. Due to space limit, below we only briefly mention how to do it for Oja-SON. Similar discussion for the other two sketches can be found in Appendix G. Note that mathematically these updates are equivalent to the non-sparse counterparts and regret guarantees are thus unchanged.

There are two ingredients to doing this for Oja-SON: (1) The eigenvectors  $V_t$  are represented as  $V_t = F_t Z_t$ , where  $Z_t \in \mathbb{R}^{m \times d}$  is a sparsely updatable direction (Step 3 in Algorithm 5) and  $F_t \in \mathbb{R}^{m \times m}$  is a matrix such that  $F_t Z_t$  is orthonormal. (2) The weights  $\mathbf{w}_t$  are split as  $\bar{\mathbf{w}}_t + Z_{t-1}^\top \mathbf{b}_t$ , where  $\mathbf{b}_t \in \mathbb{R}^m$  maintains the weights on the subspace captured by  $V_{t-1}$  (same as  $Z_{t-1}$ ), and  $\bar{\mathbf{w}}_t$  captures the weights on the complementary subspace which are again updated sparsely.

We describe the sparse updates for  $\bar{\mathbf{w}}_t$  and  $\mathbf{b}_t$  below with the details for  $F_t$  and  $Z_t$  deferred to Appendix H. Since  $S_t = (t\Lambda_t)^{\frac{1}{2}} V_t = (t\Lambda_t)^{\frac{1}{2}} F_t Z_t$  and  $\mathbf{w}_t = \bar{\mathbf{w}}_t + Z_{t-1}^\top \mathbf{b}_t$ , we know  $\mathbf{u}_{t+1}$  is

$$\mathbf{w}_t - (\mathbf{I}_d - S_t^\top H_t S_t) \frac{\mathbf{g}_t}{\alpha} = \underbrace{\bar{\mathbf{w}}_t - \frac{\mathbf{g}_t}{\alpha} - (Z_t - Z_{t-1})^\top \mathbf{b}_t}_{\stackrel{\text{def}}{=} \bar{\mathbf{u}}_{t+1}} + Z_t^\top \underbrace{\left( \mathbf{b}_t + \frac{1}{\alpha} F_t^\top (t\Lambda_t H_t) F_t Z_t \mathbf{g}_t \right)}_{\stackrel{\text{def}}{=} \mathbf{b}'_{t+1}}. \quad (5)$$

Since  $Z_t - Z_{t-1}$  is sparse by construction and the matrix operations defining  $\mathbf{b}'_{t+1}$  scale with  $m$ , overall the update can be done in  $\mathcal{O}(m^2 + ms)$ . Using the update for  $\mathbf{w}_{t+1}$  in terms of  $\mathbf{u}_{t+1}$ ,  $\mathbf{w}_{t+1}$  is equal to

$$\mathbf{u}_{t+1} - \gamma_t (\mathbf{I}_d - S_t^\top H_t S_t) \mathbf{x}_{t+1} = \underbrace{\bar{\mathbf{u}}_{t+1} - \gamma_t \mathbf{x}_{t+1}}_{\stackrel{\text{def}}{=} \bar{\mathbf{w}}_{t+1}} + Z_t^\top \underbrace{\left( \mathbf{b}'_{t+1} + \gamma_t F_t^\top (t\Lambda_t H_t) F_t Z_t \mathbf{x}_{t+1} \right)}_{\stackrel{\text{def}}{=} \mathbf{b}_{t+1}}. \quad (6)$$

Again, it is clear that all the computations scale with  $s$  and not  $d$ , so both  $\bar{\mathbf{w}}_{t+1}$  and  $\mathbf{b}_{t+1}$  require only  $\mathcal{O}(m^2 + ms)$  time to maintain. Furthermore, the prediction  $\mathbf{w}_t^\top \mathbf{x}_t = \bar{\mathbf{w}}_t^\top \mathbf{x}_t + \mathbf{b}_t^\top Z_{t-1}^\top \mathbf{x}_t$  can also be computed in  $\mathcal{O}(ms)$  time. The  $\mathcal{O}(m^3)$  in the overall complexity comes from a Gram-Schmidt step in maintaining  $F_t$  (details in Appendix H).

The pseudocode is presented in Algorithms 4 and 5 with some details deferred to Appendix H. This is the first sparse implementation of online eigenvector computation to the best of our knowledge.

## 5 Experiments

Preliminary experiments revealed that out of our three sketching options, Oja’s sketch generally has better performance (see Appendix I). For more thorough evaluation, we implemented the sparse

<sup>4</sup>For simplicity, we assume that  $V_{t-1} + \Gamma_t V_{t-1} \hat{\mathbf{g}}_t \hat{\mathbf{g}}_t^\top$  is always of full rank so that the orthonormalizing step does not reduce the dimension of  $V_t$ .

---

**Algorithm 4** Sparse Sketched Online Newton with Oja’s Algorithm

---

**Input:** Parameters  $C$ ,  $\alpha$  and  $m$ .

- 1: Initialize  $\bar{\mathbf{u}} = \mathbf{0}_{d \times 1}$  and  $\mathbf{b} = \mathbf{0}_{m \times 1}$ .
  - 2:  $(\Lambda, F, Z, H) \leftarrow \mathbf{SketchInit}(\alpha, m)$  (Algorithm 5).
  - 3: **for**  $t = 1$  **to**  $T$  **do**
  - 4:   Receive example  $\mathbf{x}_t$ .
  - 5:   **Projection step:** compute  $\hat{\mathbf{x}} = FZ\mathbf{x}_t$  and  $\gamma = \frac{\tau_C(\bar{\mathbf{u}}^\top \mathbf{x}_t + \mathbf{b}^\top Z\mathbf{x}_t)}{\mathbf{x}_t^\top \mathbf{x}_t - (t-1)\hat{\mathbf{x}}^\top \Lambda H \hat{\mathbf{x}}}$ .  
    Obtain  $\bar{\mathbf{w}} = \bar{\mathbf{u}} - \gamma \mathbf{x}_t$  and  $\mathbf{b} \leftarrow \mathbf{b} + \gamma(t-1)F^\top \Lambda H \hat{\mathbf{x}}$  (Equation 6).
  - 6:   Predict label  $y_t = \bar{\mathbf{w}}^\top \mathbf{x}_t + \mathbf{b}^\top Z\mathbf{x}_t$  and suffer loss  $\ell_t(y_t)$ .
  - 7:   Compute gradient  $\mathbf{g}_t = \ell'_t(y_t)\mathbf{x}_t$  and the *to-sket*ch vector  $\hat{\mathbf{g}} = \sqrt{\sigma_t + \eta_t}\mathbf{g}_t$ .
  - 8:    $(\Lambda, F, Z, H, \delta) \leftarrow \mathbf{SketchUpdate}(\hat{\mathbf{g}})$  (Algorithm 5).
  - 9:   **Update weight:**  $\bar{\mathbf{u}} = \bar{\mathbf{w}} - \frac{1}{\alpha}\mathbf{g}_t - (\delta^\top \mathbf{b})\hat{\mathbf{g}}$  and  $\mathbf{b} \leftarrow \mathbf{b} + \frac{1}{\alpha}tF^\top \Lambda H F Z \mathbf{g}_t$  (Equation 5).
  - 10: **end for**
- 

---

**Algorithm 5** Sparse Oja’s Sketch

---

**Internal State:**  $t, \Lambda, F, Z, H$  and  $K$ .

**SketchInit**( $\alpha, m$ )

- 1: Set  $t = 0, \Lambda = \mathbf{0}_{m \times m}, F = K = \alpha H = \mathbf{I}_m$  and  $Z$  to any  $m \times d$  matrix with orthonormal rows.
- 2: Return  $(\Lambda, F, Z, H)$ .

**SketchUpdate**( $\hat{\mathbf{g}}$ )

- 1: Update  $t \leftarrow t+1$ . Pick a diagonal stepsize matrix  $\Gamma_t$  to update  $\Lambda \leftarrow (\mathbf{I} - \Gamma_t)\Lambda + \Gamma_t \text{diag}\{FZ\hat{\mathbf{g}}\}^2$ .
  - 2: Set  $\delta = A^{-1}\Gamma_t FZ\hat{\mathbf{g}}$  and update  $K \leftarrow K + \delta\hat{\mathbf{g}}^\top Z^\top + Z\hat{\mathbf{g}}\delta^\top + (\hat{\mathbf{g}}^\top \hat{\mathbf{g}})\delta\delta^\top$ .
  - 3: Update  $Z \leftarrow Z + \delta\hat{\mathbf{g}}^\top$ .
  - 4:  $(L, Q) \leftarrow \text{Decompose}(F, K)$  (Algorithm 13), so that  $LQZ = FZ$  and  $QZ$  is orthogonal. Set  $F = Q$ .
  - 5: Set  $H \leftarrow \text{diag}\left\{\frac{1}{\alpha+t\Lambda_{1,1}}, \dots, \frac{1}{\alpha+t\Lambda_{m,m}}\right\}$ .
  - 6: Return  $(\Lambda, F, Z, H, \delta)$ .
- 

version of Oja-SON in Vowpal Wabbit.<sup>5</sup> We compare it with ADAGRAD [6, 25] on both synthetic and real-world datasets. Each algorithm takes a stepsize parameter:  $\frac{1}{\alpha}$  serves as a stepsize for Oja-SON and a scaling constant on the gradient matrix for ADAGRAD. We try both methods with the parameter set to  $2^j$  for  $j = -3, -2, \dots, 6$  and report the best results. We keep the stepsize matrix in Oja-SON fixed as  $\Gamma_t = \frac{1}{t}\mathbf{I}_m$  throughout. All methods make one online pass over data minimizing square loss.

## 5.1 Synthetic Datasets

To investigate Oja-SON’s performance in the setting it is really designed for, we generated a range of synthetic ill-conditioned datasets as follows. We picked a random Gaussian matrix  $Z \sim \mathbb{R}^{T \times d}$  ( $T = 10,000$  and  $d = 100$ ) and a random orthonormal basis  $V \in \mathbb{R}^{d \times d}$ . We chose a specific spectrum  $\lambda \in \mathbb{R}^d$  where the first  $d - 10$  coordinates are 1 and the rest increase linearly to some fixed *condition number parameter*  $\kappa$ . We let  $X = Z\text{diag}\{\lambda\}^{\frac{1}{2}}V^\top$  be our example matrix, and created a binary classification problem with labels  $y = \text{sign}(\theta^\top \mathbf{x})$ , where  $\theta \in \mathbb{R}^d$  is a random vector. We generated 20 such datasets with the same  $Z, V$  and labels  $y$  but different values of  $\kappa \in \{10, 20, \dots, 200\}$ . Note that if the algorithm is truly invariant, it would have the same behavior on these 20 datasets.

Fig. 1 (in Section 1) shows the final progressive error (i.e. fraction of misclassified examples after one pass over data) for ADAGRAD and Oja-SON (with sketch size  $m = 0, 5, 10$ ) as the condition number increases. As expected, the plot confirms the performance of first order methods such as ADAGRAD degrades when the data is ill-conditioned. The plot also shows that as the sketch size increases, Oja-SON becomes more accurate: when  $m = 0$  (no sketch at all), Oja-SON is vanilla gradient descent and is worse than ADAGRAD as expected; when  $m = 5$ , the accuracy greatly improves; and finally when  $m = 10$ , the accuracy of Oja-SON is substantially better and hardly worsens with  $\kappa$ .

---

<sup>5</sup>An open source machine learning toolkit available at <http://hunch.net/~vw>




Figure 2: Oja’s algorithm’s eigenvalue recovery error.




Figure 3: (a) Comparison of two sketch sizes on real data, and (b) Comparison against ADAGRAD on real data.

To further explain the effectiveness of Oja’s algorithm in identifying top eigenvalues and eigenvectors, the plot in Fig. 2 shows the largest relative difference between the true and estimated top 10 eigenvalues as Oja’s algorithm sees more data. This gap drops quickly after seeing just 500 examples.

## 5.2 Real-world Datasets

Next we evaluated Oja-SON on 23 benchmark datasets from the UCI and LIBSVM repository (see Appendix I for description of these datasets). Note that some datasets are very high dimensional but very sparse (e.g. for *20news*,  $d \approx 102,000$  and  $s \approx 94$ ), and consequently methods with running time quadratic (such as ONS) or even linear in dimension rather than sparsity are prohibitive.

In Fig. 3(a), we show the effect of using sketched second order information, by comparing sketch size  $m = 0$  and  $m = 10$  for Oja-SON (concrete error rates in Appendix I). We observe significant improvements in 5 datasets (*acoustic*, *census*, *heart*, *ionosphere*, *letter*), demonstrating the advantage of using second order information. However, we found that Oja-SON was outperformed by ADAGRAD on most datasets, mostly because the diagonal adaptation of ADAGRAD greatly reduces the condition number on these datasets. Moreover, one disadvantage of SON is that for the directions not in the sketch, it is essentially doing vanilla gradient descent. We expect better results using diagonal adaptation as in ADAGRAD in off-sketch directions.

To incorporate this high level idea, we performed a simple modification to Oja-SON: upon seeing example  $x_t$ , we feed  $D_t^{-\frac{1}{2}} x_t$  to our algorithm instead of  $x_t$ , where  $D_t \in \mathbb{R}^{d \times d}$  is the diagonal part of the matrix  $\sum_{\tau=1}^{t-1} g_\tau g_\tau^\top$ .<sup>6</sup> The intuition is that this diagonal rescaling first homogenizes the scales of all dimensions. Any remaining ill-conditioning is further addressed by the sketching to some degree, while the complementary subspace is no worse-off than with ADAGRAD. We believe this flexibility in picking the right vectors to sketch is an attractive aspect of our sketching-based approach.

With this modification, Oja-SON outperforms ADAGRAD on most of the datasets even for  $m = 0$ , as shown in Fig. 3(b) (concrete error rates in Appendix I). The improvement on ADAGRAD at  $m = 0$  is surprising but not impossible as the updates are not identical—our update is scale invariant like Ross et al. [33]. However, the diagonal adaptation already greatly reduces the condition number on all datasets except *splice* (see Fig. 4 in Appendix I for detailed results on this dataset), so little improvement is seen for sketch size  $m = 10$  over  $m = 0$ . For several datasets, we verified the accuracy of Oja’s method in computing the top-few eigenvalues (Appendix I), so the lack of difference between sketch sizes is due to the lack of second order information after the diagonal correction.

The average running time of our algorithm when  $m = 10$  is about 11 times slower than ADAGRAD, matching expectations. Overall, SON can significantly outperform baselines on ill-conditioned data, while maintaining a practical computational complexity.

**Acknowledgements** This work was done when Haipeng Luo and Nicolò Cesa-Bianchi were at Microsoft Research, New York.

<sup>6</sup> $D_1$  is defined as  $0.1 \times I_d$  to avoid division by zero.



## References

- [1] D. Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, 2003.
- [2] A. Balsubramani, S. Dasgupta, and Y. Freund. The fast convergence of incremental pca. In *NIPS*, 2013.
- [3] R. H. Byrd, S. Hansen, J. Nocedal, and Y. Singer. A stochastic quasi-newton method for large-scale optimization. *SIAM Journal on Optimization*, 26:1008–1031, 2016.
- [4] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [5] N. Cesa-Bianchi, A. Conconi, and C. Gentile. A second-order perceptron algorithm. *SIAM Journal on Computing*, 34(3):640–668, 2005.
- [6] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 12:2121–2159, 2011.
- [7] M. A. Erdogdu and A. Montanari. Convergence rates of sub-sampled newton methods. In *NIPS*, 2015.
- [8] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [9] W. Gao, R. Jin, S. Zhu, and Z.-H. Zhou. One-pass auc optimization. In *ICML*, 2013.
- [10] D. Garber and E. Hazan. A linearly convergent conditional gradient algorithm with applications to online and stochastic optimization. *SIAM Journal on Optimization*, 26:1493–1528, 2016.
- [11] D. Garber, E. Hazan, and T. Ma. Online learning of eigenvectors. In *ICML*, 2015.
- [12] M. Ghashami, E. Liberty, J. M. Phillips, and D. P. Woodruff. Frequent directions: Simple and deterministic matrix sketching. *SIAM Journal on Computing*, 45:1762–1792, 2015.
- [13] M. Ghashami, E. Liberty, and J. M. Phillips. Efficient frequent directions algorithm for sparse matrices. In *KDD*, 2016.
- [14] A. Gonen and S. Shalev-Shwartz. Faster sgd using sketched conditioning. *arXiv:1506.02649*, 2015.
- [15] A. Gonen, F. Orabona, and S. Shalev-Shwartz. Solving ridge regression using sketched preconditioned svrg. In *ICML*, 2016.
- [16] M. Hardt and E. Price. The noisy power method: A meta algorithm with applications. In *NIPS*, 2014.
- [17] E. Hazan and S. Kale. Projection-free online learning. In *ICML*, 2012.
- [18] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [19] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *STOC*, 1998.
- [20] M. Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML*, 2013.
- [21] D. M. Kane and J. Nelson. Sparser johnson-lindenstrauss transforms. *Journal of the ACM*, 61(1):4, 2014.
- [22] C.-L. Li, H.-T. Lin, and C.-J. Lu. Rivalry of two families of algorithms for memory-restricted streaming pca. *arXiv:1506.01490*, 2015.
- [23] E. Liberty. Simple and deterministic matrix sketching. In *KDD*, 2013.
- [24] D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [25] H. B. McMahan and M. Streeter. Adaptive bound optimization for online convex optimization. In *COLT*, 2010.
- [26] A. Mokhtari and A. Ribeiro. Global convergence of online limited memory bfgs. *JMLR*, 16:3151–3181, 2015.
- [27] P. Moritz, R. Nishihara, and M. I. Jordan. A linearly-convergent stochastic l-bfgs algorithm. In *AISTATS*, 2016.
- [28] E. Oja. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.
- [29] E. Oja and J. Karhunen. On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix. *Journal of mathematical analysis and applications*, 106(1):69–84, 1985.
- [30] F. Orabona and D. Pál. Scale-free algorithms for online linear optimization. In *ALT*, 2015.
- [31] F. Orabona, K. Crammer, and N. Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435, 2015.
- [32] M. Pilanci and M. J. Wainwright. Newton sketch: A linear-time optimization algorithm with linear-quadratic convergence. *arXiv:1505.02250*, 2015.
- [33] S. Ross, P. Mineiro, and J. Langford. Normalized online learning. In *UAI*, 2013.
- [34] N. N. Schraudolph, J. Yu, and S. Günter. A stochastic quasi-newton method for online convex optimization. In *AISTATS*, 2007.
- [35] J. Sohl-Dickstein, B. Poole, and S. Ganguli. Fast large-scale optimization by unifying stochastic gradient and quasi-newton methods. In *ICML*, 2014.
- [36] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Machine Learning*, 10(1-2):1–157, 2014.

# Supplementary material for “Efficient Second Order Online Learning by Sketching”

## A Proof of Theorem 1

*Proof.* Assuming  $T$  is a multiple of  $d$  without loss of generality, we pick  $\mathbf{x}_t$  from the basis vectors  $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$  so that each  $\mathbf{e}_i$  appears  $T/d$  times (in an arbitrary order). Note that now  $\mathcal{K}$  is just a hypercube:

$$\mathcal{K} = \{\mathbf{w} : |\mathbf{w}^\top \mathbf{x}_t| \leq C, \forall t\} = \{\mathbf{w} : \|\mathbf{w}\|_\infty \leq C\}.$$

Let  $\xi_1, \dots, \xi_T$  be independent Rademacher random variables such that  $\Pr(\xi_t = +1) = \Pr(\xi_t = -1) = \frac{1}{2}$ . For a scalar  $\theta$ , we define loss function<sup>7</sup>  $\ell_t(\theta) = (\xi_t L)\theta$ , so that Assumptions 1 and 2 are clearly satisfied with  $\sigma_t = 0$ . We show that, for any online algorithm,

$$\mathbb{E}[R_T] = \mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{w}_t^\top \mathbf{x}_t) - \inf_{\mathbf{w} \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{w}^\top \mathbf{x}_t) \right] \geq CL \sqrt{\frac{dT}{2}}$$

which implies the statement of the theorem.

First of all, note that  $\mathbb{E}[\ell_t(\mathbf{w}_t^\top \mathbf{x}_t) \mid \xi_1, \dots, \xi_{t-1}] = 0$  for any  $\mathbf{w}_t$ . Hence we have

$$\mathbb{E} \left[ \sum_{t=1}^T \ell_t(\mathbf{w}_t^\top \mathbf{x}_t) - \inf_{\mathbf{w} \in \mathcal{K}} \sum_{t=1}^T \ell_t(\mathbf{w}^\top \mathbf{x}_t) \right] = \mathbb{E} \left[ \sup_{\mathbf{w} \in \mathcal{K}} \sum_{t=1}^T -\ell_t(\mathbf{w}^\top \mathbf{x}_t) \right] = L \mathbb{E} \left[ \sup_{\mathbf{w} \in \mathcal{K}} \mathbf{w}^\top \sum_{t=1}^T \xi_t \mathbf{x}_t \right],$$

which, by the construction of  $\mathbf{x}_t$ , is

$$CL \mathbb{E} \left[ \left\| \sum_{t=1}^T \xi_t \mathbf{x}_t \right\|_1 \right] = CLd \mathbb{E} \left[ \left\| \sum_{t=1}^{T/d} \xi_t \right\| \right] \geq CLd \sqrt{\frac{T}{2d}} = CL \sqrt{\frac{dT}{2}},$$

where the final bound is due to the Khintchine inequality (see e.g. Lemma 8.2 in [4]). This concludes the proof.  $\square$

## B Projection

We prove a more general version of Lemma 1 which does not require invertibility of the matrix  $A$  here.

**Lemma 2.** For any  $\mathbf{x} \neq \mathbf{0}$ ,  $\mathbf{u} \in \mathbb{R}^{d \times 1}$  and positive semidefinite matrix  $A \in \mathbb{R}^{d \times d}$ , we have

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} : |\mathbf{w}^\top \mathbf{x}| \leq C} \|\mathbf{w} - \mathbf{u}\|_A = \begin{cases} \mathbf{u} - \frac{\tau_C(\mathbf{u}^\top \mathbf{x})}{\mathbf{x}^\top A^\dagger \mathbf{x}} A^\dagger \mathbf{x} & \text{if } \mathbf{x} \in \operatorname{range}(A) \\ \mathbf{u} - \frac{\tau_C(\mathbf{u}^\top \mathbf{x})}{\mathbf{x}^\top (\mathbf{I} - A^\dagger A) \mathbf{x}} (\mathbf{I} - A^\dagger A) \mathbf{x} & \text{if } \mathbf{x} \notin \operatorname{range}(A) \end{cases}$$

where  $\tau_C(y) = \operatorname{SGN}(y) \max\{|y| - C, 0\}$  and  $A^\dagger$  is the Moore-Penrose pseudoinverse of  $A$ . (Note that when  $A$  is rank deficient, this is one of the many possible solutions.)

*Proof.* First consider the case when  $\mathbf{x} \in \operatorname{range}(A)$ . If  $|\mathbf{u}^\top \mathbf{x}| \leq C$ , then it is trivial that  $\mathbf{w}^* = \mathbf{u}$ . We thus assume  $\mathbf{u}^\top \mathbf{x} \geq C$  below (the last case  $\mathbf{u}^\top \mathbf{x} \leq -C$  is similar). The Lagrangian of the problem is

$$L(\mathbf{w}, \lambda_1, \lambda_2) = \frac{1}{2}(\mathbf{w} - \mathbf{u})^\top A(\mathbf{w} - \mathbf{u}) + \lambda_1(\mathbf{w}^\top \mathbf{x} - C) + \lambda_2(\mathbf{w}^\top \mathbf{x} + C)$$

where  $\lambda_1 \geq 0$  and  $\lambda_2 \leq 0$  are Lagrangian multipliers. Since  $\mathbf{w}^\top \mathbf{x}$  cannot be  $C$  and  $-C$  at the same time, The complementary slackness condition implies that either  $\lambda_1 = 0$  or  $\lambda_2 = 0$ . Suppose the latter case is true, then setting the derivative with respect to  $\mathbf{w}$  to 0, we get  $\mathbf{w}^* = \mathbf{u} - \lambda_1 A^\dagger \mathbf{x} + (\mathbf{I} - A^\dagger A) \mathbf{z}$

<sup>7</sup>By adding a suitable constant, these losses can always be made nonnegative while leaving the regret unchanged.

where  $z \in R^{d \times 1}$  can be arbitrary. However, since  $A(\mathbf{I} - A^\dagger A) = 0$ , this part does not affect the objective value at all and we can simply pick  $z = 0$  so that  $w^*$  has a consistent form regardless of whether  $A$  is full rank or not. Now plugging  $w^*$  back, we have

$$L(w^*, \lambda_1, 0) = -\frac{\lambda_1^2}{2} \mathbf{x}^\top A^\dagger \mathbf{x} + \lambda_1 (\mathbf{u}^\top \mathbf{x} - C)$$

which is maximized when  $\lambda_1 = \frac{\mathbf{u}^\top \mathbf{x} - C}{\mathbf{x}^\top A^\dagger \mathbf{x}} \geq 0$ . Plugging this optimal  $\lambda_1$  into  $w^*$  gives the stated solution. On the other hand, if  $\lambda_1 = 0$  instead, we can proceed similarly and verify that it gives a smaller dual value (0 in fact), proving the previous solution is indeed optimal.

We now move on to the case when  $\mathbf{x} \notin \text{range}(A)$ . First of all the stated solution is well defined since  $\mathbf{x}^\top (\mathbf{I} - A^\dagger A) \mathbf{x}$  is nonzero in this case. Moreover, direct calculation shows that  $w^*$  is in the valid space:  $|\mathbf{w}^{*\top} \mathbf{x}| = |\mathbf{u}^\top \mathbf{x} - \tau_C(\mathbf{u}^\top \mathbf{x})| \leq C$ , and also it gives the minimal possible distance value  $\|\mathbf{w}^* - \mathbf{u}\|_A = 0$ , proving the lemma.  $\square$

## C Proof of Theorem 2

We first prove a general regret bound that holds for any choice of  $A_t$  in update 1:

$$\begin{aligned} \mathbf{u}_{t+1} &= \mathbf{w}_t - A_t^{-1} \mathbf{g}_t \\ \mathbf{w}_{t+1} &= \underset{\mathbf{w} \in \mathcal{K}_{t+1}}{\text{argmin}} \|\mathbf{w} - \mathbf{u}_{t+1}\|_{A_t} . \end{aligned}$$

This bound will also be useful in proving regret guarantees for the sketched versions.

**Proposition 1.** *For any sequence of positive definite matrices  $A_t$  and sequence of losses satisfying Assumptions 1 and 2, the regret of updates (1) against any comparator  $\mathbf{w} \in \mathcal{K}$  satisfies*

$$2R_T(\mathbf{w}) \leq \underbrace{\|\mathbf{w}\|_{A_0}^2 + \sum_{t=1}^T \mathbf{g}_t^\top A_t^{-1} \mathbf{g}_t}_{\text{"Gradient Bound"} R_G} + \underbrace{\sum_{t=1}^T (\mathbf{w}_t - \mathbf{w})^\top (A_t - A_{t-1} - \sigma_t \mathbf{g}_t \mathbf{g}_t^\top) (\mathbf{w}_t - \mathbf{w})}_{\text{"Diameter Bound"} R_D}$$

*Proof.* Since  $\mathbf{w}_{t+1}$  is the projection of  $\mathbf{u}_{t+1}$  onto  $\mathcal{K}_{t+1}$ , by the property of projections (see for example [17, Lemma 8]), the algorithm ensures

$$\|\mathbf{w}_{t+1} - \mathbf{w}\|_{A_t}^2 \leq \|\mathbf{u}_{t+1} - \mathbf{w}\|_{A_t}^2 = \|\mathbf{w}_t - \mathbf{w}\|_{A_t}^2 + \mathbf{g}_t^\top A_t^{-1} \mathbf{g}_t - 2\mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w})$$

for all  $\mathbf{w} \in \mathcal{K} \subseteq \mathcal{K}_{t+1}$ . By the curvature property in Assumption 2, we then have that

$$\begin{aligned} 2R_T(\mathbf{w}) &\leq \sum_{t=1}^T 2\mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w}) - \sigma_t (\mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w}))^2 \\ &\leq \sum_{t=1}^T \mathbf{g}_t^\top A_t^{-1} \mathbf{g}_t + \|\mathbf{w}_t - \mathbf{w}\|_{A_t}^2 - \|\mathbf{w}_{t+1} - \mathbf{w}\|_{A_t}^2 - \sigma_t (\mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w}))^2 \\ &\leq \|\mathbf{w}\|_{A_0}^2 + \sum_{t=1}^T \mathbf{g}_t^\top A_t^{-1} \mathbf{g}_t + (\mathbf{w}_t - \mathbf{w})^\top (A_t - A_{t-1} - \sigma_t \mathbf{g}_t \mathbf{g}_t^\top) (\mathbf{w}_t - \mathbf{w}), \end{aligned}$$

which completes the proof.  $\square$

*Proof of Theorem 2.* We apply Proposition 1 with the choice:  $A_0 = \alpha \mathbf{I}_d$  and  $A_t = A_{t-1} + (\sigma_t + \eta_t) \mathbf{g}_t \mathbf{g}_t^\top$ , which gives  $\|\mathbf{w}\|_{A_0}^2 = \alpha \|\mathbf{w}\|_2^2$  and

$$R_D = \sum_{t=1}^T \eta_t (\mathbf{w}_t - \mathbf{w})^\top \mathbf{g}_t \mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w}) \leq 4(CL)^2 \sum_{t=1}^T \eta_t ,$$

where the last equality uses the Lipschitz property in Assumption 1 and the boundedness of  $\mathbf{w}_t^\top \mathbf{x}_t$  and  $\mathbf{w}^\top \mathbf{x}_t$ .

For the term  $R_G$ , define  $\hat{A}_t = \frac{\alpha}{\sigma + \eta_T} \mathbf{I}_d + \sum_{s=1}^t \mathbf{g}_s \mathbf{g}_s^\top$ . Since  $\sigma_t \geq \sigma$  and  $\eta_t$  is non-increasing, we have  $\hat{A}_t \preceq \frac{1}{\sigma + \eta_T} A_t$ , and therefore:

$$\begin{aligned}
R_G &\leq \frac{1}{\sigma + \eta_T} \sum_{t=1}^T \mathbf{g}_t^\top \hat{A}_t^{-1} \mathbf{g}_t = \frac{1}{\sigma + \eta_T} \sum_{t=1}^T \langle \hat{A}_t - \hat{A}_{t-1}, \hat{A}_t^{-1} \rangle \\
&\leq \frac{1}{\sigma + \eta_T} \sum_{t=1}^T \ln \frac{|\hat{A}_t|}{|\hat{A}_{t-1}|} = \frac{1}{\sigma + \eta_T} \ln \frac{|\hat{A}_T|}{|\hat{A}_0|} \\
&= \frac{1}{\sigma + \eta_T} \sum_{i=1}^d \ln \left( 1 + \frac{(\sigma + \eta_T) \lambda_i \left( \sum_{t=1}^T \mathbf{g}_t \mathbf{g}_t^\top \right)}{\alpha} \right) \\
&\leq \frac{d}{\sigma + \eta_T} \ln \left( 1 + \frac{(\sigma + \eta_T) \sum_{i=1}^d \lambda_i \left( \sum_{t=1}^T \mathbf{g}_t \mathbf{g}_t^\top \right)}{d\alpha} \right) \\
&= \frac{d}{\sigma + \eta_T} \ln \left( 1 + \frac{(\sigma + \eta_T) \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{d\alpha} \right)
\end{aligned}$$

where the second inequality is by the concavity of the function  $\ln |X|$  (see [18, Lemma 12] for an alternative proof), and the last one is by Jensen's inequality. This concludes the proof.  $\square$

## D A Truly Invariant Algorithm

In this section we discuss how to make our adaptive online Newton algorithm truly invariant to invertible linear transformations. To achieve this, we set  $\alpha = 0$  and replace  $A_t^{-1}$  with the Moore-Penrose pseudoinverse  $A_t^\dagger$ :<sup>8</sup>

$$\begin{aligned}
\mathbf{u}_{t+1} &= \mathbf{w}_t - A_t^\dagger \mathbf{g}_t, \\
\mathbf{w}_{t+1} &= \operatorname{argmin}_{\mathbf{w} \in \mathcal{K}_{t+1}} \|\mathbf{w} - \mathbf{u}_{t+1}\|_{A_t}.
\end{aligned} \tag{7}$$

When written in this form, it is not immediately clear that the algorithm has the invariant property. However, one can rewrite the algorithm in a mirror descent form:

$$\begin{aligned}
\mathbf{w}_{t+1} &= \operatorname{argmin}_{\mathbf{w} \in \mathcal{K}_{t+1}} \left\| \mathbf{w} - \mathbf{w}_t + A_t^\dagger \mathbf{g}_t \right\|_{A_t}^2 \\
&= \operatorname{argmin}_{\mathbf{w} \in \mathcal{K}_{t+1}} \|\mathbf{w} - \mathbf{w}_t\|_{A_t}^2 + 2(\mathbf{w} - \mathbf{w}_t)^\top A_t A_t^\dagger \mathbf{g}_t \\
&= \operatorname{argmin}_{\mathbf{w} \in \mathcal{K}_{t+1}} \|\mathbf{w} - \mathbf{w}_t\|_{A_t}^2 + 2\mathbf{w}^\top \mathbf{g}_t
\end{aligned}$$

where we use the fact that  $\mathbf{g}_t$  is in the range of  $A_t$  in the last step. Now suppose all the data  $\mathbf{x}_t$  are transformed to  $M\mathbf{x}_t$  for some unknown and invertible matrix  $M$ , then one can verify that all the weights will be transformed to  $M^{-T}\mathbf{w}_t$  accordingly, ensuring the prediction to remain the same.

Moreover, the regret bound of this algorithm can be bounded as below. First notice that even when  $A_t$  is rank deficient, the projection step still ensures the following:  $\|\mathbf{w}_{t+1} - \mathbf{w}\|_{A_t}^2 \leq \|\mathbf{u}_{t+1} - \mathbf{w}\|_{A_t}^2$ , which is proven in [18, Lemma 8]. Therefore, the entire proof of Theorem 2 still holds after replacing  $A_t^{-1}$  with  $A_t^\dagger$ , giving the regret bound:

$$\frac{1}{2} \sum_{t=1}^T \mathbf{g}_t^\top A_t^\dagger \mathbf{g}_t + 2(CL)^2 \eta_t. \tag{8}$$

The key now is to bound the term  $\sum_{t=1}^T \mathbf{g}_t^\top A_t^\dagger \mathbf{g}_t$  where we define  $\hat{A}_t = \sum_{s=1}^t \mathbf{g}_s \mathbf{g}_s^\top$ . In order to do this, we proceed similarly to the proof of [5, Theorem 4.2] to show that this term is of order  $\mathcal{O}(d^2 \ln T)$  in the worst case.

<sup>8</sup>See Appendix B for the closed form of the projection step.

**Theorem 4.** Let  $\lambda^*$  be the minimum among the smallest nonzero eigenvalues of  $\widehat{A}_t$  ( $t = 1, \dots, T$ ) and  $r$  be the rank of  $\widehat{A}_T$ . We have

$$\sum_{t=1}^T \mathbf{g}_t^\top \widehat{A}_t^\dagger \mathbf{g}_t \leq r + \frac{(1+r)r}{2} \ln \left( 1 + \frac{2 \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{(1+r)r\lambda^*} \right).$$

*Proof.* First by Cesa-Bianchi et al. [5, Lemma D.1], we have

$$\mathbf{g}_t^\top \widehat{A}_t^\dagger \mathbf{g}_t = \begin{cases} 1 & \text{if } \mathbf{g}_t \notin \text{range}(\widehat{A}_{t-1}) \\ 1 - \frac{\det_+(\widehat{A}_{t-1})}{\det_+(\widehat{A}_t)} < 1 & \text{if } \mathbf{g}_t \in \text{range}(\widehat{A}_{t-1}) \end{cases}$$

where  $\det_+(M)$  denotes the product of the nonzero eigenvalues of matrix  $M$ . We thus separate the steps  $t$  such that  $\mathbf{g}_t \in \text{range}(\widehat{A}_{t-1})$  from those where  $\mathbf{g}_t \notin \text{range}(\widehat{A}_{t-1})$ . For each  $k = 1, \dots, r$  let  $T_k$  be the first time step  $t$  in which the rank of  $A_t$  is  $k$  (so that  $T_1 = 1$ ). Also let  $T_{r+1} = T + 1$  for convenience. With this notation, we have

$$\begin{aligned} \sum_{t=1}^T \mathbf{g}_t^\top \widehat{A}_t^\dagger \mathbf{g}_t &= \sum_{k=1}^r \left( \mathbf{g}_{T_k}^\top \widehat{A}_{T_k}^\dagger \mathbf{g}_{T_k} + \sum_{t=T_k+1}^{T_{k+1}-1} \mathbf{g}_t^\top \widehat{A}_t^\dagger \mathbf{g}_t \right) \\ &= \sum_{k=1}^r \left( 1 + \sum_{t=T_k+1}^{T_{k+1}-1} \left( 1 - \frac{\det_+(\widehat{A}_{t-1})}{\det_+(\widehat{A}_t)} \right) \right) \\ &= r + \sum_{k=1}^r \sum_{t=T_k+1}^{T_{k+1}-1} \left( 1 - \frac{\det_+(\widehat{A}_{t-1})}{\det_+(\widehat{A}_t)} \right) \\ &\leq r + \sum_{k=1}^r \sum_{t=T_k+1}^{T_{k+1}-1} \ln \frac{\det_+(\widehat{A}_t)}{\det_+(\widehat{A}_{t-1})} \\ &= r + \sum_{k=1}^r \ln \frac{\det_+(\widehat{A}_{T_{k+1}-1})}{\det_+(\widehat{A}_{T_k})}. \end{aligned}$$

Fix any  $k$  and let  $\lambda_{k,1}, \dots, \lambda_{k,k}$  be the nonzero eigenvalues of  $\widehat{A}_{T_k}$  and  $\lambda_{k,1} + \mu_{k,1}, \dots, \lambda_{k,k} + \mu_{k,k}$  be the nonzero eigenvalues of  $\widehat{A}_{T_{k+1}-1}$ . Then

$$\ln \frac{\det_+(\widehat{A}_{T_{k+1}-1})}{\det_+(\widehat{A}_{T_k})} = \ln \prod_{i=1}^k \frac{\lambda_{k,i} + \mu_{k,i}}{\lambda_{k,i}} = \sum_{i=1}^k \ln \left( 1 + \frac{\mu_{k,i}}{\lambda_{k,i}} \right).$$

Hence, we arrive at

$$\sum_{t=1}^T \mathbf{g}_t^\top \widehat{A}_t^\dagger \mathbf{g}_t \leq r + \sum_{k=1}^r \sum_{i=1}^k \ln \left( 1 + \frac{\mu_{k,i}}{\lambda_{k,i}} \right).$$

To further bound the latter quantity, we use  $\lambda^* \leq \lambda_{k,i}$  and Jensen's inequality :

$$\begin{aligned} \sum_{k=1}^r \sum_{i=1}^k \ln \left( 1 + \frac{\mu_{k,i}}{\lambda_{k,i}} \right) &\leq \sum_{k=1}^r \sum_{i=1}^k \ln \left( 1 + \frac{\mu_{k,i}}{\lambda^*} \right) \\ &\leq \frac{(1+r)r}{2} \ln \left( 1 + \frac{2 \sum_{k=1}^r \sum_{i=1}^k \mu_{k,i}}{(1+r)r\lambda^*} \right). \end{aligned}$$

Finally noticing that

$$\sum_{i=1}^k \mu_{k,i} = \text{TR}(\widehat{A}_{T_{k+1}-1}) - \text{TR}(\widehat{A}_{T_k}) = \sum_{t=T_k+1}^{T_{k+1}-1} \text{TR}(\mathbf{g}_t \mathbf{g}_t^\top) = \sum_{t=T_k+1}^{T_{k+1}-1} \|\mathbf{g}_t\|_2^2$$

completes the proof.  $\square$

---

**Algorithm 6** Random Projection Sketch for RP-SON

---

**Internal State:**  $S$  and  $H$ .**SketchInit**( $\alpha, m$ )

- 1: Set  $S = \mathbf{0}_{m \times d}$  and  $H = \frac{1}{\alpha} \mathbf{I}_m$ .
- 2: Return  $(S, H)$ .

**SketchUpdate**( $\hat{\mathbf{g}}$ )

- 1: Draw  $\mathbf{r} \sim \mathcal{N}(0, \frac{1}{\sqrt{m}})$  and update  $S \leftarrow S + \mathbf{r} \hat{\mathbf{g}}^\top$ .
  - 2: Compute  $\mathbf{q} = S \hat{\mathbf{g}} - \frac{\hat{\mathbf{g}}^\top \hat{\mathbf{g}}}{2} \mathbf{r}$ , update  $H \leftarrow H - \frac{H \mathbf{q} \mathbf{r}^\top H}{1 + \mathbf{r}^\top H \mathbf{q}}$  and  $H \leftarrow H - \frac{H \mathbf{r} \mathbf{q}^\top H}{1 + \mathbf{q}^\top H \mathbf{r}}$ .
  - 3: Return  $(S, H)$ .
- 

Taken together, Eq. (8) and Theorem 4 lead to the following regret bounds (recall the definitions of  $\lambda^*$  and  $r$  from Theorem 4).

**Corollary 1.** *If  $\sigma_t = 0$  for all  $t$  and  $\eta_t$  is set to be  $\frac{1}{CL} \sqrt{\frac{d}{t}}$ , then the regret of the algorithm defined by Eq. (7) is at most*

$$\frac{CL}{2} \sqrt{\frac{T}{d}} \left( r + \frac{(1+r)r}{2} \ln \left( 1 + \frac{2 \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{(1+r)r\lambda^*} \right) \right) + 4CL\sqrt{Td}.$$

On the other hand, if  $\sigma_t \geq \sigma > 0$  for all  $t$  and  $\eta_t$  is set to be 0, then the regret is at most

$$\frac{1}{2\sigma} \left( r + \frac{(1+r)r}{2} \ln \left( 1 + \frac{2 \sum_{t=1}^T \|\mathbf{g}_t\|_2^2}{(1+r)r\lambda^*} \right) \right).$$

## E Regret Bound for RP-SON

The pseudocode of the RP sketch is presented in Algorithm 6. Recall the notation  $R_G$  and  $R_D$  in Proposition 1 and let  $r$  be the rank of  $G_T$ , we prove the following regret bound:

**Theorem 5.** *Under Assumptions 1 and 2, if the sketch size  $m = \Omega((r + \ln(T/\delta))\epsilon^{-2})$ , then RP-SON ensures*

- (1)  $\mathbb{E}[R_D] \leq 4(CL)^2 \sum_{t=1}^T \eta_t$ , and
- (2)  $R_G \leq \frac{1}{1-\epsilon} \sum_{t=1}^T \mathbf{g}_t^\top (\alpha \mathbf{I}_d + G_t^\top G_t)^{-1} \mathbf{g}_t$  with probability at least  $1 - \delta$ .

*Proof.* We apply the property of the random projection method (see for example [36, Theorem 2.3]): as long as  $m = \Omega((r + \ln(T/\delta))\epsilon^{-2})$ , with probability at least  $1 - \delta$ ,

$$(1 - \epsilon)G_t^\top G_t \preceq S_t^\top S_t \preceq (1 + \epsilon)G_t^\top G_t \quad \text{for all } t = 1, \dots, T$$

which implies  $A_t^{-1} \preceq \frac{1}{1-\epsilon}(\alpha \mathbf{I}_d + G_t^\top G_t)^{-1}$  and thus  $R_G \leq \frac{1}{1-\epsilon} \sum_{t=1}^T \mathbf{g}_t^\top (\alpha \mathbf{I}_d + G_t^\top G_t)^{-1} \mathbf{g}_t$ . For  $R_D$ , first fix all the randomness before drawing  $\mathbf{r}_t$  and let  $\mathbb{E}_t$  be the corresponding conditional expectation, then we have

$$\mathbb{E}_t[A_t - A_{t-1}] = \mathbb{E}_t \left[ S_{t-1}^\top \mathbf{r}_t \hat{\mathbf{g}}_t^\top + \hat{\mathbf{g}}_t \mathbf{r}_t^\top S_{t-1} + \|\mathbf{r}_t\|_2^2 \hat{\mathbf{g}}_t \hat{\mathbf{g}}_t^\top \right] = (\sigma_t + \eta_t) \mathbf{g}_t \mathbf{g}_t^\top.$$

Since  $\mathbf{w}_t, \mathbf{w}$  and  $\mathbf{g}_t$  are fixed, we continue with

$$\mathbb{E}_t \left[ (\mathbf{w}_t - \mathbf{w})^\top (A_t - A_{t-1} - \sigma_t \mathbf{g}_t \mathbf{g}_t^\top) (\mathbf{w}_t - \mathbf{w}) \right] = \eta_t (\mathbf{w}_t - \mathbf{w})^\top \mathbf{g}_t \mathbf{g}_t^\top (\mathbf{w}_t - \mathbf{w}) \leq 4(CL)^2 \eta_t.$$

Therefore, taking the overall expectation gives  $\mathbb{E}[R_D] \leq 4(CL)^2 \sum_{t=1}^T \eta_t$ .  $\square$

This theorem implies that the bound on  $R_D$  is the same as the one without using sketch, and the term  $R_G$  is only constant larger.

## F Proof of Theorem 3

*Proof.* We again first apply Proposition 1 (recall the notation  $R_G$  and  $R_D$  stated in the proposition). By the construction of the sketch, we have

$$A_t - A_{t-1} = S_t^\top S_t - S_{t-1}^\top S_{t-1} = \widehat{\mathbf{g}}_t \widehat{\mathbf{g}}_t^\top - \rho_t V_t^\top V_t \preceq \widehat{\mathbf{g}}_t \widehat{\mathbf{g}}_t^\top.$$

It follows immediately that  $R_D$  is again at most  $4(CL)^2 \sum_{t=1}^T \eta_t$ . For the term  $R_G$ , we will apply the following guarantee of Frequent Directions (see the proof of Theorem 1.1 of [12]):  $\sum_{t=1}^T \rho_t \leq \frac{\Omega_k}{m-k}$ . Specifically, since  $\text{TR}(V_t A_t^{-1} V_t^\top) \leq \frac{1}{\alpha} \text{TR}(V_t V_t^\top) = \frac{m}{\alpha}$  we have

$$\begin{aligned} R_G &= \sum_{t=1}^T \frac{1}{\sigma_t + \eta_t} \langle A_t^{-1}, A_t - A_{t-1} + \rho_t V_t^\top V_t \rangle \\ &\leq \frac{1}{\sigma + \eta_T} \sum_{t=1}^T (\langle A_t^{-1}, A_t - A_{t-1} + \rho_t V_t^\top V_t \rangle) \\ &= \frac{1}{\sigma + \eta_T} \sum_{t=1}^T (\langle A_t^{-1}, A_t - A_{t-1} \rangle + \rho_t \text{TR}(V_t A_t^{-1} V_t^\top)) \\ &\leq \frac{1}{(\sigma + \eta_T)} \sum_{t=1}^T \langle A_t^{-1}, A_t - A_{t-1} \rangle + \frac{m\Omega_k}{(m-k)(\sigma + \eta_T)\alpha}. \end{aligned}$$

Finally for the term  $\sum_{t=1}^T \langle A_t^{-1}, A_t - A_{t-1} \rangle$ , we proceed similarly to the proof of Theorem 2:

$$\begin{aligned} \sum_{t=1}^T \langle A_t^{-1}, A_t - A_{t-1} \rangle &\leq \sum_{t=1}^T \ln \frac{|A_t|}{|A_{t-1}|} = \ln \frac{|A_T|}{|A_0|} = \sum_{i=1}^d \ln \left( 1 + \frac{\lambda_i(S_T^\top S_T)}{\alpha} \right) \\ &= \sum_{i=1}^m \ln \left( 1 + \frac{\lambda_i(S_T^\top S_T)}{\alpha} \right) \leq m \ln \left( 1 + \frac{\text{TR}(S_T^\top S_T)}{m\alpha} \right) \end{aligned}$$

where the first inequality is by the concavity of the function  $\ln |X|$ , the second one is by Jensen's inequality, and the last equality is by the fact that  $S_T^\top S_T$  is of rank  $m$  and thus  $\lambda_i(S_T^\top S_T) = 0$  for any  $i > m$ . This concludes the proof.  $\square$

## G Sparse updates for RP and FD sketches

### G.1 Random Projection

We recall the updates of RP sketch. Since  $\widehat{\mathbf{g}}_t$  is sparse,  $S_t = S_{t-1} + r\widehat{\mathbf{g}}_t^\top$  is easily updated in  $\mathcal{O}(ms)$  time.  $H_t$  can also be updated in  $\mathcal{O}(m^2 + ms)$  time clearly. However, since the sketch  $S_t$  is getting denser and denser, direct update of the weight vector is a dense operation too. The solution is to represent and store  $\mathbf{w}_t$  in the form of  $\bar{\mathbf{w}}_t + S_{t-1}^\top \mathbf{b}_t$  for some  $\bar{\mathbf{w}}_t \in \mathbb{R}^d$  and  $\mathbf{b}_t \in \mathbb{R}^m$ . Note that now computing the prediction  $\mathbf{w}_t^\top \mathbf{x}_t$  needs  $\mathcal{O}(ms)$  time. Rewriting the update rules we have

$$\begin{aligned} \mathbf{u}_{t+1} &= \mathbf{w}_t - \frac{1}{\alpha} \mathbf{g}_t + \frac{1}{\alpha} S_t^\top H_t S_t \mathbf{g}_t = \bar{\mathbf{w}}_t + S_{t-1}^\top \mathbf{b}_t - \frac{1}{\alpha} \mathbf{g}_t + \frac{1}{\alpha} S_t^\top H_t S_t \mathbf{g}_t \\ &= \underbrace{\bar{\mathbf{w}}_t - \widehat{\mathbf{g}}_t r^\top \mathbf{b}_t}_{\stackrel{\text{def}}{=} \bar{\mathbf{u}}_{t+1}} - \frac{1}{\alpha} \mathbf{g}_t + S_t^\top \underbrace{\left( \mathbf{b}_t + \frac{1}{\alpha} H_t S_t \mathbf{g}_t \right)}_{\stackrel{\text{def}}{=} \mathbf{b}'_{t+1}}. \end{aligned}$$

Since  $\mathbf{g}_t$  and  $\widehat{\mathbf{g}}_t$  are sparse, computing  $\bar{\mathbf{u}}_{t+1}$  and  $\mathbf{b}'_{t+1}$  needs  $\mathcal{O}(m^2 + ms)$  time. Finally, for the projection step,  $c_t$  can clearly be computed in  $\mathcal{O}(m^2 + ms)$  time, and the update rule of  $\bar{\mathbf{w}}_{t+1}$  and

---

**Algorithm 7** Sparse Sketched Online Newton with Random Projection
 

---

**Input:** Parameters  $C$ ,  $\alpha$  and  $m$ .

- 1: Initialize  $\bar{\mathbf{u}} = \mathbf{0}_{d \times 1}$ ,  $\mathbf{b} = \mathbf{0}_{m \times 1}$  and  $(S, H) \leftarrow \mathbf{SketchInit}(\alpha, m)$  (Algorithm 6).
  - 2: **for**  $t = 1$  **to**  $T$  **do**
  - 3:   Receive example  $\mathbf{x}_t$ .
  - 4:   Projection step: compute  $\hat{\mathbf{x}} = S\mathbf{x}_t$ ,  $\gamma = \frac{\tau_C(\bar{\mathbf{u}}^\top \mathbf{x}_t + \mathbf{b}^\top \hat{\mathbf{x}})}{\mathbf{x}_t^\top \mathbf{x}_t - \hat{\mathbf{x}}^\top H \hat{\mathbf{x}}}$ ,  $\bar{\mathbf{w}} = \bar{\mathbf{u}} - \gamma \mathbf{x}_t$  and  $\mathbf{b} \leftarrow \mathbf{b} + cH\hat{\mathbf{x}}$ .
  - 5:   Predict label  $y_t = \bar{\mathbf{w}}^\top \mathbf{x}_t + \mathbf{b}^\top \hat{\mathbf{x}}$  and suffer loss  $\ell_t(y_t)$ .
  - 6:   Compute gradient  $\mathbf{g}_t = \ell'_t(y_t)\mathbf{x}_t$  and the to-sketch vector  $\hat{\mathbf{g}} = \sqrt{\sigma_t + \eta_t}\mathbf{g}_t$ .
  - 7:    $(S, H) \leftarrow \mathbf{SketchUpdate}(\hat{\mathbf{g}})$  (Algorithm 6).
  - 8:   Update  $\bar{\mathbf{u}} = \bar{\mathbf{w}} - (r^\top \mathbf{b})\hat{\mathbf{g}} - \frac{1}{\alpha}\mathbf{g}_t$  and  $\mathbf{b} \leftarrow \mathbf{b} + \frac{1}{\alpha}HS\mathbf{g}_t$ .
  - 9: **end for**
- 

$\mathbf{b}_{t+1}$  is thus derived as follows:

$$\begin{aligned}
 \mathbf{w}_{t+1} &= \mathbf{u}_{t+1} - \gamma_t(\mathbf{x}_{t+1} - S_t^\top H_t S_t \mathbf{x}_{t+1}) \\
 &= \bar{\mathbf{u}}_{t+1} + S_t^\top \mathbf{b}'_{t+1} - \gamma_t(\mathbf{x}_{t+1} - S_t^\top H_t S_t \mathbf{x}_{t+1}) \\
 &= \underbrace{\bar{\mathbf{u}}_{t+1} - \gamma_t \mathbf{x}_{t+1}}_{\stackrel{\text{def}}{=} \bar{\mathbf{w}}_{t+1}} + S_t^\top \underbrace{(\mathbf{b}'_{t+1} + \gamma_t H_t S_t \mathbf{x}_{t+1})}_{\stackrel{\text{def}}{=} \mathbf{b}_{t+1}}
 \end{aligned}$$

which again takes  $\mathcal{O}(m^2 + ms)$  time. Taken together, the total time complexity per round is  $\mathcal{O}(m^2 + ms)$ . The pseudocode for this version of the algorithm is presented in Algorithm 7.

## G.2 Frequent Directions

The sparse version of our algorithm with the Frequent Directions option is much more involved. We begin by taking a detour and introducing a fast and epoch-based variant of the Frequent Directions algorithm proposed in [12]. The idea is the following: instead of doing an eigendecomposition immediately after inserting a new  $\hat{\mathbf{g}}$  every round, we double the size of the sketch (to  $2m$ ), keep up to  $m$  recent  $\hat{\mathbf{g}}$ 's, do the decomposition only at the end of every  $m$  rounds and finally keep the top  $m$  eigenvectors with shrunk eigenvalues. The advantage of this variant is that it can be implemented straightforwardly in  $\mathcal{O}(md)$  time on average without doing a complicated rank-one SVD update, while still ensuring the exact same guarantee with the only price of doubling the sketch size.

Algorithm 8 shows the details of this variant and how we maintain  $H$ . The sketch  $S$  is always represented by two parts: the top part ( $DV$ ) comes from the last eigendecomposition, and the bottom part ( $G$ ) collects the recent to-sketch vector  $\hat{\mathbf{g}}$ 's. Note that within each epoch, the update of  $H^{-1}$  is a rank-two update and thus  $H$  can be updated similarly to the case of random projection (Lines 3 and 4 of Algorithm 8).

Although we can use any available algorithm that runs in  $\mathcal{O}(m^2d)$  time to do the eigendecomposition (Line 7 in Algorithm 8), we explicitly write down the procedure of reducing this problem to eigendecomposing a small square matrix in Algorithm 9, which will be important for deriving the sparse version of the algorithm. Lemma 3 proves that Algorithm 9 works correctly for finding the top  $m$  eigenvector and eigenvalues.

**Lemma 3.** *The outputs of Algorithm 9 are such that the  $i$ -th row of  $V'$  and the  $i$ -th entry of the diagonal of  $\Sigma$  are the  $i$ -th eigenvector and eigenvalue of  $S^\top S$  respectively.*

*Proof.* Let  $W^\top \in \mathbb{R}^{d \times (d-m-r)}$  be an orthonormal basis of the null space of  $\begin{pmatrix} V \\ Q \end{pmatrix}$ . By Line 2, we know that  $GW^\top = \mathbf{0}$  and  $E = (V^\top \ Q^\top \ W^\top)$  forms an orthonormal basis of  $\mathbb{R}^d$ . Therefore, we



---

**Algorithm 8** Frequent Direction Sketch (epoch version)

---

**Internal State:**  $\tau, D, V, G$  and  $H$ .**SketchInit**( $\alpha, m$ )

- 1: Set  $\tau = 1, D = \mathbf{0}_{m \times m}, G = \mathbf{0}_{m \times d}, H = \frac{1}{\alpha} \mathbf{I}_{2m}$  and let  $V$  be any  $m \times d$  matrix whose rows are orthonormal.
- 2: Return  $(\mathbf{0}_{2m \times d}, H)$ .

**SketchUpdate**( $\hat{g}$ )

- 1: Insert  $\hat{g}$  into the  $\tau$ -th row of  $G$ .
  - 2: **if**  $\tau < m$  **then**
    - 3: Let  $e$  be the  $2m \times 1$  basis vector whose  $(m + \tau)$ -th entry is 1 and  $\mathbf{q} = S\hat{g} - \frac{\hat{g}^\top \hat{g}}{2} e$ .
    - 4: Update  $H \leftarrow H - \frac{Hqe^\top H}{1+e^\top Hq}$  and  $H \leftarrow H - \frac{Heq^\top H}{1+q^\top He}$ .
    - 5: Update  $\tau \leftarrow \tau + 1$ .
  - 6: **else**
    - 7:  $(V, \Sigma) \leftarrow \mathbf{ComputeEigenSystem} \left( \begin{pmatrix} DV \\ G \end{pmatrix} \right)$  (Algorithm 9).
    - 8: Set  $D$  to be a diagonal matrix with  $D_{i,i} = \sqrt{\Sigma_{i,i} - \Sigma_{m,m}}, \forall i \in [m]$ .
    - 9: Set  $H \leftarrow \text{diag} \left\{ \frac{1}{\alpha + D_{1,1}^2}, \dots, \frac{1}{\alpha + D_{m,m}^2}, \frac{1}{\alpha}, \dots, \frac{1}{\alpha} \right\}$ .
    - 10: Set  $G = \mathbf{0}_{m \times d}$ .
    - 11: Set  $\tau = 1$ .
  - 12: **end if**
  - 13: Return  $\left( \begin{pmatrix} DV \\ G \end{pmatrix}, H \right)$ .
- 

---

**Algorithm 9** ComputeEigenSystem( $S$ )

---

**Input:**  $S = \begin{pmatrix} DV \\ G \end{pmatrix}$ .**Output:**  $V' \in \mathbb{R}^{m \times d}$  and diagonal matrix  $\Sigma \in \mathbb{R}^{m \times m}$  such that the  $i$ -th row of  $V'$  and the  $i$ -th entry of the diagonal of  $\Sigma$  are the  $i$ -th eigenvector and eigenvalue of  $S^\top S$  respectively.

- 1: Compute  $M = GV^\top$ .
  - 2: Decompose  $G - MV$  into the form  $LQ$  where  $L \in \mathbb{R}^{m \times r}, Q$  is a  $r \times d$  matrix whose rows are orthonormal and  $r$  is the rank of  $G - MV$  (e.g. by a Gram-Schmidt process).
  - 3: Compute the top  $m$  eigenvectors ( $U \in \mathbb{R}^{m \times (m+r)}$ ) and eigenvalues ( $\Sigma \in \mathbb{R}^{m \times m}$ ) of the matrix  $\begin{pmatrix} D^2 & \mathbf{0}_{m \times r} \\ \mathbf{0}_{r \times m} & \mathbf{0}_{r \times r} \end{pmatrix} + \begin{pmatrix} M^\top \\ L^\top \end{pmatrix} \begin{pmatrix} M & L \end{pmatrix}$ .
  - 4: Return  $(V', \Sigma)$  where  $V' = U \begin{pmatrix} V \\ Q \end{pmatrix}$ .
- 

have

$$\begin{aligned} S^\top S &= V^\top D^2 V + G^\top G \\ &= E \begin{pmatrix} D^2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix} E^\top + EE^\top G^\top GEE^\top \\ &= E \left( \begin{pmatrix} D^2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix} + \begin{pmatrix} VG^\top \\ QG^\top \\ WG^\top \end{pmatrix} (GV^\top GQ^\top GW^\top) \right) E^\top \\ &= (V^\top Q^\top) \underbrace{\left( \begin{pmatrix} D^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} + \begin{pmatrix} M^\top \\ L^\top \end{pmatrix} \begin{pmatrix} M & L \end{pmatrix} \right)}_{=C} \begin{pmatrix} V \\ Q \end{pmatrix} \end{aligned}$$

where in the last step we use the fact  $GQ^\top = (MV + LQ)Q^\top = L$ . Now it is clear that the eigenvalue of  $C$  will be the eigenvalue of  $S^\top S$  and the eigenvector of  $C$  will be the eigenvector of  $S^\top S$  after left multiplied by matrix  $(V^\top Q^\top)$ , completing the proof.  $\square$

We are now ready to present the sparse version of SON with Frequent Direction sketch (Algorithm 10). The key point is that we represent  $V_t$  as  $F_t Z_t$  for some  $F_t \in \mathbb{R}^{m \times m}$  and  $Z_t \in \mathbb{R}^{m \times d}$ , and the weight vector  $\mathbf{w}_t$  as  $\bar{\mathbf{w}}_t + Z_{t-1}^\top \mathbf{b}_t$  and ensure that the update of  $Z_t$  and  $\bar{\mathbf{w}}_t$  will always be sparse. To see this, denote the sketch  $S_t$  by  $\begin{pmatrix} D_t F_t Z_t \\ G_t \end{pmatrix}$  and let  $H_{t,1}$  and  $H_{t,2}$  be the top and bottom half of  $H_t$ . Now the update rule of  $\mathbf{u}_{t+1}$  can be rewritten as

$$\begin{aligned} \mathbf{u}_{t+1} &= \mathbf{w}_t - (\mathbf{I}_d - S_t^\top H_t S_t) \frac{\mathbf{g}_t}{\alpha} \\ &= \bar{\mathbf{w}}_t + Z_{t-1}^\top \mathbf{b}_t - \frac{1}{\alpha} \mathbf{g}_t + \frac{1}{\alpha} (Z_t^\top F_t^\top D_t, G_t^\top) \begin{pmatrix} H_{t,1} S_t \mathbf{g}_t \\ H_{t,2} S_t \mathbf{g}_t \end{pmatrix} \\ &= \underbrace{\bar{\mathbf{w}}_t + \frac{1}{\alpha} (G_t^\top H_{t,2} S_t \mathbf{g}_t - \mathbf{g}_t)}_{\bar{\mathbf{u}}_{t+1}} - (Z_t - Z_{t-1})^\top \mathbf{b}_t + \underbrace{Z_t^\top \left( \mathbf{b}_t + \frac{1}{\alpha} F_t^\top D_t H_{t,1} S_t \mathbf{g}_t \right)}_{\mathbf{b}'_{t+1}} \end{aligned}$$

We will show that  $Z_t - Z_{t-1} = \Delta_t G_t$  for some  $\Delta_t \in \mathbb{R}^{m \times m}$  shortly, and thus the above update is efficient due to the fact that the rows of  $G_t$  are collections of previous sparse vectors  $\hat{\mathbf{g}}$ .

Similarly, the update of  $\mathbf{w}_{t+1}$  can be written as

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{u}_{t+1} - \gamma_t (\mathbf{x}_{t+1} - S_t^\top H_t S_t \mathbf{x}_{t+1}) \\ &= \bar{\mathbf{u}}_{t+1} + Z_t^\top \mathbf{b}'_{t+1} - \gamma_t \mathbf{x}_{t+1} + \gamma_t (Z_t^\top F_t^\top D_t, G_t^\top) \begin{pmatrix} H_{t,1} S_t \mathbf{x}_{t+1} \\ H_{t,2} S_t \mathbf{x}_{t+1} \end{pmatrix} \\ &= \underbrace{\bar{\mathbf{u}}_{t+1} + \gamma_t (G_t^\top H_{t,2} S_t \mathbf{x}_{t+1} - \mathbf{x}_{t+1})}_{\bar{\mathbf{w}}_{t+1}} + Z_t^\top \underbrace{(\mathbf{b}'_{t+1} + \gamma_t F_t^\top D_t H_{t,1} S_t \mathbf{x}_{t+1})}_{\mathbf{b}_{t+1}}. \end{aligned}$$

It is clear that  $\gamma_t$  can be computed efficiently, and thus the update of  $\mathbf{w}_{t+1}$  is also efficient. These updates correspond to Line 6 and 10 of Algorithm 10.

It remains to perform the sketch update efficiently. Algorithm 11 is the sparse version of Algorithm 8. The challenging part is to compute eigenvectors and eigenvalues efficiently. Fortunately, in light of Algorithm 9, using the new representation  $V = FZ$  one can directly translate the process to Algorithm 12 and find that the eigenvectors can be expressed in the form  $N_1 Z + N_2 G$ . To see this, first note that Line 1 of both algorithms compute the same matrix  $M = GV^\top = GZ^\top F^\top$ . Then Line 2 decomposes the matrix

$$G - MV = G - MFZ = \begin{pmatrix} -MF & \mathbf{I}_m \end{pmatrix} \begin{pmatrix} Z \\ G \end{pmatrix} \stackrel{\text{def}}{=} PR$$

using Gram-Schmidt into the form  $LQR$  such that the rows of  $QR$  are orthonormal (that is,  $QR$  corresponds to  $Q$  in Algorithm 9). While directly applying Gram-Schmidt to  $PR$  would take  $\mathcal{O}(m^2 d)$  time, this step can in fact be efficiently implemented by performing Gram-Schmidt to  $P$  (instead of  $PR$ ) in a Banach space where inner product is defined as  $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top K \mathbf{b}$  with

$$K = RR^\top = \begin{pmatrix} ZZ^\top & ZG^\top \\ GZ^\top & GG^\top \end{pmatrix}$$

being the Gram matrix of  $R$ . Since we can efficiently maintain the Gram matrix of  $Z$  (see Line 10 of Algorithm 11) and  $GZ^\top$  and  $GG^\top$  can be computed sparsely, this decomposing step can be done efficiently too. This modified Gram-Schmidt algorithm is presented in Algorithm 13 (which will also be used in sparse Oja's sketch), where Line 4 is the key difference compared to standard Gram-Schmidt (see Lemma 4 below for a formal proof of correctness).

Line 3 of Algorithms 9 and 12 are exactly the same. Finally the eigenvectors  $U \begin{pmatrix} V \\ Q \end{pmatrix}$  in Algorithm 9 now becomes (with  $U_1, U_2, Q_1, Q_2, N_1, N_2$  defined in Line 4 of Algorithm 12)

$$\begin{aligned} U \begin{pmatrix} FZ \\ QR \end{pmatrix} &= (U_1, U_2) \begin{pmatrix} FZ \\ QR \end{pmatrix} = U_1 FZ + U_2 (Q_1, Q_2) \begin{pmatrix} Z \\ G \end{pmatrix} \\ &= (U_1 FZ + U_2 Q_1) Z + U_2 Q_2 G = N_1 Z + N_2 G. \end{aligned}$$

---

**Algorithm 10** Sparse Sketched Online Newton with Frequent Directions

---

**Input:** Parameters  $C$ ,  $\alpha$  and  $m$ .

- 1: Initialize  $\bar{\mathbf{u}} = \mathbf{0}_{d \times 1}$ ,  $\mathbf{b} = \mathbf{0}_{m \times 1}$  and  $(D, F, Z, G, H) \leftarrow \mathbf{SketchInit}(\alpha, m)$  (Algorithm 11).
  - 2: Let  $S$  denote the matrix  $\begin{pmatrix} DFZ \\ G \end{pmatrix}$  throughout the algorithm (without actually computing it).
  - 3: Let  $H_1$  and  $H_2$  denote the upper and lower half of  $H$ , i.e.  $H = \begin{pmatrix} H_1 \\ H_2 \end{pmatrix}$ .
  - 4: **for**  $t = 1$  **to**  $T$  **do**
  - 5:   Receive example  $\mathbf{x}_t$ .
  - 6:   Projection step: compute  $\hat{\mathbf{x}} = S\mathbf{x}_t$  and  $\gamma = \frac{\tau C(\bar{\mathbf{u}}^\top \mathbf{x}_t + \mathbf{b}^\top Z\mathbf{x}_t)}{\mathbf{x}_t^\top \mathbf{x}_t - \hat{\mathbf{x}}^\top H\hat{\mathbf{x}}}$ .  
    Obtain  $\bar{\mathbf{w}} = \bar{\mathbf{u}} + \gamma(G^\top H_2 \hat{\mathbf{x}} - \mathbf{x}_t)$  and  $\mathbf{b} \leftarrow \mathbf{b} + \gamma F^\top D H_1 \hat{\mathbf{x}}$ .
  - 7:   Predict label  $y_t = \bar{\mathbf{w}}^\top \mathbf{x}_t + \mathbf{b}^\top Z\mathbf{x}_t$  and suffer loss  $\ell_t(y_t)$ .
  - 8:   Compute gradient  $\mathbf{g}_t = \ell'_t(y_t)\mathbf{x}_t$  and the to-sketch vector  $\hat{\mathbf{g}} = \sqrt{\sigma_t + \eta_t}\mathbf{g}_t$ .
  - 9:    $(D, F, Z, G, H, \Delta) \leftarrow \mathbf{SketchUpdate}(\hat{\mathbf{g}})$  (Algorithm 11).
  - 10:   Update  $\bar{\mathbf{u}} = \bar{\mathbf{w}} + \frac{1}{\alpha}(G^\top H_2 S\mathbf{g} - \mathbf{g}) - G^\top \Delta^\top \mathbf{b}$  and  $\mathbf{b} \leftarrow \mathbf{b} + \frac{1}{\alpha}F^\top D H_1 S\mathbf{g}$ .
  - 11: **end for**
- 

---

**Algorithm 11** Sparse Frequent Direction Sketch

---

**Internal State:**  $\tau, D, F, Z, G, H$  and  $K$ .

**SketchInit**( $\alpha, m$ )

- 1: Set  $\tau = 1$ ,  $D = \mathbf{0}_{m \times m}$ ,  $F = K = \mathbf{I}_m$ ,  $H = \frac{1}{\alpha}\mathbf{I}_{2m}$ ,  $G = \mathbf{0}_{m \times d}$ , and let  $Z$  be any  $m \times d$  matrix whose rows are orthonormal.
- 2: Return  $(D, F, Z, G, H)$ .

**SketchUpdate**( $\hat{\mathbf{g}}$ )

- 1: Insert  $\hat{\mathbf{g}}$  into the  $\tau$ -th row of  $G$ .
  - 2: **if**  $\tau < m$  **then**
  - 3:   Let  $\mathbf{e}$  be the  $2m \times 1$  basic vector whose  $(m + \tau)$ -th entry is 1 and compute  $\mathbf{q} = S\hat{\mathbf{g}} - \frac{\hat{\mathbf{g}}^\top \hat{\mathbf{g}}}{2}\mathbf{e}$ .
  - 4:   Update  $H \leftarrow H - \frac{H\mathbf{q}\mathbf{e}^\top H}{1 + \mathbf{e}^\top H\mathbf{q}}$  and  $H \leftarrow H - \frac{H\mathbf{e}\mathbf{q}^\top H}{1 + \mathbf{q}^\top H\mathbf{e}}$ .
  - 5:   Set  $\Delta = \mathbf{0}_{m \times m}$ .
  - 6:   Set  $\tau \leftarrow \tau + 1$ .
  - 7: **else**
  - 8:    $(N_1, N_2, \Sigma) \leftarrow \mathbf{ComputeSparseEigenSystem}\left(\begin{pmatrix} DFZ \\ G \end{pmatrix}, K\right)$  (Algorithm 12).
  - 9:   Compute  $\Delta = N_1^{-1}N_2$ .
  - 10:   Update Gram matrix  $K \leftarrow K + \Delta G Z^\top + Z G^\top \Delta^\top + \Delta G G^\top \Delta^\top$ .
  - 11:   Update  $F = N_1$ ,  $Z \leftarrow Z + \Delta G$ , and let  $D$  be such that  $D_{i,i} = \sqrt{\Sigma_{i,i} - \Sigma_{m,m}}$ ,  $\forall i \in [m]$ .
  - 12:   Set  $H \leftarrow \text{diag}\left\{\frac{1}{\alpha + D_{1,1}^2}, \dots, \frac{1}{\alpha + D_{m,m}^2}, \frac{1}{\alpha}, \dots, \frac{1}{\alpha}\right\}$ .
  - 13:   Set  $G = \mathbf{0}_{m \times d}$ .
  - 14:   Set  $\tau = 1$ .
  - 15: **end if**
  - 16: Return  $(D, F, Z, G, H, \Delta)$ .
- 

Therefore, having the eigenvectors in the form  $N_1 Z + N_2 G$ , we can simply update  $F$  as  $N_1$  and  $Z$  as  $Z + N_1^{-1}N_2 G$  so that the invariant  $V = FZ$  still holds (see Line 11 of Algorithm 11). The update of  $Z$  is sparse since  $G$  is sparse.

We finally summarize the results of this section in the following theorem.

**Theorem 6.** *The average running time of Algorithm 10 is  $\mathcal{O}(m^2 + ms)$  per round, and the regret bound is exactly the same as the one stated in Theorem 3.*

**Lemma 4.** *The output of Algorithm 13 ensures that  $LQR = PR$  and the rows of  $QR$  are orthonormal.*

---

**Algorithm 12** ComputeSparseEigenSystem( $S, K$ )
 

---

**Input:**  $S = \begin{pmatrix} DFZ \\ G \end{pmatrix}$  and Gram matrix  $K = ZZ^\top$ .

**Output:**  $N_1, N_2 \in \mathbb{R}^{m \times m}$  and diagonal matrix  $\Sigma \in \mathbb{R}^{m \times m}$  such that the  $i$ -th row of  $N_1 Z + N_2 G$  and the  $i$ -th entry of the diagonal of  $\Sigma$  are the  $i$ -th eigenvector and eigenvalue of the matrix  $S^\top S$ .

- 1: Compute  $M = GZ^\top F^\top$ .
  - 2:  $(L, Q) \leftarrow$  Decompose  $\left( \begin{pmatrix} -MF & \mathbf{I}_m \end{pmatrix}, \begin{pmatrix} K & ZG^\top \\ GZ^\top & GG^\top \end{pmatrix} \right)$  (Algorithm 13).
  - 3: Let  $r$  be the number of columns of  $L$ . Compute the top  $m$  eigenvectors ( $U \in \mathbb{R}^{m \times (m+r)}$ ) and eigenvalues ( $\Sigma \in \mathbb{R}^{m \times m}$ ) of the matrix  $\begin{pmatrix} D^2 & \mathbf{0}_{m \times r} \\ \mathbf{0}_{r \times m} & \mathbf{0}_{r \times r} \end{pmatrix} + \begin{pmatrix} M^\top \\ L^\top \end{pmatrix} \begin{pmatrix} M & L \end{pmatrix}$ .
  - 4: Set  $N_1 = U_1 F + U_2 Q_1$  and  $N_2 = U_2 Q_2$  where  $U_1$  and  $U_2$  are the first  $m$  and last  $r$  columns of  $U$  respectively, and  $Q_1$  and  $Q_2$  are the left and right half of  $Q$  respectively.
  - 5: Return  $(N_1, N_2, \Sigma)$ .
- 

*Proof.* It suffices to prove that Algorithm 13 is exactly the same as using the standard Gram-Schmidt to decompose the matrix  $PR$  into  $L$  and an orthonormal matrix which can be written as  $QR$ . First note that when  $K = \mathbf{I}_n$ , Algorithm 13 is simply the standard Gram-Schmidt algorithm applied to  $P$ . We will thus go through Line 1-10 of Algorithm 13 with  $P$  replaced by  $PR$  and  $K$  by  $\mathbf{I}_n$  and show that it leads to the exact same calculations as running Algorithm 13 directly. For clarity, we add “ $\tilde{\cdot}$ ” to symbols to distinguish the two cases (so  $\tilde{P} = PR$  and  $\tilde{K} = \mathbf{I}_n$ ). We will inductively prove the invariance  $\tilde{Q} = QR$  and  $\tilde{L} = L$ . The base case  $\tilde{Q} = QR = \mathbf{0}$  and  $\tilde{L} = L = \mathbf{0}$  is trivial. Now assume it holds for iteration  $i - 1$  and consider iteration  $i$ . We have

$$\begin{aligned} \tilde{\alpha} &= \tilde{Q}\tilde{K}\tilde{p} = QR R^\top p = QKp = \alpha, \\ \tilde{\beta} &= \tilde{p} - \tilde{Q}^\top \tilde{\alpha} = R^\top p - (QR)^\top \alpha = R^\top (p - Q^\top \alpha) = R^\top \beta, \\ \tilde{c} &= \sqrt{\tilde{\beta}^\top \tilde{K}\tilde{\beta}} = \sqrt{(R^\top \beta)^\top (R^\top \beta)} = \sqrt{\beta^\top K\beta} = c, \end{aligned}$$

which clearly implies that after execution of Line 5-9, we again have  $\tilde{Q} = QR$  and  $\tilde{L} = L$ , finishing the induction.  $\square$

## H Details for sparse Oja’s algorithm

We finally provide the missing details for the sparse version of the Oja’s algorithm. Since we already discussed the updates for  $\tilde{w}_t$  and  $\tilde{b}_t$  in Section 4, we just need to describe how the updates for  $F_t$  and  $Z_t$  work. Recall that the dense Oja’s updates can be written in terms of  $F$  and  $Z$  as

$$\begin{aligned} \Lambda_t &= (\mathbf{I}_m - \Gamma_t)\Lambda_{t-1} + \Gamma_t \text{diag}\{F_{t-1}Z_{t-1}\hat{\mathbf{g}}_t\}^2 \\ F_t Z_t &\stackrel{\text{orth}}{\leftarrow} F_{t-1}Z_{t-1} + \Gamma_t F_{t-1}Z_{t-1}\hat{\mathbf{g}}_t\hat{\mathbf{g}}_t^\top = F_{t-1}(Z_{t-1} + F_{t-1}^{-1}\Gamma_t F_{t-1}Z_{t-1}\hat{\mathbf{g}}_t\hat{\mathbf{g}}_t^\top). \end{aligned} \tag{9}$$

Here, the update for the eigenvalues is straightforward. For the update of eigenvectors, first we let  $Z_t = Z_{t-1} + \delta_t \hat{\mathbf{g}}_t^\top$  where  $\delta_t = F_{t-1}^{-1}\Gamma_t F_{t-1}Z_{t-1}\hat{\mathbf{g}}_t$  (note that under the assumption of Footnote 4,  $F_t$  is always invertible). Now it is clear that  $Z_t - Z_{t-1}$  is a sparse rank-one matrix and the update of  $\tilde{u}_{t+1}$  is efficient. Finally it remains to update  $F_t$  so that  $F_t Z_t$  is the same as orthonormalizing  $F_{t-1}Z_t$ , which can in fact be achieved by applying the Gram-Schmidt algorithm to  $F_{t-1}$  in a Banach space where inner product is defined as  $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top K_t \mathbf{b}$  where  $K_t$  is the Gram matrix  $Z_t Z_t^\top$  (see Algorithm 13). Since we can maintain  $K_t$  efficiently based on the update of  $Z_t$ :

$$K_t = K_{t-1} + \delta_t \hat{\mathbf{g}}_t^\top Z_{t-1}^\top + Z_{t-1} \hat{\mathbf{g}}_t \delta_t^\top + (\hat{\mathbf{g}}_t^\top \hat{\mathbf{g}}_t) \delta_t \delta_t^\top,$$

the update of  $F_t$  can therefore be implemented in  $\mathcal{O}(m^3)$  time.

---

**Algorithm 13** Decompose( $P, K$ )

---

**Input:**  $P \in \mathbb{R}^{m \times n}$ ,  $K \in \mathbb{R}^{m \times m}$  such that  $K$  is the Gram matrix  $K = RR^\top$  for some matrix  $R \in \mathbb{R}^{n \times d}$  where  $n \geq m, d \geq m$ ,

**Output:**  $L \in \mathbb{R}^{m \times r}$  and  $Q \in \mathbb{R}^{r \times n}$  such that  $LQR = PR$  where  $r$  is the rank of  $PR$  and the rows of  $QR$  are orthonormal.

- 1: Initialize  $L = \mathbf{0}_{m \times m}$  and  $Q = \mathbf{0}_{m \times n}$ .
  - 2: **for**  $i = 1$  **to**  $m$  **do**
  - 3:   Let  $\mathbf{p}^\top$  be the  $i$ -th row of  $P$ .
  - 4:   Compute  $\alpha = QK\mathbf{p}, \beta = \mathbf{p} - Q^\top\alpha$  and  $c = \sqrt{\beta^\top K\beta}$ .
  - 5:   **if**  $c \neq 0$  **then**
  - 6:     Insert  $\frac{1}{c}\beta^\top$  to the  $i$ -th row of  $Q$ .
  - 7:   **end if**
  - 8:   Set the  $i$ -th entry of  $\alpha$  to be  $c$  and insert  $\alpha$  to the  $i$ -th row of  $L$ .
  - 9: **end for**
  - 10: Delete the all-zero columns of  $L$  and all-zero rows of  $Q$ .
  - 11: Return  $(L, Q)$ .
- 




Figure 4: Error rates for Oja-SON with different sketch sizes on splice dataset

## I Experiment Details

This section reports some detailed experimental results omitted from Section 5.2. Table 1 includes the description of benchmark datasets; Table 2 reports error rates on relatively small datasets to show that Oja-SON generally has better performance; Table 3 reports concrete error rates for the experiments described in Section 5.2; finally Table 4 shows that Oja’s algorithm estimates the eigenvalues accurately.

As mentioned in Section 5.2, we see substantial improvement for the *splice* dataset when using Oja’s sketch even after the diagonal adaptation. We verify that the condition number for this dataset before and after the diagonal adaptation are very close (682 and 668 respectively), explaining why a large improvement is seen using Oja’s sketch. Fig. 4 shows the decrease of error rates as Oja-SON with different sketch sizes sees more examples. One can see that even with  $m = 1$  Oja-SON already performs very well. This also matches our expectation since there is a huge gap between the top and second eigenvalues of this dataset (50.7 and 0.4 respectively).

Table 1: Datasets used in experiments

Dataset	#examples	avg. sparsity	#features
20news	18845	93.89	101631
a9a	48841	13.87	123
acoustic	78823	50.00	50
adult	48842	12.00	105
australian	690	11.19	14
breast-cancer	683	10.00	10
census	299284	32.01	401
cod-rna	271617	8.00	8
covtype	581011	11.88	54
diabetes	768	7.01	8
gisette	1000	4971.00	5000
heart	270	9.76	13
ijcnn1	91701	13.00	22
ionosphere	351	30.06	34
letter	20000	15.58	16
magic04	19020	9.99	10
mnist	11791	142.43	780
mushrooms	8124	21.00	112
rcv1	781265	75.72	43001
real-sim	72309	51.30	20958
splice	1000	60.00	60
w1a	2477	11.47	300
w8a	49749	11.65	300

Table 2: Error rates for Sketched Online Newton with different sketching algorithms

Dataset	RP-SON	FD-SON	Oja-SON
australian	<b>15.6</b>	16.0	15.8
breast-cancer	4.8	5.3	<b>3.7</b>
diabetes	35.5	35.4	<b>32.8</b>
mushrooms	0.5	0.5	<b>0.2</b>
splice	22.9	<b>22.6</b>	22.9

Table 3: Error rates for different algorithms (with best results bolded)

Dataset	Oja-SON				ADAGRAD
	Without Diagonal Adaptation		With Diagonal Adaptation		
	$m = 0$	$m = 10$	$m = 0$	$m = 10$	
20news	0.121338	0.121338	<b>0.049590</b>	<b>0.049590</b>	0.068020
a9a	0.204447	0.195203	<b>0.155953</b>	<b>0.155953</b>	0.156414
acoustic	0.305824	0.260241	<b>0.257894</b>	<b>0.257894</b>	0.259493
adult	0.199763	0.199803	<b>0.150830</b>	<b>0.150830</b>	0.181582
australian	0.366667	0.366667	0.162319	<b>0.157971</b>	0.289855
breast-cancer	0.374817	0.374817	<b>0.036603</b>	<b>0.036603</b>	0.358712
census	0.093610	0.062038	0.051479	<b>0.051439</b>	0.069629
cod-rna	0.175107	0.175107	0.049710	<b>0.049643</b>	0.081066
covtype	<b>0.042304</b>	<b>0.042312</b>	0.050827	0.050818	0.045507
diabetes	0.433594	0.433594	0.329427	<b>0.328125</b>	0.391927
gisette	0.208000	0.208000	<b>0.152000</b>	<b>0.152000</b>	0.154000
heart	0.477778	0.388889	<b>0.244444</b>	<b>0.244444</b>	0.362963
ijcnn1	0.046826	0.046826	<b>0.034536</b>	<b>0.034645</b>	0.036913
ionosphere	0.188034	<b>0.148148</b>	0.182336	0.182336	0.190883
letter	0.306650	0.232300	0.233250	<b>0.230450</b>	0.237350
magic04	0.000263	0.000263	<b>0.000158</b>	<b>0.000158</b>	0.000210
mnist	0.062336	0.062336	0.040031	<b>0.039182</b>	0.046561
mushrooms	0.003323	0.002339	0.002462	0.002462	<b>0.001969</b>
rcv1	0.055976	0.052694	0.052764	0.052766	<b>0.050938</b>
real-sim	0.045140	0.043577	<b>0.029498</b>	<b>0.029498</b>	0.031670
splice	0.343000	0.343000	0.294000	<b>0.229000</b>	0.301000
w1a	0.001615	0.001615	0.004845	0.004845	<b>0.003633</b>
w8a	<b>0.000101</b>	<b>0.000101</b>	0.000422	0.000422	0.000221

Table 4: Largest relative error between true and estimated top 10 eigenvalues using Oja's rule.

Dataset	Relative eigenvalue difference
a9a	0.90
australian	0.85
breast-cancer	5.38
diabetes	5.13
heart	4.36
ijcnn1	0.57
magic04	11.48
mushrooms	0.91
splice	8.23
w8a	0.95