

DYNAMIC LEAST-SQUARES REGRESSION

Anonymous authors

Paper under double-blind review

ABSTRACT

In large-scale supervised learning, after a model is trained with an initial dataset, a common challenge is how to exploit new incremental data without re-training the model from scratch. Motivated by this problem, we revisit the canonical problem of *dynamic* least-squares regression (LSR), where the goal is to learn a linear model over incremental training data. In this setup, data and labels $(\mathbf{A}^{(t)}, \mathbf{b}^{(t)}) \in \mathbb{R}^{t \times d} \times \mathbb{R}^t$ evolve in an online fashion ($t \gg d$), and the goal is to efficiently maintain an (approximate) solution of $\min_{\mathbf{x}^{(t)}} \|\mathbf{A}^{(t)} \mathbf{x}^{(t)} - \mathbf{b}^{(t)}\|_2$ for all $t \in [T]$. Our main result is a dynamic data structure which maintains an arbitrarily small constant approximate solution to dynamic LSR with amortized update time $O(d^{1+o(1)})$, almost matching the running time of the *static* (sketching-based) solution. By contrast, for *exact* (or $1/\text{poly}(n)$ -accuracy) solutions, we show a separation between the models, namely, that dynamic LSR requires $\Omega(d^{2-o(1)})$ amortized update time under the OMv Conjecture (Henzinger et al., STOC’15). Our data structure is fast, conceptually simple, easy to implement, and our experiments demonstrate their practicality on both synthetic and real-world datasets.

1 INTRODUCTION

The problem of least-squares regression (LSR) dates back to Gauss in 1821 (Stigler, 1981), and is the backbone of statistical inference (Hastie et al., 2001), signal processing (Rabiner & Gold, 1975), convex optimization (Bubeck, 2015), control theory (Chui, 1990) and network routing (Lee & Sidford, 2014; Madry, 2013). Given an overdetermined ($n \gg d$) linear system $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$, the goal is to find the solution vector \mathbf{x} that minimizes the mean squared error (MSE)

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2. \quad (1)$$

Among many other loss functions (e.g., ℓ_p) that have been studied for linear regression, ℓ_2 -regression has been the most popular choice as it is at the same time robust to outliers, and admits a *high-accuracy* efficient solution.

The computational task of least-squares regression arises naturally in high-dimensional statistics and has been the central of focus. The exact closed-form solution is given by the well-known Normal equation $\mathbf{x}^* = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$, which requires $O(nd^2)$ time to compute using naive matrix-multiplication, or $O(nd^{\omega-1}) \approx O(nd^{1.37})$ time using fast matrix-multiplication (FMM) (Strassen, 1969) for the current FMM exponent of $\omega \approx 2.37$ (Le Gall, 2014; Alman & Williams, 2021).

Despite the elegance and simplicity of this closed-form solution, in practice the latter runtime is often too slow, especially in modern data analysis applications where both the dimension of the feature space (d) and the size of datasets (n) are overwhelmingly large. A more modest objective in attempt to circumvent this computational overhead, is to seek an ϵ -accurate solution that satisfies

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq (1 + \epsilon) \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2.$$

This was the primary motivation behind the development of the *sketch-and-solve* paradigm, where the idea is to first compress the matrix into one with fewer ($\sim d/\epsilon^2$) rows and then to compute the standard LSR solution but over the smaller matrix. A long line of developments on this framework culminates in algorithms that run in close to *input-sparsity* time (Sarlos, 2006; Clarkson & Woodruff, 2017; Nelson & Nguyen, 2013; Chepurko et al., 2021). In particular, a direct application of sketch-and-solve yields an algorithm runs in $\tilde{O}(\text{nnz}(A)\epsilon^{-1} + d^\omega)^1$, which is near optimal in the “low

¹In this paper we use $\tilde{O}(\cdot)$ to hide polylogarithmic terms, and we use $O_\epsilon(\cdot)$ to hide $\text{poly}(\log d, \epsilon^{-1})$ terms.

precision” regime ($\epsilon = 1/\text{poly}(\log d)$). Interestingly, when combined with more sophisticated ideas of preconditioning and (conjugate) gradient descent, the runtime of this algorithm in terms of the error ϵ can be further improved to $\tilde{O}(\text{nnz}(A) \log(1/\epsilon) + d^\omega)$, which yields a *high precision* algorithm, i.e., it can efficiently solve the problem to within *polynomial accuracy* $\epsilon = 1/\text{poly}(d)$.

Dynamic least-squares In many real-world scenarios of the aforementioned applications, data is evolving in an online fashion either by nature or by design, and such applications require *maintaining* the solution (1) adaptively, where rows of the data matrix and their corresponding labels ($\mathbf{A}^{(t)}, \mathbf{b}^{(t)}$) arrive one-by-one *incrementally*. This is known as the dynamic least-squares regression problem.

The origins of dynamic least-squares regression was in control theory of the 1950’s (Plackett, 1950), in the context of *dynamical linear systems*. In this setup, the data matrix $[\mathbf{A}^{(t)}, \mathbf{b}^{(t)}]$ corresponds to the set of measurement and it evolves in an online (incremental) fashion, and the goal is to efficiently maintain the (exact) solution to a noisy linear system $\mathbf{b} := \mathbf{A}^{(t)}\mathbf{x}^{(t)} + \xi^{(t)}$ *without* recomputing the LSR solution from scratch. The *recursive least-squares* (RLS) framework and the celebrated *Kalman filter* (Kalman, 1960) provide a rather simple update rule for maintaining an *exact* solution for this problem, by maintaining the sample covariance matrix and using *Woodbury’s identity* (which assert that an incremental update to $[\mathbf{A}^{(t)}, \mathbf{b}^{(t)}]$ translates into a rank-1 update to the sample covariance matrix), and hence each update can be computed in $O(d^2)$ time (Kalman, 1960).

Beyond this classic motivation for dynamic LSR, a more timely motivation comes from modern deep learning applications: Most neural networks need to be frequently re-trained upon arrival on new training data, in order to improve prediction accuracy, and it is desirable to *avoid recomputing* weights from scratch. This problem of efficient incremental training of DNNs has been studied before in *elastic machine learning* (Liberty et al., 2020) and in the context of *continual learning* (Parisi et al., 2019). Our work sheds light on this question by analyzing the minimal computational resources required for ℓ_2 loss-minimization.

Despite the rich and versatile literature on static LSR, the understanding of the dynamic counterpart was so far quite limited: The previous best known result requires $O(d^2)$ amortized update time (by a direct application of the Woodbury identity). The basic questions we address in this papers are:

Is it possible to achieve faster update time for maintaining an exact solution? How about a small-approximate solution – Is it then possible to achieve amortized $O(d)$ or even input-sparsity time?

In this paper, we settle both of these questions and present an essentially complete characterization of the dynamic complexity of LSR.

1.1 OVERVIEW OF OUR RESULTS

Our first result is a negative answer to the first question above of maintaining exact (or polynomial-accuracy) LSR solutions in the dynamic setting – We prove that Kalman’s approach is essentially optimal, assuming the popular *Online Matrix-Vector* (OMv) Conjecture (Henzinger et al., 2015)²:

Theorem 1.1 (Hardness of exact dynamic LSR, informal). *Assuming the OMv Conjecture, any dynamic algorithm that maintains an $\epsilon = 1/\text{poly}(d)$ -approximate solution for the dynamic LSR problem over $T = \text{poly}(d)$ iterations, must have $\Omega(d^{2-o(1)})$ amortized update time per iteration.*

Theorem 1.1 separates the static and the dynamic complexities of the exact LSR problem: As mentioned above, the static problem can be solved by *batching* rows together using FMM in time $O(Td^{\omega-1})$, whereas the dynamic problem requires $\Omega(Td^2)$ by Theorem 1.1. Indeed, the implication Theorem 1.1 is stronger, it also separates the static and dynamic complexity of approximate LSR problem under the high precision regime, it asserts that a polylogarithmic dependence on the precision (i.e. $d \text{poly}(\log(1/\epsilon))$) on update time is impossible (assuming OMv), in sharp contrast to the static case.

We next focus on an *approximate* version of this classic online problem, *dynamic ϵ -LSR*, where the goal is to efficiently maintain, during *all* iterations $t \in [T]$, an ϵ -approximate solution under incremental row-updates to $\mathbf{A}^{(t)}$ and labels $\mathbf{b}^{(t)}$, where efficiency is measured by the *amortized*

²This conjecture postulates that multiplying a fixed $d \times d$ matrix A with an *online* matrix B , column-by-column (AB_t), requires $d^{3-o(1)}$ time, in sharp contrast to the batch setting where this can be done using FMM in $d^\omega \ll d^3$ time. See Section 4.

update time for inserting a new row. A natural complexity benchmark for this dynamic problem is the aforementioned best *static* sketch-and-solve solution, which for $n = T$ is $\tilde{O}(\text{nnz}(\mathbf{A}^{(T)})\epsilon^{-1} + d^\omega) = \tilde{O}(\text{nnz}(\mathbf{A})\epsilon^{-1})$ for $T \gg d$. Our main result is a provably efficient and practical dynamic data structure, whose total running time essentially matches the complexity of the offline problem:

Theorem 1.2 (Main result, informal version of Theorem 3.1). *For any accuracy parameter $\epsilon > 0$, there is a randomized dynamic data structure which, with probability at least 0.9, maintains an ϵ -approximate solution to the dynamic LSR problem simultaneously for all iterations $t \in [T]$, with total update time*

$$O(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(T)).$$

Theorem 1.2 almost matches the fastest static sketching-based solution, up to polylogarithmic terms and the additive FMM term. When $T \gg d$, this theorem shows that amortized update time of our algorithm is $O(d^{1+o(1)})$.

1.2 RELATED WORK

Sketching and sampling The least squares regression as a fundamental problem has been extensively studied in the literature. A long line of work (Ailon & Chazelle, 2006; Clarkson & Woodruff, 2017; Nelson & Nguyen, 2013; Avron et al., 2017; Cohen et al., 2015; Woodruff, 2014; 2021) have focused on using dimension reduction technique (sketching or sampling) to speedup the computation task, culminates into algorithms that run in $\tilde{O}(\text{nnz}(\mathbf{A}) \log(1/\epsilon) + d^\omega)$ time (Clarkson & Woodruff, 2017). See Appendix A for detailed discussions.

Regression in online, streaming, and sliding window models Least-squares regressions have also been studied in various computational models, though the focus of these models are generally not the (amortized) running time. Our algorithm uses techniques developed by Cohen et al. (2020), where they study the regression problem in the online model, with the goal of maintaining a spectral approximation of data matrix in the online stream. Their algorithm only needs to store $O_\epsilon(d)$ rows but the amortized running time is still $\Omega(d^2)$ (see Section 3.1 for detailed discussions). In the streaming model, the main focus is the space complexity, and a direct application of random Gaussian sketch or count sketch reduces the space complexity to $\tilde{O}(d^2\epsilon^{-1})$ and it is shown to be tight (Clarkson & Woodruff, 2009). Recent work of (Braverman et al., 2020) studies regressions and other numerical linear algebra tasks in the sliding window model, where data come in an online stream and only the most recent updates form the underlying data set. The major focus of a sliding window model is still the space complexity, and there is no amortized running time guarantee.

Disparity from online learning Our work crucially differs from online learning literature (Hazan, 2019), in that the main bottleneck in online regret-minimization and bandit problems is information-theoretic, whereas the challenge in our loss-minimization problem is purely computational. See Appendix A for detailed discussions.

2 PROBLEM FORMULATION

In a dynamic least-squares regression problem, initially, we are given a matrix $\mathbf{A}^{(0)} \in \mathbb{R}^{n_0 \times d}$ together with a vector $\mathbf{b}^{(0)} \in \mathbb{R}^{n_0}$. At the t -th step, a new data of form $((\mathbf{a}^{(t)})^\top, \beta^{(t)}) \in \mathbb{R}^d \times \mathbb{R}$ arrives, and the goal is to maintain an ϵ -approximate solution. A formal description is provided below, where we assume $n_0 = d + 1$ for simplicity (see Remark 2.3).

Definition 2.1 (Dynamic least-squares regression). *Let $d \in \mathbb{N}_+$ and $\epsilon \in [0, 1)$ be two fixed parameters. We say an algorithm solves ϵ -approximate dynamic least squares regression if*

- *The data structure is given a matrix $\mathbf{A}^{(0)} \in \mathbb{R}^{(d+1) \times d}$ and a vector $\mathbf{b}^{(0)} \in \mathbb{R}^{d+1}$ in the preprocessing phase.*
- *For each iteration $t \in [T]$, the algorithm receives updates $\mathbf{a}^{(t)} \in \mathbb{R}^d$ and $\beta^{(t)} \in \mathbb{R}$. Define $\mathbf{A}^{(t)} := [(\mathbf{A}^{(t-1)})^\top, \mathbf{a}^{(t)}]^\top \in \mathbb{R}^{(d+t+1) \times d}$ to be $\mathbf{A}^{(t-1)}$ appended with a new row $(\mathbf{a}^{(t)})^\top$, and $\mathbf{b}^{(t)} := [(\mathbf{b}^{(t-1)})^\top, \beta^{(t)}]^\top \in \mathbb{R}^{d+t+1}$ to be $\mathbf{b}^{(t-1)}$ appended with a new entry $\beta^{(t)}$. After this update, the algorithm outputs an ϵ -approximate solution $\mathbf{x}^{(t)} \in \mathbb{R}^d$:*

$$\|\mathbf{A}^{(t)} \mathbf{x}^{(t)} - \mathbf{b}^{(t)}\|_2 \leq (1 + \epsilon) \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}^{(t)} \mathbf{x} - \mathbf{b}^{(t)}\|_2.$$

We write $[0 : T] = \{0, 1, \dots, T\}$, and for any $t \in [0 : T]$, we denote $\mathbf{M}^{(t)} := [\mathbf{A}^{(t)}, \mathbf{b}^{(t)}] \in \mathbb{R}^{(d+t+1) \times (d+1)}$. We make the following assumptions.

Assumption 2.2. We assume 1. Each data have bounded ℓ_2 norm, i.e., $\forall i \in [T + d + 1]$, the i -th row of $\mathbf{M}^{(T)}$ satisfies $\|\mathbf{M}_{i,*}^{(T)}\|_2 \leq D$. 2. The initial matrix $\mathbf{M}^{(0)}$ has full rank, and its smallest singular value is bounded by $\sigma_{d+1}(\mathbf{M}^{(0)}) \geq \sigma$ for some polynomially small $\sigma \in (0, 1)$.

Remark 2.3. We remark that these assumptions are essentially w.l.o.g. for the following reasons: 1. Real world data inherently have bounded ℓ_2 norm, and in applications like machine learning, data are often normalized. 2. We can assume the initial matrix $\mathbf{A}^{(0)}$ has $d + 1$ rows because brute-forcely adding these $d + 1$ initial rows would only take $O(d^3)$ time, and this is within our desired total running time of $O_\epsilon(\text{nnz}(\mathbf{A}^{(T)}) + d^3)$. 3. To satisfy the assumption that $\sigma_{d+1}(\mathbf{M}^{(0)}) \geq \sigma$ for some polynomially small σ , we could let the initial matrix $\mathbf{M}^{(0)} = \sigma \cdot \mathbf{I}_{d+1}$. This is equivalent to adding a small regularization term of $\sigma \cdot \|\mathbf{x}\|_2$ and this incurs only a polynomially small additive error.

3 DYNAMIC ϵ -LSR DATA STRUCTURE

In this section, we provide an approximation algorithm for the dynamic least squares regression. Notably, our algorithm maintains an ϵ -approximate solution in near input sparsity time.

Theorem 3.1 (Data structure for dynamic least squares regression). *Let $\epsilon > 0$, $d, T \in \mathbb{N}$. There exists a randomized algorithm for dynamic least-squares regression (Algorithm 1–4). With probability at least 0.9, the algorithm maintains an ϵ -approximation solution for all iterations $t \in [T]$ and the total update time over T iterations is at most $O(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(TD/\sigma))$. Our data structure uses at most $O(\epsilon^{-2} d^2 \cdot \log^2(TD/\sigma))$ space.*

Notations We use a superscript (t) to denote the matrix/vector/scalar maintained by the data structure at the end of the t -th iterations. In particular, the superscript (0) represents the variables after the preprocessing step. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, $i \in [n]$ we define its leverage score $\tau(\mathbf{A}) \in \mathbb{R}^n$ as $\tau_i(\mathbf{A}) := \mathbf{a}_i^\top (\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{a}_i$. We define the generalized leverage score (same as Cohen et al. (2015)) of $\mathbf{A} \in \mathbb{R}^{n \times d}$ with respect to another matrix $\mathbf{B} \in \mathbb{R}^{n' \times d}$ as: $\tau_i^{\mathbf{B}}(\mathbf{A}) := \mathbf{a}_i^\top (\mathbf{B}^\top \mathbf{B})^\dagger \mathbf{a}_i$. For more properties of the leverage scores see Section B.1.

3.1 TECHNIQUE OVERVIEW

Our approach is formally described in Algorithm 1–4, we first overview the ideas behind it.

From a high level view, our approach follows the online row sampling framework (Cohen et al., 2015; 2020; Braverman et al., 2020): When a new row arrives, we sample and keep the new row with probability proportional to the (approximated version of) *online leverage score*

$$\tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) = (\mathbf{m}^{(t)})^\top ((\mathbf{M}^{(t-1)})^\top \mathbf{M}^{(t-1)})^{-1} \mathbf{m}^{(t)}.$$

The sampled matrix is a spectral approximation to the true data matrix. We maintain an approximate least-squares regression solution using this sampled matrix.

Naively computing the online leverage score takes $O(d^2)$ time. In order to accelerate this computation, we use two approximations:

1. Similar to Cohen et al. (2020), we compute the online leverage scores with respect to the sampled matrix instead of the true data matrix. However, this idea alone is still not enough to achieve sub-quadratic time.
2. We use a JL-embedding³ trick (Spielman & Srivastava, 2011) to compress the size of the $d \times d$ matrix to $\approx \epsilon^{-2} \times d$. In this way, in each iteration it only takes $O_\epsilon(d)$ time to compute the approximate online leverage score.

We further use an inductive analysis to bound the overall error (Lemma 3.4).

Finally, we adopt a similar strategy as Cohen et al. (2020) to prove that sampling according to the approximate online leverage score still keeps at most $O_\epsilon(d)$ sampled rows (Lemma 3.7). Whenever a row is sampled, it takes $O(d^2)$ time to update the maintained matrices using Woodbury identity. Hence, the amortized update time of the sampled rows is $O_\epsilon(d^3/T) = o(d)$ when $T \gg d$.

³Johnson-Lindenstrauss (JL) Lemma shows a way to embed high-dimensional vectors to low-dimensional space while preserving the distances between the vectors. A rigorous statement is shown in Lemma B.4.

Remark 3.2 (Difference from sketching-based solutions). *Our approach crucially differs from the sketching-based solutions, which do not provide any speedup over the direct application of Woodbury identity ($O(d^2)$ time per iteration). A sketching-based solution maintains a sketched matrix $\mathbf{SM} \in \mathbb{R}^{O_\epsilon(d) \times d}$, where \mathbf{S} is a sketching matrix (e.g. SRHT (Ailon & Chazelle, 2006) or Count Sketch (Clarkson & Woodruff, 2017)) that mixes the rows of \mathbf{M} . When a new row of \mathbf{M} arrives, at least one row of the sketched matrix \mathbf{SM} needs to be updated, in contrast to our sampling-based approach where the sampled matrix is not updated in most of the iterations.*

Implementation We explain the detailed implementation of our algorithm. At the beginning of the t -th iteration, a sampling matrix $\mathbf{D}^{(t-1)}$ is derived based on the online leverage score, and the sub-sampled matrix $\mathbf{N}^{(t-1)} = \mathbf{D}^{(t-1)}\mathbf{M}^{(t-1)}$ maintains a spectral approximation on the column space of $\mathbf{M}^{(t-1)} = [\mathbf{A}^{(t-1)}, \mathbf{b}^{(t-1)}]$. Let $s^{(t-1)}$ denote the number of sampled rows. To obtain spectral approximation, we maintain the approximate covariance matrices $\mathbf{H}^{(t-1)} = ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1}$ and $\mathbf{B}^{(t-1)} = \mathbf{N}^{(t-1)}\mathbf{H}^{(t-1)}$. The online leverage score $\tau^{(t)}$ of a new row $\mathbf{m}^{(t)} = ((\mathbf{a}^{(t)})^\top, \beta^{(t)})$ can be approximated as $\|\mathbf{B}^{(t-1)}\mathbf{m}^{(t-1)}\|_2$. To efficiently compute the leverage score of a new row, we left-multiply by a JL matrix $\mathbf{J}^{(t-1)}$ and maintain a proxy $\tilde{\mathbf{B}}^{(t-1)} = \mathbf{J}^{(t-1)}\mathbf{B}^{(t-1)}$, with the guarantee that $\|\tilde{\mathbf{B}}^{(t-1)}\mathbf{m}\|_2 \approx \|\mathbf{B}^{(t-1)}\mathbf{m}\|_2$ for any $\mathbf{m} \in \mathbb{R}^{d+1}$ with high probability.

When a new row $\mathbf{m}^{(t)}$ is inserted at the t -th iteration, we sample it via the approximate online leverage score (Line 3 of SAMPLE). We only perform update if the new row is sampled. In that case, we renew the JL matrix and perform a series of careful updates on all the variables that we maintain (See UPDATEMEMBERS). To obtain the final solution $\mathbf{x}^{(t)} \in \mathbb{R}^d$, we solve $\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{D}^{(t)}\mathbf{A}^{(t)}\mathbf{x} - \mathbf{D}^{(t)}\mathbf{b}^{(t)}\|_2$, which has the closed-form solution of $\mathbf{x}^{(t)} = \mathbf{G}^{(t)} \cdot \mathbf{u}^{(t)}$ and can be efficiently maintained by taking $\mathbf{G}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \in \mathbb{R}^{d \times d}$ and $\mathbf{u}^{(t)} = (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)}$.

Algorithm 1 PREPROCESS ($\mathbf{A}, \mathbf{b}, \epsilon, T$)

```

1:  $\mathbf{M} \leftarrow [\mathbf{A}, \mathbf{b}]$   $\triangleright \mathbf{M} \in \mathbb{R}^{(d+1) \times (d+1)}$ 
2:  $\mathbf{D} \leftarrow \mathbf{I}_{d+1}$ 
   # Spectral approximation
3:  $s \leftarrow d + 1$ 
4:  $\mathbf{N} \leftarrow \mathbf{D} \cdot \mathbf{M}$   $\triangleright \mathbf{N} \in \mathbb{R}^{s \times (d+1)}$ 
5:  $\mathbf{H} \leftarrow ((\mathbf{N})^\top \mathbf{N})^{-1}$   $\triangleright \mathbf{H} \in \mathbb{R}^{(d+1) \times (d+1)}$ 
6:  $\mathbf{B} \leftarrow \mathbf{N} \cdot \mathbf{H}$   $\triangleright \mathbf{B} \in \mathbb{R}^{s \times (d+1)}$ 
   # JL approximation
7:  $\delta \leftarrow O(1/T^2), k \leftarrow O(\epsilon^{-2} \log(T/\delta))$ 
8:  $\mathbf{J} \leftarrow \text{JL}(s, \epsilon, \delta, T)$   $\triangleright$  JL matrix  $\mathbf{J} \in \mathbb{R}^{k \times s}$ 
9:  $\tilde{\mathbf{B}} \leftarrow \mathbf{J} \cdot \mathbf{B}$   $\triangleright \tilde{\mathbf{B}} \in \mathbb{R}^{k \times (d+1)}$ 
   # Maintain solution
10:  $\mathbf{G} \leftarrow (\mathbf{A}^\top \mathbf{D}^\top \mathbf{D} \mathbf{A})^{-1}$   $\triangleright \mathbf{G} \in \mathbb{R}^{d \times d}$ 
11:  $\mathbf{u} \leftarrow \mathbf{A}^\top \mathbf{D}^2 \mathbf{b}$   $\triangleright \mathbf{u} \in \mathbb{R}^d$ 
12:  $\mathbf{x} \leftarrow \mathbf{G} \cdot \mathbf{u}$   $\triangleright \mathbf{x} \in \mathbb{R}^d$ 

```

Algorithm 2 UPDATE (\mathbf{a}, β)

```

1:  $\mathbf{m} \leftarrow [\mathbf{a}^\top, \beta]^\top$   $\triangleright \mathbf{m} \in \mathbb{R}^{d+1}$ 
2:  $\nu \leftarrow \text{SAMPLE}(\mathbf{m})$   $\triangleright \nu \in \mathbb{R}$ 
3:  $\mathbf{D} \leftarrow \begin{bmatrix} \mathbf{D} & 0 \\ 0 & \nu \end{bmatrix}$ 
4: if  $\nu \neq 0$  then UPDATEMEMBERS( $\mathbf{m}$ )
5: return  $\mathbf{x}$ 

```

Algorithm 3 SAMPLE (\mathbf{m})

```

1:  $\tau \leftarrow \|\tilde{\mathbf{B}} \cdot \mathbf{m}\|_2^2$ 
2:  $p \leftarrow \min\{3(1 + \epsilon)^2 \epsilon^{-2} \tau \log(1/\delta), 1\}$ 
3:  $\nu \leftarrow 1/\sqrt{p}$  with probability  $p$ , and  $\nu \leftarrow 0$  otherwise

```

Algorithm 4 UPDATEMEMBERS (\mathbf{m})

```

   # Update spectral approximation
1:  $s \leftarrow s + 1$ 
2:  $\Delta \mathbf{H} \leftarrow -\frac{\mathbf{H} \mathbf{m} \mathbf{m}^\top \mathbf{H} / p}{1 + \mathbf{m}^\top \mathbf{H} \mathbf{m} / p}$ 
3:  $\mathbf{H} \leftarrow \mathbf{H} + \Delta \mathbf{H}$ 
4:  $\mathbf{B} \leftarrow [(\mathbf{B} + \mathbf{N} \cdot \Delta \mathbf{H})^\top, \mathbf{H} \cdot \mathbf{m} / \sqrt{p}]^\top$ 
5:  $\mathbf{N} \leftarrow [\mathbf{N}^\top, \mathbf{m} / \sqrt{p}]^\top$ 
   # Update JL approximation
6:  $\mathbf{J} \leftarrow \text{JL}(s, \epsilon, \delta, T)$ 
7:  $\tilde{\mathbf{B}} \leftarrow \mathbf{J} \cdot \mathbf{B}$ 
   # Update solution
8:  $\mathbf{G} \leftarrow \mathbf{G} - \frac{\mathbf{G} \mathbf{a} \mathbf{a}^\top \mathbf{G} / p}{1 + \mathbf{a}^\top \mathbf{G} \mathbf{a} / p}$ 
9:  $\mathbf{u} \leftarrow \mathbf{u} + \beta \cdot \mathbf{a} / p$ 
10:  $\mathbf{x} \leftarrow \mathbf{G} \cdot \mathbf{u}$ 

```

We outline the proof of Theorem 3.1, and defer the detailed proof to Appendix C due to space limits.

3.2 CORRECTNESS

We show the correctness of our algorithm and prove it maintains an ϵ -approximate solution for all iterations with high probability. We start with closed-form formulas for all the variables we maintain.

Lemma 3.3 (Closed-form formulas). *At the t -th iteration of Algorithm 1–4, we have*

1. $\mathbf{M}^{(t)} = [\mathbf{A}^{(t)}, \mathbf{b}^{(t)}] \in \mathbb{R}^{(d+t+1) \times (d+1)}$.
2. $\mathbf{D}^{(t)} \in \mathbb{R}^{(d+t+1) \times (d+t+1)}$ is a diagonal matrix with $s^{(t)}$ non-zero entries.
3. $\mathbf{N}^{(t)} = (\mathbf{D}^{(t)} \mathbf{M}^{(t)})_{S^{(t)},*} \in \mathbb{R}^{s^{(t)} \times (d+1)}$, where $S^{(t)} \subset [d+t+1]$ is defined as the set of non-zero entries of $\mathbf{D}^{(t)}$.
4. $\mathbf{H}^{(t)} = ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1} \in \mathbb{R}^{(d+1) \times (d+1)}$.
5. $\mathbf{B}^{(t)} = \mathbf{N}^{(t)} \mathbf{H}^{(t)} \in \mathbb{R}^{s^{(t)} \times (d+1)}$.
6. $\tilde{\mathbf{B}}^{(t)} = \mathbf{J}^{(t)} \cdot \mathbf{B}^{(t)} \in \mathbb{R}^{k \times (d+1)}$, where $k = O(\epsilon^{-2} \log(T/\delta))$.
7. $\mathbf{G}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \in \mathbb{R}^{d \times d}$.
8. $\mathbf{u}^{(t)} = (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)} \in \mathbb{R}^d$.
9. $\mathbf{x}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \cdot (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)} \in \mathbb{R}^d$.

The following lemma is key for our correctness analysis. It shows that we maintain a good approximation on online leverage scores and a spectral approximation of $\mathbf{M}^{(t)}$ throughout all iterations.

Lemma 3.4 (Spectral approximation via leverage score maintenance). *With probability at least $1 - 2T\delta$,*

$$(1 - \epsilon)^2 \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \tau^{(t)} \leq (1 + \epsilon)^2 \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}), \quad \forall t \in [T], \quad (2)$$

and

$$(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)} \approx_\epsilon (\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}, \quad \forall t \in [0 : T]. \quad (3)$$

Proof Sketch. We prove by induction and show that with probability $1 - 2t\delta$, Eq. (3) holds for all $t' \in [0 : t]$ and Eq. (2) holds for all $t' \in [t]$. The base case $t = 0$ holds trivially, as $\mathbf{D}^{(0)} = \mathbf{I}_{d+1}$, and therefore, $(\mathbf{M}^{(0)})^\top (\mathbf{D}^{(0)})^2 \mathbf{M}^{(0)} = (\mathbf{M}^{(0)})^\top \mathbf{M}^{(0)}$. Given the induction hypothesis upon $t - 1$, we proceed in the following three steps.

- We first use the induction hypothesis to prove that $\|\mathbf{B}^{(t-1)} \mathbf{m}^{(t)}\|$ is a good estimate on the online leverage score, that is

$$(1 - \epsilon) \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \|\mathbf{B}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq (1 + \epsilon) \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}).$$

- We then use the JL lemma (Lemma B.4) to show that with probability $1 - \delta$, the sketched covariance matrix $\|\tilde{\mathbf{B}}^{(t-1)} \mathbf{m}^{(t)}\|$ returns good estimation $\|\mathbf{B}^{(t-1)} \mathbf{m}^{(t)}\|$. That is

$$(1 - \epsilon)^2 \cdot \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \|\tilde{\mathbf{B}}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq (1 + \epsilon)^2 \cdot \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}). \quad (4)$$

- Finally, we wrap up the proof by proving the second part of induction. In particular, we show that conditioned on Eq. (4) holds, we have

$$(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)} \approx_\epsilon (\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}.$$

The proof then follows by an union bound over failure events. \square

It is well known that spectral approximations of $(\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}$ give approximate solutions to least squares regressions (Woodruff, 2014), so we have proved the correctness of our algorithm.

Lemma 3.5 (Correctness of Algorithm 1–4). *With probability at least $1 - O(1/T)$, in each iteration, UPDATE of Algorithm 2 outputs a vector $\mathbf{x}^{(t)} \in \mathbb{R}^d$ such that*

$$\|\mathbf{A}^{(t)} \mathbf{x}^{(t)} - \mathbf{b}^{(t)}\|_2 \leq (1 + \epsilon) \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}^{(t)} \mathbf{x} - \mathbf{b}^{(t)}\|_2.$$

3.3 TIME ANALYSIS

Next, we bound the overall update time of our algorithm. We first compute the worst case update time of Algorithm 2. When $\nu^{(t)} = 0$, i.e., the t -th row is not sampled, the UPDATE procedure only needs to compute the approximate leverage score $\tau^{(t)}$ (SAMPLE, Algorithm 3), and it takes $O(k \cdot \text{nnz}(\mathbf{m}^{(t)}))$ time. When $\nu^{(t)} \neq 0$, i.e., the t -th row is sampled, the UPDATE procedure makes a call to UPDATEMEMBERS (Algorithm 4), and it takes $O(k \cdot s^{(t)} \cdot d)$ time. Plugging in the value of k , we have the following lemma.

Lemma 3.6 (Worst case update time). *At the t -th iteration of the UPDATE procedure (Algorithm 2),*

- *If $\nu^{(t)} = 0$, then UPDATE takes $O(\epsilon^{-2} \log(T/\delta) \cdot \text{nnz}(\mathbf{a}^{(t)}))$ time.*
- *If $\nu^{(t)} \neq 0$, then UPDATE takes $O(\epsilon^{-2} s^{(t)} d \log(T/\delta))$ time.*

To bound the amortized update time, we need to bound the total number of sampled rows, and this is closely related to the sum of online leverage scores. Such an upper bound was already established by Cohen et al. (2020), here we present a slightly generalized version of it.

Lemma 3.7 (Sum of online leverage scores, generalization of Theorem 2.2 of (Cohen et al., 2020)). *If the matrix $\mathbf{M}^{(T)}$ satisfy Assumption 2.2, then*

$$\sum_{t=1}^T \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq O(d \log(TD/\sigma)).$$

Now we are ready to bound the amortized update time of our algorithm.

Lemma 3.8 (Amortized update time). *With probability at least 0.99, the total running time of UPDATE over T iterations is at most $O(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(TD/\sigma))$.*

Proof Sketch. In this proof sketch we simplify the second term as $d^3 \cdot \text{poly}(\epsilon^{-1} \log(TD/\sigma))$. The first term comes from the computation cost of querying leverage score, which takes $O(\epsilon^{-2} \log(T/\delta) \cdot \text{nnz}(\mathbf{a}^{(t)}))$ time in the t -th iteration even if the t -th row is not sampled. The second term bounds the total update time for the sampled rows:

- From Lemma 3.4 and Lemma 3.7, with high probability the sum of the approximate online leverage scores $\tau^{(t)}$ are bounded by $O(d \log(TD/\sigma))$.
- Using Markov inequality, the total number of sampled rows is bounded by

$$s^{(T)} = O\left(\sum_{i=1}^T p^{(i)}\right) = O(\epsilon^{-2} \cdot \log(1/\delta) \cdot \sum_{t=1}^T \tau^{(t)}) \leq O(d \cdot \text{poly}(\epsilon^{-1} \log(TD/\sigma))).$$

- Since there are $s^{(T)}$ sampled rows, and for each sampled row we update data structure members in $\leq O(s^{(T)} d \cdot \text{poly}(\epsilon^{-1} \log(TD/\sigma)))$ time, the total update time for sampled rows is

$$s^{(T)} \cdot O(s^{(T)} d \cdot \text{poly}(\epsilon^{-1} \log(TD/\sigma))) = O(d^3 \cdot \text{poly}(\epsilon^{-1} \log(TD/\sigma))). \quad \square$$

4 HARDNESS RESULT

We prove a $\Omega(d^{2-o(1)})$ amortized time lower bound for dynamic least squares regression with high precision, assuming the OMv conjecture. The OMv conjecture was originally proposed by Henzinger et al. (2015), and it is widely accepted in the theoretical computer science community.

Conjecture 4.1 (OMv conjecture, (Henzinger et al., 2015)). *Let $d \in \mathbb{N}, T = \text{poly}(d)$. Let $\gamma > 0$ be any constant. $\mathbf{B} \in \{0, 1\}^{d \times d}$ is a Boolean matrix. $\forall t \in [T]$, a Boolean vector $\mathbf{z}^{(t)} \in \{0, 1\}^d$ is revealed at the t -th step. We say an algorithm solves the OMv problem if it returns the Boolean matrix-vector product $\mathbf{Bz}^{(t)} \in \mathbb{R}^d$ at every time step. The conjecture states that there is no algorithm that solves the OMv problem using $\text{poly}(d)$ preprocessing time and $O(d^{2-\gamma})$ amortized running time, and has an error probability $\leq 1/3$.*

The results in this section are all under the Word RAM model where the word size $w = O(\log d)$. Our main result is formally stated below.

Theorem 4.2 (Hardness of dynamic-least squares regression with high precision). *Let $d \in \mathbb{N}$, $T = \text{poly}(d)$, $\epsilon = \frac{1}{d^{8T^2}} = 1/\text{poly}(d)$, and let $\gamma > 0$ be any constant. Assuming the OMv conjecture is true, any dynamic algorithm that maintains an ϵ -approximate solution of the least squares regression requires at least $\Omega(d^{2-\gamma})$ amortized time per update.*

Our lower bound is proved by first reducing the standard OMv conjecture for Boolean matrices to OMv-hardness for well-conditioned positive semidefinite (PSD) matrices over real numbers. Then we use this new OMv-hardness result to prove our lower bound for dynamic least squares regression. We only provide a proof sketch here and detailed proof are delayed to Section D.

OMv-hardness for well-conditioned PSD matrix The OMv conjecture asserts the hardness of solving online Boolean matrix-vector product *exactly*. We extend it to solving online real-valued matrix-vector product for well-conditioned PSD matrices, while allowing polynomially small error.

Lemma 4.3 (Hardness of approximate real-valued OMv). *Let $d \in \mathbb{N}$, $T = \text{poly}(d)$. Let $\gamma > 0$ be any constant. Let $\mathbf{H} \in \mathbb{R}^{d \times d}$ be a symmetric matrix whose eigenvalues satisfy $1 \leq \lambda_d(\mathbf{H}) \leq \dots \leq \lambda_1(\mathbf{H}) \leq 3$. For any $t \in [T]$, $\mathbf{z}^{(t)} \in \mathbb{R}^d$ is revealed at the t -th step, and $\|\mathbf{z}^{(t)}\|_2 \leq 1$. Assuming the OMv conjecture is true, then there is no algorithm with $\text{poly}(d)$ preprocessing time and $O(d^{2-\gamma})$ amortized running time that can return an $O(1/d^2)$ -approximate answer to $\mathbf{H}\mathbf{z}^{(t)}$ for all t , i.e., a vector $\mathbf{y}^{(t)} \in \mathbb{R}^d$ s.t. $\|\mathbf{y}^{(t)} - \mathbf{H}\mathbf{z}^{(t)}\|_2 \leq \epsilon$, and has an error probability $\leq 1/3$.*

Proof Sketch. Given a Boolean matrix $\mathbf{B} \in \{0, 1\}^{d \times d}$ in the OMv conjecture, we construct a PSD matrix $\mathbf{H} = \begin{bmatrix} 2\mathbf{I}_d & \frac{1}{d}\mathbf{B} \\ \frac{1}{d}\mathbf{B}^\top & 2\mathbf{I}_d \end{bmatrix} \in \mathbb{R}^{2d \times 2d}$. We note that \mathbf{H} is symmetric and $1 \leq \lambda_d(\mathbf{H}) \leq \lambda_1(\mathbf{H}) \leq 3$. Given a binary OMv query vector $\mathbf{z}^{(t)}$, we construct $\bar{\mathbf{z}}^{(t)} = (\mathbf{0}_d, \mathbf{z}^{(t)}) \in \mathbb{R}^{2d}$. Since \mathbf{H} has a constant condition number, we can prove that rounding an $\epsilon \sim 1/d^2$ -approximate answer $\hat{\mathbf{y}} \approx_\epsilon \mathbf{H} \cdot \bar{\mathbf{z}}^{(t)}$ still gives the correct binary answer to $\mathbf{B}\mathbf{z}^{(t)}$. \square

Reducing OMv to dynamic least-squares regression We next wrap up the proof of Theorem 4.2 by reducing OMv to dynamic ϵ -LSR.

Proof Sketch of Theorem 4.2. Given a PSD matrix H and a sequence of query $\{\mathbf{H}\mathbf{z}^{(t)}\}_{t=1}^T$ of the problem in Lemma 4.3, we reduce it to a dynamic ϵ -LSR, where the initial \mathbf{A} is such that $\mathbf{A}^\top \mathbf{A} = \mathbf{H}^{-1}$ (this preprocessing step of the reduction takes $\sim d^\omega$ time), and the label is 0 for the initial d data. For each $t \in [T]$, the incoming row $\mathbf{a}^{(t)}$ is a small scaled version of $\mathbf{z}^{(t)}$, i.e., $\mathbf{a}^{(t)} = \frac{1}{d^2\sqrt{T}} \cdot \mathbf{z}^{(t)} \in \mathbb{R}^d$, and the label is 1. Let $\mathbf{x}_*^{(t)}$ be the optimal solution at the t -th step, $\mathbf{H}^{(t)} := ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-2}$, the reduction is complete via the following three steps:

- Step 1. $\mathbf{x}^{(t)}$ and $\mathbf{x}_*^{(t)}$ are close, i.e., $\mathbf{x}^{(t)} = \mathbf{x}_*^{(t)} \pm O(\frac{1}{d^4\sqrt{T}})$.
- Step 2. $\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}$ recovers $\mathbf{H}^{(t-1)}\mathbf{a}^{(t)}$, i.e., $\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)} = \mathbf{H}^{(t-1)}\mathbf{a}^{(t)} \pm O(\frac{1}{d^4\sqrt{T}})$.
- Step 3. $\mathbf{H}^{(t-1)}\mathbf{a}^{(t)}$ is close to $\mathbf{H}\mathbf{a}^{(t)}$, i.e., $\mathbf{H}^{(t-1)}\mathbf{a}^{(t)} = \mathbf{H}\mathbf{a}^{(t)} \pm O(\frac{1}{d^6\sqrt{T}})$.

In particular, let $\mathbf{y}^{(t)} = d^2\sqrt{T}(\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)})$, Step 2 and 3 directly implies $\|\mathbf{y} - \mathbf{H}\mathbf{z}^{(t)}\|_2 \leq O(1/d^2)$. This completes the proof. \square

5 EXPERIMENTS

Our method is most suitable for data distributions that are non-uniform. Indeed, if the data has low coherence (they are all similar to each other), then the naive uniform sampling is as good as leverage score sampling. We perform empirical evaluations on our algorithm over both synthetic and real-world datasets.

Synthetic dataset We follow the empirical study of (Dobriban & Liu, 2019) and generate data from the *elliptical model*. In this model $\mathbf{a}^{(t)} = w^{(t)}\Sigma\mathbf{z}^{(t)}$, where $\mathbf{z}^{(t)} \sim N(0, \mathbf{I}_d)$ is a random Gaussian vector, $\Sigma \in \mathbb{R}^{d \times d}$ is a PSD matrix, and $w^{(t)}$ is a scalar. The label is generated as $b^{(t)} =$

$\langle \mathbf{a}^{(t)}, \mathbf{x}^* \rangle + w^{(t)} \xi^{(t)}$, where $\mathbf{x}^* \in \mathbb{R}^d$ is a hidden vector and $\xi \sim N(0, 1)$ is standard Gaussian noise. This model has a long history in multivariate statistics, see e.g. (Martin & Maes, 1979). In our experiments, we set $\Sigma = \mathbf{I}_d$ for simplicity. In order to make the dataset non-trivial, we set $w^{(t)}$ to be large ($= \sqrt{T}$) for a few ($= d/10$) iterations, and small ($= 1$) for the rest of the iterations. We set $T = 400000$ and $d = 500$.

Real-world dataset We use the VirusShare dataset from the UCI Machine Learning Repository⁴. We select this dataset because it has a large number of features and data points, and has low errors when fitted by a linear model. The dataset is collected from Nov 2010 to Jul 2014 by VirusShare (an online platform for malware detection). It has $T = 107888$ data points and $d = 482$ features.

Baseline algorithms We compare with three baseline methods. 1. *Kalman’s approach* makes use of Woodbury identity and gives an exact solution. 2. The *uniform sampling* approach samples new rows uniformly at random. 3. The *row sampling* approach samples new rows according to the exact online leverage scores. (Cohen et al., 2020)

Our experiments are executed on an Apple M1 CPU with codes written in MATLAB. We repeat all experiments for at least 5 times and take the mean. On both datasets, we initiate the model based on the first 10% of the data. The experiment results are formally presented in Figure 1 and more details can be found in Appendix E. Our algorithm consistently outperforms baseline methods: Our algorithm runs faster when achieving comparable error rates.

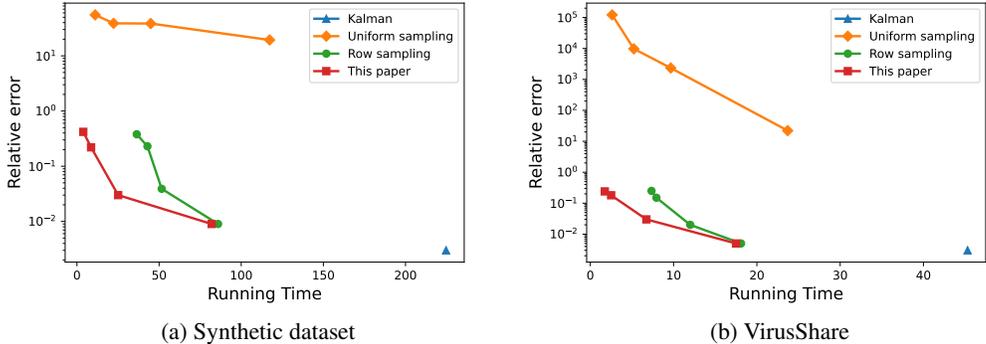


Figure 1: Experiment results. The x -axis shows the running time (unit: seconds), and the y -axis shows the relative error $(\text{err}/\text{err}_{\text{std}} - 1)$, where err is the error of the particular approach, and err_{std} is the error of the static Normal equation. The y -axis is on a log scale. For uniform sampling, we take sampling probability $p = 0.05, 0.1, 0.2, 0.5$. For row sampling and our algorithm, we take the error parameter $\epsilon = 0.1, 0.2, 0.5, 1$. Kalman’s approach has a relative error of 0.

6 CONCLUSION

We provide the first practical and *provably* fast data structure for dynamic least-squares regression, obtaining nearly tight upper and lower bounds for this fundamental problem. On the algorithmic side, we design an ϵ -approximation dynamic algorithm whose total update time almost matches the *input sparsity* of the (online) matrix. On the lower bound side, we prove that it is impossible to maintain an exact (or even high-accuracy) solution with $\ll d^{2-o(1)}$ amortized update time under the OMv conjecture. As such, this result exhibits the first separation between the static and the dynamic LSR problems.

Our paper sets forth several interesting future directions. On the theoretical side, a very interesting question is whether it is possible to reduce the additive term d^3 of our algorithm to matrix-multiplication time d^ω ? A second open problem—of interest in both theory and practice—is whether it is possible to achieve input-sparsity amortized update time in the *fully dynamic* setting, i.e., when allowing both addition and *deletion* of data rows? Finally, it would be interesting to find connections between dynamic least-squares regression and incremental training of more complicated models, such as dynamic Kernel-ridge regression and to deep neural networks.

⁴<https://archive.ics.uci.edu/ml/datasets.php>

REFERENCES

- Thomas D Ahle, Michael Kapralov, Jakob BT Knudsen, Rasmus Pagh, Ameya Velingker, David P Woodruff, and Amir Zandieh. Oblivious sketching of high-degree polynomial kernels. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 141–160. SIAM, 2020.
- Nir Ailon and Bernard Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp. 557–563, 2006.
- Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 522–539. SIAM, 2021.
- Haim Avron, Kenneth L Clarkson, and David P Woodruff. Faster kernel ridge regression using sketching and preconditioning. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1116–1138, 2017.
- Christos Boutsidis, David P Woodruff, and Peilin Zhong. Optimal principal component analysis in distributed and streaming models. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 236–249, 2016.
- van den Jan Brand, Yin-Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Bipartite matching in nearly-linear time on moderately dense graphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 919–930. IEEE, 2020a.
- van den Jan Brand, Yin Tat Lee, Aaron Sidford, and Zhao Song. Solving tall dense linear programs in nearly linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 775–788, 2020b.
- van den Jan Brand, Yin Tat Lee, Yang P Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Minimum cost flows, mdps, and ℓ_1 -regression in nearly linear time for dense instances. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 859–869, 2021a.
- van den Jan Brand, Binghui Peng, Zhao Song, and Omri Weinstein. Training (overparametrized) neural networks in near-linear time. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, 2021b.
- Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 517–528. IEEE, 2020.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- Nadiia Chepurko, Kenneth L Clarkson, Praneeth Kacham, and David P Woodruff. Near-optimal algorithms for linear algebra in the current matrix multiplication time. *arXiv preprint arXiv:2107.08090*, 2021.
- Charles K. Chui. Estimation, control, and the discrete kalman filter (donald e. calin). *SIAM Review*, 32(3):493–494, 1990. doi: 10.1137/1032097. URL <https://doi.org/10.1137/1032097>.
- Kenneth L Clarkson and David P Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 205–214, 2009.
- Kenneth L Clarkson and David P Woodruff. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)*, 63(6):1–45, 2017.

- Michael B Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 181–190. ACM, 2015.
- Michael B Cohen, Cameron Musco, and Jakub Pachocki. Online row sampling. *Theory OF Computing*, 16(15):1–25, 2020.
- Edgar Dobriban and Sifan Liu. Asymptotics for sketching in least squares. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 3675–3685, 2019.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Sampling algorithms for l2 regression and applications. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm (SODA)*, pp. 1127–1136, 2006a.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-based methods. In *Approximation, Randomization, and Combinatorial Optimization. (APPROX-RANDOM)*, pp. 316–326, 2006b.
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-row-based methods. In *European Symposium on Algorithms (ESA)*, pp. 304–314, 2006c.
- Trevor Hastie, Jerome H. Friedman, and Robert Tibshirani. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer, 2001. ISBN 978-1-4899-0519-2. doi: 10.1007/978-0-387-21606-5. URL <https://doi.org/10.1007/978-0-387-21606-5>.
- Elad Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.
- Monika Henzinger, Sebastian Krininger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 21–30, 2015.
- Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. A faster algorithm for solving general lps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 823–832, 2021.
- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pp. 296–303, 2014.
- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 424–433. IEEE, 2014.
- Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix multiplication time. In *Annual Conference on Learning Theory (COLT)*, 2019.
- Edo Liberty, Zohar Karnin, Bing Xiang, Laurence Rouesnel, Baris Coskun, Ramesh Nallapati, Julio Delgado, Amir Sadoughi, Yury Astashonok, Piali Das, et al. Elastic machine learning algorithms in amazon sagemaker. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pp. 731–737, 2020.
- Aleksander Madry. Navigating central path with electrical flows: From flows to matchings, and back. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 253–262. IEEE, 2013.
- Nick Martin and Hermine Maes. *Multivariate analysis*. Academic press London, 1979.

- Xiangrui Meng and Michael W Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 91–100, 2013.
- Jelani Nelson and Huy L Nguyễn. Osnap: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *2013 IEEE 54th annual symposium on foundations of computer science*, pp. 117–126. IEEE, 2013.
- Rasmus Pagh. Compressed matrix multiplication. *ACM Transactions on Computation Theory (TOCT)*, 5(3):1–17, 2013.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- R. L. Plackett. Some theorems in least squares. *Biometrika*, 37(1/2):149–157, 1950. ISSN 00063444. URL <http://www.jstor.org/stable/2332158>.
- Lawrence R Rabiner and Bernard Gold. Theory and application of digital signal processing. *Englewood Cliffs: Prentice-Hall*, 1975.
- Ilya Razenshteyn, Zhao Song, and David P Woodruff. Weighted low rank approximations with provable guarantees. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 250–263, 2016.
- Vladimir Rokhlin and Mark Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. *Proceedings of the National Academy of Sciences*, 105(36):13212–13217, 2008.
- Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pp. 143–152. IEEE, 2006.
- Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pp. 81–90, 2004.
- Stephen M. Stigler. Gauss and the Invention of Least Squares. *The Annals of Statistics*, 9(3): 465–474, 1981. doi: 10.1214/aos/1176345451. URL <https://doi.org/10.1214/aos/1176345451>.
- Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.
- David Woodruff. A very sketchy talk (invited talk). In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2021.
- David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014.

Roadmap We present additional related work in Section A, and describe notations and prerequisite results in Section B. We provide missing proof for our algorithm in Section C and Section D is devoted to prove the lower bound. We include detailed experiment results in Section E.

A ADDITIONAL RELATED WORK

Sketching and sampling are essential tools for speeding up computation costs of regression problems, we give a concise overview on the literature

Sketching The celebrated *sketch-and-solve* paradigm is developed to speed up the computation of numerical linear algebra tasks, including regressions (Ailon & Chazelle, 2006; Clarkson & Woodruff, 2017; Meng & Mahoney, 2013; Nelson & Nguyễn, 2013; Avron et al., 2017; Pagh, 2013; Ahle et al., 2020; Woodruff, 2014; 2021), low rank approximation (Clarkson & Woodruff, 2017; Boutsidis et al., 2016; Razenshteyn et al., 2016; Woodruff, 2014), and more generally, sketching techniques have been useful in optimizations (Lee et al., 2019; Jiang et al., 2021) and machines learning applications (Ahle et al., 2020; Brand et al., 2021b). The most direct application of the sketch-and-solve paradigm is the (overconstrained) least squares regression problem, which aims to solve $\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$, where $\mathbf{A} \in \mathbb{R}^{n \times d}$, $\mathbf{b} \in \mathbb{R}^n$ with $n \gg d$. The sketch-and-solve paradigm tackles the problem by first reducing the dimension via a subspace embedding S and then solves $\min_{\mathbf{x} \in \mathbb{R}^{n \times d}} \|\mathbf{S}\mathbf{A}\mathbf{x} - \mathbf{S}\mathbf{b}\|_2^2$ in a black box way. The subspace embedding and the reduced problem can be solved in $\tilde{O}(nd + \epsilon^{-2}d^\omega)$ time via the fast JL sketch (Ailon & Chazelle, 2006) and $\tilde{O}(\epsilon^{-1} \cdot \text{nnz}(\mathbf{A}) + \epsilon^{-2}d^\omega)$ time via the count sketch (Clarkson & Woodruff, 2017; Meng & Mahoney, 2013; Nelson & Nguyễn, 2013). The dependence on ϵ can be improved via preconditioning (Spielman & Teng, 2004; Rokhlin & Tygert, 2008; Clarkson & Woodruff, 2017; Avron et al., 2017; Brand et al., 2021b): Instead of solving the regression via a black box sketch-and-solve reduction, one can first precondition the matrix \mathbf{A} and turn it into a well-conditioned matrix, then run gradient descent or conjugate gradient descent, resulting in a total running time of $\tilde{O}(\text{nnz}(\mathbf{A}) \log(1/\epsilon) + d^\omega)$.

Leverage score sampling An alternative way to achieve subspace embedding (equivalently spectral approximation) is row sampling according to leverage scores (Drineas et al., 2006a;b;c; Spielman & Srivastava, 2011). When the matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is the Laplacian matrix of a graph, the leverage scores are equivalent to the effective resistances of the graph, and they can be approximately computed in $\tilde{O}(n)$ time using efficient Laplacian solvers and JL sketch Spielman & Srivastava (2011). For general matrices $\mathbf{A} \in \mathbb{R}^{n \times d}$, using iterative uniform row sampling and JL sketch, the leverage scores can be approximately computed in $\tilde{O}(\text{nnz}(\mathbf{A}) + d^\omega)$ time (Cohen et al., 2015). Dynamic maintenance of leverage scores has been studied by (Brand et al., 2020b;a; 2021a) in their specific applications of optimization problems.

Online linear regression The online linear regression, or more general, the online convex optimization (OCO) problem (Hazan, 2019), aims to minimize a sequence of convex functions f_1, \dots, f_T in an online stream. The performance is measured in terms of *regret*, defined as $\sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x}^* \in \mathcal{D}} \sum_{t=1}^T f_t(\mathbf{x}^*)$, where \mathcal{D} is the feasible region and \mathbf{x}_t is the output in the t -th iteration. The OCO is fundamentally different from the dynamic counterpart, and they are incomparable in general. In OCO, the loss function f_t is equipped with different output \mathbf{x}_t , whereas the dynamic problems seeks for a single solution $\mathbf{x}^{(T)}$ that achieves low error for all f_1, \dots, f_T . The difficulty of OCO are mainly *information-theoretic*, as the function f_t is revealed after fixing the choice of \mathbf{x}_t , whereas the challenge in our loss-minimization problem is purely *computational*.

B PRELIMINARY

Notations For an integer $n > 0$, let $[n] = \{1, 2, \dots, n\}$ and $[n_1 : n_2] = \{n_1, \dots, n_2\}$. Let $\mathbf{1}_n$ (resp. $\mathbf{0}_n$) denote the all one (resp. all zero) vectors of length n , and let \mathbf{I}_n denote the identity matrix of size $n \times n$.

Consider a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. We use $\text{nnz}(\mathbf{A})$ to denote the number of non-zero entries of a matrix \mathbf{A} . We use $\det(\mathbf{A})$ to denote the determinant of a matrix \mathbf{A} . We use $\ker[\mathbf{A}]$ and $\text{Im}[\mathbf{A}]$ to denote the kernel space and the column space of \mathbf{A} . We use $\mathbf{A}_{i,*}$ and $\mathbf{A}_{*,j}$ to denote the i -th row and the

j -th column of \mathbf{A} . For two sets $S \subseteq [m], T \subseteq [n]$, we use $\mathbf{A}_{S,*}$ and $\mathbf{A}_{*,T}$ to denote the submatrix of \mathbf{A} obtained by taking the rows in S or taking the columns in T .

We use \odot to denote the coordinate-wise multiplication of matrices: For two matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, $\mathbf{A} \odot \mathbf{B} = \sum_{i=1}^m \sum_{j=1}^n \mathbf{A}_{i,j} \cdot \mathbf{B}_{i,j}$.

For a vector $\mathbf{v} \in \mathbb{R}^n$, denote $\text{diag}(\mathbf{v}) \in \mathbb{R}^{n \times n}$ as a diagonal matrix whose diagonal entries are \mathbf{v} ; for a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, denote $\text{diag}(\mathbf{A}) \in \mathbb{R}^n$ as the vector of the diagonal entries of \mathbf{A} .

For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ ($n > d$), let $\sigma(\mathbf{A}) = (\sigma_1(\mathbf{A}), \dots, \sigma_d(\mathbf{A})) \in \mathbb{R}^d$ be the singular values of \mathbf{A} , and $\sigma_1(\mathbf{A}) \geq \dots \geq \sigma_d(\mathbf{A}) \geq 0$. The condition number of \mathbf{A} is $\kappa(\mathbf{A}) = \frac{\sigma_1(\mathbf{A})}{\sigma_d(\mathbf{A})}$. The pseudo-inverse of \mathbf{A} is denoted as $\mathbf{A}^\dagger \in \mathbb{R}^{d \times n}$. When $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric matrix, we use $\lambda(\mathbf{A}) = (\lambda_1(\mathbf{A}), \dots, \lambda_d(\mathbf{A})) \in \mathbb{R}^n$ to denote the eigenvalues of \mathbf{A} , and $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$.

The spectral norm of \mathbf{A} is defined as $\|\mathbf{A}\| = \max_{\mathbf{x} \in \mathbb{R}^d, \|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$. The Frobenius norm of \mathbf{A} is defined as $\|\mathbf{A}\|_F = \sqrt{\sum_{i \in [n]} \sum_{j \in [d]} \mathbf{A}_{i,j}^2}$.

B.1 BASICS ON NUMERICAL LINEAR ALGEBRA

PSD matrix A positive semidefinite (PSD) matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric and satisfies $\mathbf{x}^\top \mathbf{A} \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. We write $\mathbf{A} \succeq 0$ to denote that \mathbf{A} is PSD, and we write $\mathbf{B} \succeq \mathbf{A}$ to denote that $\mathbf{B} - \mathbf{A}$ is PSD.

Spectral approximation For two symmetric matrices $\mathbf{A}, \tilde{\mathbf{A}} \in \mathbb{R}^{n \times n}$, we say that $\tilde{\mathbf{A}}$ and \mathbf{A} are ϵ -spectral approximations of each other (denoted as $\tilde{\mathbf{A}} \approx_\epsilon \mathbf{A}$) if

$$(1 - \epsilon) \cdot \mathbf{A} \preceq \tilde{\mathbf{A}} \preceq (1 + \epsilon) \cdot \mathbf{A}.$$

Leverage scores For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we define its leverage score $\tau(\mathbf{A}) \in \mathbb{R}^n$ as follows: $\forall i \in [n]$, let \mathbf{a}_i^\top be the i -th row of \mathbf{A} , then

$$\tau_i(\mathbf{A}) := \mathbf{a}_i^\top (\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{a}_i.$$

The summation of all leverage scores equals to the rank of the matrix \mathbf{A} . The leverage score of a row measures its importance with respect to other rows. If a row is orthogonal to all other rows, it has a large leverage score of 1. If all rows are the same, they all have the same leverage score of $\frac{1}{n}$.

We also define a generalized leverage score (same as Cohen et al. (2015)) of $\mathbf{A} \in \mathbb{R}^{n \times d}$ with respect to another matrix $\mathbf{B} \in \mathbb{R}^{n' \times d}$ as follows:

$$\tau_i^{\mathbf{B}}(\mathbf{A}) := \mathbf{a}_i^\top (\mathbf{B}^\top \mathbf{B})^\dagger \mathbf{a}_i, \quad \forall i \in [n].$$

The following fact shows that the leverage score of a row w.r.t. a submatrix is larger than the leverage score of the original matrix.

Fact B.1 (Leverage score with respect to submatrix). *Let $\mathbf{A} \in \mathbb{R}^{n_1 \times d}$ and $\mathbf{B} \in \mathbb{R}^{n_2 \times d}$ be two matrices, and let $\mathbf{M} := [\mathbf{A}^\top, \mathbf{B}^\top]^\top \in \mathbb{R}^{(n_1+n_2) \times d}$. If \mathbf{A} has rank d , then for any vector \mathbf{v} ,*

$$\mathbf{v}^\top (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{v} \geq \mathbf{v}^\top (\mathbf{M}^\top \mathbf{M})^{-1} \mathbf{v}.$$

Proof. Since \mathbf{A} has rank d , \mathbf{M} also has rank d . We have $(\mathbf{A}^\top \mathbf{A})^{-1} \succeq (\mathbf{M}^\top \mathbf{M})^{-1}$ since $\mathbf{M}^\top \mathbf{M} = (\mathbf{A}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{B}) \succeq \mathbf{A}^\top \mathbf{A}$. Thus $\mathbf{v}^\top (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{v} \geq \mathbf{v}^\top (\mathbf{M}^\top \mathbf{M})^{-1} \mathbf{v}$. \square

Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we can obtain a spectral approximation of \mathbf{A} by sampling a few rows of \mathbf{A} according to its leverage scores.

Lemma B.2 (Spectral approximation by leverage score sampling Drineas et al. (2006a); Spielman & Srivastava (2011); Cohen et al. (2015)). *Let $\epsilon, \delta \in (0, 1)$ be two parameters. Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be a matrix. Let $\mathbf{p} \in \mathbb{R}^n$ satisfy*

$$p_i \geq \min\{3\epsilon^{-2}\tau_i(\mathbf{A}) \log(1/\delta), 1\}$$

where $\tau_i(\mathbf{A})$ is the leverage score of the i -th row of \mathbf{A} . Define a random diagonal matrix $D \in \mathbb{R}^{n \times n}$ as follows:

$$D_{i,i} = \begin{cases} \frac{1}{\sqrt{p_i}}, & \text{with probability } p_i, \\ 0, & \text{otherwise.} \end{cases}$$

Then with probability at least $1 - \delta$, $\mathbf{A}^\top D^2 \mathbf{A}$ is an ϵ -spectral approximation of $\mathbf{A}^\top \mathbf{A}$.

For the proof of this lemma, see e.g. Lemma 4 of Cohen et al. (2015).

Least squares regressions can be solved approximately using spectral approximations.

Lemma B.3 (Approximate least squares regression from spectral approximation). *Given $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{b} \in \mathbb{R}^n$, define a matrix $\mathbf{M} := [\mathbf{A}, \mathbf{b}] \in \mathbb{R}^{n \times (d+1)}$. Let $\mathbf{D} \in \mathbb{R}^{n' \times n}$ be a matrix that satisfies $\mathbf{M}^\top \mathbf{D}^\top \mathbf{D} \mathbf{M} \approx_\epsilon \mathbf{M}^\top \mathbf{M}$. Define $\mathbf{x} \in \mathbb{R}^d$ to be*

$$\mathbf{x} := \arg \min_{\mathbf{x}' \in \mathbb{R}^d} \|\mathbf{D} \mathbf{A} \mathbf{x}' - \mathbf{D} \mathbf{b}\|_2.$$

Then with probability at least $1 - \delta$, \mathbf{x} satisfies

$$\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2 \leq (1 + \epsilon) \min_{\mathbf{x}' \in \mathbb{R}^d} \|\mathbf{A} \mathbf{x}' - \mathbf{b}\|_2.$$

See Section 2.5 of Woodruff (2014) for details.

We make use of the Johnson-Lindenstrauss (JL) Lemma in our data structure.

Lemma B.4 (Johnson-Lindenstrauss Lemma (Johnson & Lindenstrauss, 1984)). *There exists a function $\text{JL}(n, \epsilon, \delta, m)$ that returns a random matrix $\mathbf{J} \in \mathbb{R}^{k \times n}$ where $k = O(\epsilon^{-2} \log(m/\delta))$, and \mathbf{J} satisfies that for any fixed m -element subset $V \subset \mathbb{R}^n$,*

$$\Pr [\forall \mathbf{v} \in V, (1 - \epsilon)\|\mathbf{v}\|_2 \leq \|\mathbf{J} \mathbf{v}\|_2 \leq (1 + \epsilon)\|\mathbf{v}\|_2] \geq 1 - \delta.$$

Furthermore, the function JL runs in $O(kn)$ time.

B.2 KALMAN'S METHOD

We review Kalman's approach for dynamic least squares regression, which maintains an exact solution with $O(d^2)$ amortized update time per iteration. First, we note the standard least squares regression has a closed-form solution.

Fact B.5 (Closed-form formula for least squares regression). *Let $n \geq d$ be two integers. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank d , and any vector $\mathbf{b} \in \mathbb{R}^n$, the vector $\mathbf{x}^* := \mathbf{A}^\dagger \cdot \mathbf{b} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \cdot \mathbf{b}$ satisfies*

$$\|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_2 = \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2.$$

We will use the Woodbury identity to compute the changes of the exact solution.

Fact B.6 (Woodbury identity). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{C} \in \mathbb{R}^{k \times k}$, $\mathbf{U} \in \mathbb{R}^{n \times k}$, $\mathbf{V} \in \mathbb{R}^{k \times n}$, one has*

$$(\mathbf{A} + \mathbf{U} \mathbf{C} \mathbf{V})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{U} (\mathbf{C}^{-1} + \mathbf{V} \mathbf{A}^{-1} \mathbf{U})^{-1} \mathbf{V} \mathbf{A}^{-1}.$$

In particular, when \mathbf{U}, \mathbf{V} are vectors, i.e., $\mathbf{U} = \mathbf{a}$, $\mathbf{V} = \mathbf{a}^\top$, and $\mathbf{c} = 1$, one has

$$(\mathbf{A} + \mathbf{a} \mathbf{a}^\top)^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{a} (1 + \mathbf{a}^\top \mathbf{A}^{-1} \mathbf{a})^{-1} \mathbf{a}^\top \mathbf{A}^{-1}.$$

Theorem B.7. *There is an data structure with $O(d^\omega)$ preprocessing time and $O(d^2)$ update time that maintains an exact solution of the dynamic least squares regression problem.*

Proof. For any $t \in [T]$, the data structure maintains $\mathbf{H}^{(t)} := ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1}$ and $\mathbf{u}^{(t)} := (\mathbf{A}^{(t)})^\top \mathbf{b}^{(t)}$. The later one costs $O(d)$ time to update. By Woodbury identity, the former one satisfies

$$\begin{aligned} \mathbf{H}^{(t)} &= ((\mathbf{A}^{(t-1)})^\top \mathbf{A}^{(t-1)} + \mathbf{a}^{(t)} (\mathbf{a}^{(t)})^\top)^{-1} \\ &= \mathbf{H}^{(t-1)} - \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)}, \end{aligned}$$

and it can be updated in $O(d^2)$ time.

The optimal solution $\mathbf{x}^{(t)}$ at step t satisfies $\mathbf{x}^{(t)} = ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1} (\mathbf{A}^{(t)})^\top \mathbf{b}^{(t)} = \mathbf{H}^{(t)} \mathbf{u}^{(t)}$, and it can be computed in $O(d^2)$ time given $\mathbf{H}^{(t)}$ and $\mathbf{u}^{(t)}$. \square

C MISSING PROOFS FROM SECTION 3

We devote to prove the following result.

Theorem C.1 (Restatement of Theorem 3.1). *Let $\epsilon > 0$, $d, T \in \mathbb{N}$. There exists a randomized algorithm for dynamic least-squares regression (Algorithm 1–4). With probability at least 0.9, the algorithm maintains an ϵ -approximation solution for all iterations $t \in [T]$ and the total update time over T iterations is at most $O(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(TD/\sigma))$. Our data structure uses at most $O(\epsilon^{-2} d^2 \cdot \log^2(TD/\sigma))$ space.*

Adversarial input We remark that our data structure works against adaptive adversaries. This is because we only re-use the random JL matrix when the new row is not sampled, and in this case the adversary cannot learn any information since our data structure behaves in exactly the same manner as a deterministic data structure that rejects all new rows. Whenever a new row is sampled, we immediately use fresh randomness to generate a new JL matrix, and any information the adversary might learn about the old JL matrix becomes useless.

C.1 MISSING PROOFS FROM SECTION 3.2

We prove the correctness of algorithm. We start by providing closed-form formulas for all the variables we maintain in each iteration.

Lemma C.2 (Restatement of Lemma 3.3). *At the t -th iteration of Algorithm 1–4, we have*

1. $\mathbf{M}^{(t)} = [\mathbf{A}^{(t)}, \mathbf{b}^{(t)}] \in \mathbb{R}^{(d+t+1) \times (d+1)}$.
2. $\mathbf{D}^{(t)} \in \mathbb{R}^{(d+t+1) \times (d+t+1)}$ is a diagonal matrix with $s^{(t)}$ non-zero entries.
3. $\mathbf{N}^{(t)} = (\mathbf{D}^{(t)} \mathbf{M}^{(t)})_{S^{(t)}, *}, \in \mathbb{R}^{s^{(t)} \times (d+1)}$, where $S^{(t)} \subset [d+t+1]$ is defined as the set of non-zero entries of $\mathbf{D}^{(t)}$.
4. $\mathbf{H}^{(t)} = ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1} \in \mathbb{R}^{(d+1) \times (d+1)}$.
5. $\mathbf{B}^{(t)} = \mathbf{N}^{(t)} \mathbf{H}^{(t)} \in \mathbb{R}^{s^{(t)} \times (d+1)}$.
6. $\tilde{\mathbf{B}}^{(t)} = \mathbf{J}^{(t)} \cdot \mathbf{B}^{(t)} \in \mathbb{R}^{k \times (d+1)}$, where $k = O(\epsilon^{-2} \log(T/\delta))$.
7. $\mathbf{G}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \in \mathbb{R}^{d \times d}$.
8. $\mathbf{u}^{(t)} = (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)} \in \mathbb{R}^d$.
9. $\mathbf{x}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \cdot (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)} \in \mathbb{R}^d$.

Proof. First note that all claims of the lemma hold for $t = 0$ in PREPROCESS.

We assume that the claims hold for $t-1$, and inductively prove them for t . We only prove the lemma for $\mathbf{H}^{(t)}$ (Part 4), $\mathbf{B}^{(t)}$ (Part 5), $\mathbf{G}^{(t)}$ (Part 7) and $\mathbf{u}^{(t)}$ (Part 8). The rest follows directly from the algorithm description.

Part 4 ($\mathbf{H}^{(t)} = ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1}$) If $\nu^{(t)} = 0$, then $\mathbf{N}^{(t)} = \mathbf{N}^{(t-1)}$ and $\mathbf{H}^{(t)} = \mathbf{H}^{(t-1)} = ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1}$. Otherwise when $\nu^{(t)} \neq 0$, one has $\mathbf{N}^{(t)} = [(\mathbf{N}^{(t-1)})^\top, \mathbf{m}^{(t)}/\sqrt{p^{(t)}}]^\top$. Using Woodbury identity, we have

$$\begin{aligned}
& ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1} \\
&= ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)} + (\mathbf{m}^{(t)})^\top \mathbf{m}^{(t)}/p^{(t)})^{-1} \\
&= ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1} - \frac{((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1} \mathbf{m}^{(t)} (\mathbf{m}^{(t)})^\top ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1} / p^{(t)}}{1 + (\mathbf{m}^{(t)})^\top ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1} \mathbf{m}^{(t)}/p^{(t)}} \\
&= \mathbf{H}^{(t-1)} - \frac{\mathbf{H}^{(t-1)} \mathbf{m}^{(t)} (\mathbf{m}^{(t)})^\top \mathbf{H}^{(t-1)} / p^{(t)}}{1 + (\mathbf{m}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{m}^{(t)}/p^{(t)}},
\end{aligned}$$

The second term is exactly the $\Delta\mathbf{H}$ term when calling UPDATEMEMBERS (Line 2). Hence, $\mathbf{H}^{(t)} = \mathbf{H}^{(t-1)} + \Delta\mathbf{H} = ((\mathbf{N}^{(t)})^\top \mathbf{N}^{(t)})^{-1}$.

Part 5 ($\mathbf{B}^{(t)} = \mathbf{N}^{(t)}\mathbf{H}^{(t)}$) If $\nu^{(t)} = 0$, then $\mathbf{N}^{(t)} = \mathbf{N}^{(t-1)}$ and $\mathbf{H}^{(t)} = \mathbf{H}^{(t-1)}$, so $\mathbf{B}^{(t)} = \mathbf{B}^{(t-1)} = \mathbf{N}^{(t)}\mathbf{H}^{(t)}$. Otherwise when $\nu^{(t)} \neq 0$, one has $\mathbf{N}^{(t)} = [(\mathbf{N}^{(t-1)})^\top, \mathbf{m}^{(t)}/\sqrt{p^{(t)}}]^\top$ and $\mathbf{H}^{(t)} = \mathbf{H}^{(t-1)} + \Delta\mathbf{H}$. We have

$$\begin{aligned}\mathbf{N}^{(t)}\mathbf{H}^{(t)} &= \begin{bmatrix} \mathbf{N}^{(t-1)} \cdot \mathbf{H}^{(t)} \\ (\mathbf{m}^{(t)})^\top \cdot \mathbf{H}^{(t)}/\sqrt{p^{(t)}} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{B}^{(t-1)} + \mathbf{N}^{(t-1)} \cdot \Delta\mathbf{H} \\ (\mathbf{m}^{(t)})^\top \cdot \mathbf{H}^{(t)}/\sqrt{p^{(t)}} \end{bmatrix}\end{aligned}$$

This is exactly what we compute in Line 4 of UPDATEMEMBERS.

Part 7 ($\mathbf{G}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1}$) The proof is analogous to that of Part 4 ($\mathbf{H}^{(t)}$). If $\nu^{(t)} = 0$, then $\mathbf{G}^{(t)} = \mathbf{G}^{(t-1)} = ((\mathbf{A}^{(t-1)})^\top (\mathbf{D}^{(t-1)})^2 \mathbf{A}^{(t-1)})^{-1} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1}$.

On the other hand, when $\nu^{(t)} \neq 0$, using Woodbury identity, one has

$$\begin{aligned}((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} &= ((\mathbf{A}^{(t-1)})^\top (\mathbf{D}^{(t-1)})^2 \mathbf{A}^{(t-1)} + (\mathbf{a}^{(t)})^\top \mathbf{a}^{(t)}/p^{(t)})^{-1} \\ &= \mathbf{G}^{(t-1)} - \frac{\mathbf{G}^{(t-1)} \mathbf{a}^{(t)} (\mathbf{a}^{(t)})^\top \mathbf{G}^{(t-1)}/p^{(t)}}{1 + (\mathbf{a}^{(t)})^\top \mathbf{G}^{(t-1)} \mathbf{a}^{(t)}/p^{(t)}} = \mathbf{G}^{(t)}.\end{aligned}$$

This is exactly what we compute in Line 8 of UPDATEMEMBERS.

Part 8 ($\mathbf{u}^{(t)} = (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)}$) We focus on the case $\nu^{(t)} \neq 0$, and we have

$$\mathbf{u}^{(t)} = \mathbf{u}^{(t-1)} + \beta^{(t)} \cdot \mathbf{a}^{(t)}/p^{(t)} = \mathbf{A}^{(t-1)} (\mathbf{D}^{(t-1)})^2 \mathbf{b}^{(t-1)} + (\nu^{(t)})^2 \cdot \mathbf{a}^{(t)} \cdot \beta^{(t)} = \mathbf{A}^{(t)} (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)}$$

The first step follows from the updating rule of the data structure (Line 9 of UPDATEMEMBERS), and the second step follows from $\mathbf{u}^{(t-1)} = \mathbf{A}^{(t-1)} (\mathbf{D}^{(t-1)})^2 \mathbf{b}^{(t-1)}$ and $\nu^{(t)} = 1/\sqrt{p^{(t)}}$, the last step follows from the definition of $\mathbf{A}^{(t)}$, $\mathbf{b}^{(t)}$ and $\mathbf{D}^{(t)}$. We conclude the proof. \square

The following lemma is the key to our proof, it asserts that $\tau^{(t)}$ gives a suitable estimation on the leverage score, and consequently, the sub-sampled matrix $\mathbf{N}^{(t)} = (\mathbf{D}^{(t)}\mathbf{M}^{(t)})_{S^{(t)},*}$ obtains good spectral approximation of $\mathbf{M}^{(t)}$, with high probability.

Lemma C.3 (Restatement of Lemma 3.4). *With probability at least $1 - 2T\delta$,*

$$(1 - \epsilon)^2 \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \tau^{(t)} \leq (1 + \epsilon)^2 \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}), \quad \forall t \in [T], \quad (5)$$

and

$$(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)} \approx_\epsilon (\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}, \quad \forall t \in [0 : T]. \quad (6)$$

Proof. We prove the claim inductively. The induction hypothesis is that with probability $1 - 2t\delta$, Eq. (6) holds for all $t' \in [0 : t]$ and Eq. (5) holds for all $t' \in [t]$. The base case $t = 0$ holds trivially, as $\mathbf{D}^{(0)} = \mathbf{I}_{d+1}$, and therefore, $(\mathbf{M}^{(0)})^\top (\mathbf{D}^{(0)})^2 \mathbf{M}^{(0)} = (\mathbf{M}^{(0)})^\top \mathbf{M}^{(0)}$.

We next assume the induction hypothesis holds for $t - 1$, and prove it for t . We assume that Eq. (6) holds for all $t' \in [0 : t - 1]$ and Eq. (5) holds for all $t' \in [t - 1]$, and later add the $2(t - 1)\delta$ failure probability using Union bound.

Step 1 We first prove that

$$(1 - \epsilon) \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \|\mathbf{B}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq (1 + \epsilon) \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}). \quad (7)$$

Using the closed-form formulas of Lemma 3.3, we have

$$\begin{aligned}
\|\mathbf{B}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 &= (\mathbf{m}^{(t)})^\top \cdot (\mathbf{H}^{(t-1)})^\top (\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)} \mathbf{H}^{(t-1)} \cdot \mathbf{m}^{(t)} \\
&= (\mathbf{m}^{(t)})^\top \cdot ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1} \cdot \mathbf{m}^{(t)} \\
&= (\mathbf{m}^{(t)})^\top \cdot ((\mathbf{M}^{(t-1)})^\top (\mathbf{D}^{(t-1)})^2 \mathbf{M}^{(t-1)})^{-1} \cdot \mathbf{m}^{(t)} \\
&\leq (1 + \epsilon) \cdot (\mathbf{m}^{(t)})^\top \cdot ((\mathbf{M}^{(t-1)})^\top \mathbf{M}^{(t-1)})^{-1} \cdot \mathbf{m}^{(t)} \\
&= (1 + \epsilon) \cdot \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}),
\end{aligned}$$

where the first step follows from $\mathbf{B}^{(t-1)} = \mathbf{N}^{(t-1)} \mathbf{H}^{(t-1)}$, the second step follows from $\mathbf{H}^{(t-1)} = ((\mathbf{N}^{(t-1)})^\top \mathbf{N}^{(t-1)})^{-1}$, the third step follows from $\mathbf{N}^{(t-1)} = (\mathbf{D}^{(t-1)} \mathbf{M}^{(t-1)})_{S^{(t-1)},*}$ with $S^{(t-1)}$ being the set of non-zero entries of $\mathbf{D}^{(t-1)}$, the fourth step follows from the inductive hypothesis, and the last step follows from the definition of generalized leverage scores that $\tau_i^{\mathbf{B}}(\mathbf{A}) = a_i^\top (\mathbf{B}^\top \mathbf{B})^{-1} a_i$.

The other direction of Eq. (7) follows from a similar argument.

Step 2 We next prove that with probability at least $1 - \delta$,

$$(1 - \epsilon)^2 \cdot \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq \|\tilde{\mathbf{B}}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq (1 + \epsilon)^2 \cdot \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}). \quad (8)$$

By the closed-form formulas of Lemma 3.3, we have $\tilde{\mathbf{B}}^{(t-1)} = \mathbf{J}^{(t-1)} \cdot \mathbf{B}^{(t-1)}$, where $\mathbf{J}^{(t-1)} \in \mathbb{R}^{k \times s^{(t)}}$ and $k = O(\epsilon^{-2} \log(T/\delta))$. Using the JL Lemma (Lemma B.4), we have that with probability at least $1 - \delta$,

$$(1 - \epsilon) \|\mathbf{B}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq \|\tilde{\mathbf{B}}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2 \leq (1 + \epsilon) \|\mathbf{B}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2.$$

Combining with Eq. (7) and using an union bound, we have that Eq. (8) holds with probability at least $1 - 2\delta$. We remark that we invoke the JL property of the matrix $\mathbf{J}^{(t-1)}$ at most T times, and since its dimension is $k = O(\epsilon^{-2} \log(T/\delta))$, the above argument is correct for all T iterations.

Since $\tau^{(t)} = \|\tilde{\mathbf{B}}^{(t-1)} \cdot \mathbf{m}^{(t)}\|_2^2$, this proves the first part of the lemma.

Step 3 Finally, we conclude the proof by proving the second part of induction, and show that conditioned on Eq. (8) holds, with probability at least $1 - \delta$, we have

$$(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)} \approx_\epsilon (\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}. \quad (9)$$

For $i \in [d+1]$, we have $\mathbf{D}_{i,i}^{(t)} = 1$. For $i \in [t]$, we have $\mathbf{D}_{d+i+1, d+i+1}^{(t)} = 1/\sqrt{p^{(i)}}$ with probability $p^{(i)}$ and otherwise $\mathbf{D}_{d+i+1, d+i+1}^{(t)} = 0$, where the sampling probability $p^{(i)}$ satisfies

$$\begin{aligned}
p^{(i)} &= \min\{3(1 + \epsilon)^2 \epsilon^{-2} \tau^{(i)} \log(1/\delta), 1\} \\
&\geq \min\{3\epsilon^{-2} \cdot \tau_{d+i+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(i)}) \cdot \log(1/\delta), 1\} \\
&\geq \min\{3\epsilon^{-2} \cdot \tau_{d+i+1}(\mathbf{M}^{(t-1)}) \cdot \log(1/\delta), 1\}
\end{aligned}$$

where the first step follows from the definition, the second step follows from the induction hypothesis and the last step follows from Fact B.1.

Thus the sampling matrix $\mathbf{D}^{(t)}$ satisfies the constraints of Lemma B.2, so with probability at least $1 - \delta$, one has that $(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)}$ is an ϵ -spectral approximation of $(\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}$. This proves Eq. (9).

Finally, taking an union bound over the δ failure probability of Step 2, the δ failure probability of Step 5, and the $2(t-1)\delta$ failure probability of the induction hypothesis, we complete the proof. \square

We are now ready to prove the correctness of our data structure.

Lemma C.4 (Restatement of Lemma 3.5). *With probability at least $1 - O(1/T)$, in each iteration, UPDATE of Algorithm 2 outputs a vector $\mathbf{x}^{(t)} \in \mathbb{R}^d$ such that*

$$\|\mathbf{A}^{(t)} \mathbf{x}^{(t)} - \mathbf{b}^{(t)}\|_2 \leq (1 + \epsilon) \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{A}^{(t)} \mathbf{x} - \mathbf{b}^{(t)}\|_2.$$

Proof. Note by Lemma 3.3, one has $\mathbf{x}^{(t)} = ((\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{A}^{(t)})^{-1} \cdot (\mathbf{A}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{b}^{(t)}$, which is the closed-form minimizer of $\|\mathbf{D}^{(t)} \mathbf{A}^{(t)} \mathbf{x} - \mathbf{D} \mathbf{b}^{(t)}\|_2$. From Eq. (6) in Lemma 3.4, we know that with probability at least $1 - 2T\delta = 1 - O(1/T)$ (since $\delta = O(1/T^2)$), we have $(\mathbf{M}^{(t)})^\top (\mathbf{D}^{(t)})^2 \mathbf{M}^{(t)} \approx_\epsilon (\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}$. Using this spectral approximation and Lemma B.3, we conclude the proof of this lemma. \square

C.2 MISSING PROOFS FORM SECTION 3.3

We first bound the worst-case running time of the UPDATE procedure at the t -th iteration. Without loss of generality, we may assume $\text{nnz}(\mathbf{a}^{(t)}) > 0$, otherwise we can ignore the t -th update.

Lemma C.5 (Restatement of Lemma 3.6). *At the t -th iteration of the UPDATE procedure,*

- *If $\nu^{(t)} = 0$, then UPDATE takes $O(\epsilon^{-2} \log(T/\delta) \cdot \text{nnz}(\mathbf{a}^{(t)}))$ time.*
- *If $\nu^{(t)} \neq 0$, then UPDATE takes $O(\epsilon^{-2} s^{(t)} d \log(T/\delta))$ time.*

Proof. If $\nu^{(t)} = 0$, we only need to invoke the SAMPLE procedure. The most time-consuming step of SAMPLE is to compute $\tilde{\mathbf{B}}^{(t-1)} \cdot \mathbf{m}^{(t)}$ when computing $\tau^{(t)}$ (Line 1). Since $\tilde{\mathbf{B}}^{(t-1)} \in \mathbb{R}^{k \times (d+1)}$ where $k = O(\epsilon^{-2} \log(T/\delta))$ and $\mathbf{m}^{(t)} = [(\mathbf{a}^{(t)})^\top, \beta^{(t)}]^\top$, this takes $O(\epsilon^{-2} \log(T/\delta) \cdot \text{nnz}(\mathbf{a}^{(t)}))$ time.

If $\nu^{(t)} \neq 0$, besides the SAMPLE procedure, the data structure also needs to invoke the UPDATERMEMBERS procedure. The most time-consuming step is to compute $\tilde{\mathbf{B}}^{(t)} = \mathbf{J}^{(t)} \cdot \mathbf{B}^{(t)}$ on Line 7. Indeed, it's easy to see that all other computations only involve matrix-vector multiplications and matrix additions, and they can be computed in $O(s^{(t)} d)$ time. Since $\mathbf{J}^{(t)} \in \mathbb{R}^{k \times s^{(t)}}$ for $k = O(\epsilon^{-2} \log(T/\delta))$, and $\mathbf{B}^{(t)} \in \mathbb{R}^{s^{(t)} \times (d+1)}$, computing $\tilde{\mathbf{B}}^{(t)} = \mathbf{J}^{(t)} \cdot \mathbf{B}^{(t)}$ takes $O(\epsilon^{-2} s^{(t)} d \log(T/\delta))$ time. \square

In order to bound the amortized update time, we need to bound the sum of $\tau^{(t)}$. We use a slightly generalized version of Theorem 2.2 of Cohen et al. (2020).

Lemma C.6 (Restatement of Lemma 3.7). *Let $\mathbf{M}^{(0)} \in \mathbb{R}^{(d+1) \times (d+1)}$ be a full rank matrix that satisfy $\sigma_{d+1}(\mathbf{M}^{(0)}) \geq \sigma$ for some small $\sigma \in (0, 1)$. For T vectors $\mathbf{m}^{(1)}, \mathbf{m}^{(2)}, \dots, \mathbf{m}^{(T)} \in \mathbb{R}^{d+1}$, define matrix $\mathbf{M}^{(t)} \in \mathbb{R}^{(d+t+1) \times (d+1)}$ to be $\mathbf{M}^{(t)}$ appended with a new row $(\mathbf{m}^{(t)})^\top$. Assume that $\|\mathbf{M}_{i,*}^{(T)}\|_2 \leq D$ for all $i \in [d + T + 1]$. Then*

$$\sum_{t=1}^T \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq O(d \log(TD/\sigma)).$$

For completeness we include a proof here. The proof slightly generalizes that of Theorem 2.2 of Cohen et al. (2020).

Proof. For simplicity, $\forall t \in [T]$ we denote

$$l_t := \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) = (\mathbf{m}^{(t)})^\top ((\mathbf{M}^{(t-1)})^\top \mathbf{M}^{(t-1)})^{-1} \mathbf{m}^{(t)}.$$

For any $t \in [T - 1]$, we have

$$\begin{aligned} \det((\mathbf{M}^{(t+1)})^\top \mathbf{M}^{(t+1)}) &= \det((\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}) \cdot (1 + (\mathbf{m}^{(t+1)})^\top ((\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)})^{-1} \mathbf{m}^{(t+1)}) \\ &\geq \det((\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}) \cdot (1 + l_{t+1}) \\ &\geq \det((\mathbf{M}^{(t)})^\top \mathbf{M}^{(t)}) \cdot e^{l_{t+1}/2}, \end{aligned}$$

where the first step follows from $\det(\mathbf{S} + \mathbf{u}\mathbf{u}^\top) = \det(\mathbf{S}) \cdot (1 + \mathbf{u}^\top \mathbf{S}^{-1} \mathbf{u})$ for any invertible square matrix \mathbf{S} and any vector \mathbf{u} .

Hence,

$$\begin{aligned} \det((\mathbf{M}^{(T)})^\top \mathbf{M}^{(T)}) &\geq \det((\mathbf{M}^{(0)})^\top \mathbf{M}^{(0)}) \cdot e^{\sum_{t=1}^T l_{t+1}/2} \\ &\geq \sigma^{2(d+1)} \cdot e^{\sum_{t=1}^T l_{t+1}/2}. \end{aligned}$$

Since $\det((\mathbf{M}^{(T)})^\top \mathbf{M}^{(T)}) \leq \|(\mathbf{M}^{(T)})^\top \mathbf{M}^{(T)}\|_2^{d+1} \leq (D^2(T+d+1))^{d+1}$, we have

$$(D^2(T+d+1))^{d+1} \geq \sigma^{2(d+1)} \cdot e^{\sum_{t=1}^T l_{t+1}/2}.$$

Taking logarithms of both sides, we have

$$\begin{aligned} (d+1) \log(D^2(T+d+1)) &\geq 2(d+1) \log(\sigma) + \sum_{t=1}^T l_{t+1}/2 \\ \implies \sum_{t=1}^T l_{t+1} &\leq O(d \log(TD/\sigma)). \quad \square \end{aligned}$$

Lemma C.7 (Upper bound on sum of $\tau^{(t)}$). *With probability at least $1 - 2T\delta$,*

$$\sum_{t=1}^T \tau^{(t)} \leq O(d \log(TD/\sigma)).$$

Proof. From Lemma 3.4, with probability at least $1 - 2T\delta$, $\tau^{(t)} \leq (1+\epsilon)^2 \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)})$ holds for all $t \in [T]$, so using Theorem 3.7 we have

$$\sum_{t=1}^T \tau^{(t)} \leq (1+\epsilon)^2 \sum_{t=1}^T \tau_{d+t+1}^{\mathbf{M}^{(t-1)}}(\mathbf{M}^{(t)}) \leq O(d \log(TD/\sigma)). \quad \square$$

We can now bound the amortized update time of Algorithm 2.

Lemma C.8 (Restatement of Lemma 3.8). *With probability at least 0.99, the total running time of UPDATE over T iterations is at most $O(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(TD/\sigma))$.*

Proof. Consider the sampling matrix $\mathbf{D}^{(T)} \in \mathbb{R}^{(d+T+1) \times (d+T+1)}$. Recall that $s^{(T)}$ is the total number of non-zero diagonal entries of $\mathbf{D}^{(T)}$. For $i \in [d+1]$, $\mathbf{D}_{i,i}^{(T)} = 1$, and for $t \in [T]$, $\mathbf{D}_{d+t+1, d+t+1}^{(T)} \neq 0$ with probability $p^{(t)} = \min\{3(1+\epsilon)^2 \epsilon^{-2} \tau^{(t)} \log(1/\delta), 1\}$. So the expectation of $s^{(T)}$ is bounded by

$$\begin{aligned} \mathbb{E}[s^{(T)}] &= d+1 + \sum_{t=1}^T p^{(t)} \\ &\leq d+1 + \min \left\{ O \left(\epsilon^{-2} \cdot \log(1/\delta) \cdot \sum_{t=1}^T \tau^{(t)} \right), T \right\} \\ &\leq d+1 + O(\epsilon^{-2} d \cdot \log(TD/\sigma) \log(1/\delta)) \cdot (1 - 2T\delta) + 2T\delta \cdot T \\ &\leq O(\epsilon^{-2} d \cdot \log(TD/\sigma) \log(1/\delta)). \end{aligned}$$

The second step follows from $p^{(t)} = \min\{3(1+\epsilon)^2 \epsilon^{-2} \tau^{(t)} \log(1/\delta), 1\}$, the third step follows from the fact that $\sum_{t=1}^T \tau^{(t)} \leq O(d \log(TD/\sigma))$ holds with probability $1 - 2T\delta$ (Lemma C.7), and the last step uses $\delta = O(1/T^2)$.

By Markov inequality, with probability at least 0.99, one has $s^{(T)} \leq O(\epsilon^{-2} d \cdot \log(TD/\sigma) \log(1/\delta))$. Using Lemma 3.6, the total running time is bounded by

$$\begin{aligned} &\underbrace{\sum_{t=1}^T O(\epsilon^{-2} \log(T/\delta) \cdot \text{nnz}(\mathbf{a}^{(t)}))}_{\text{when } \nu^{(t)}=0} + \underbrace{s^{(T)} \cdot O(\epsilon^{-2} s^{(T)} d \log(T/\delta))}_{\text{when } \nu^{(t)} \neq 0} \\ &\leq O \left(\epsilon^{-2} \text{nnz}(\mathbf{A}^{(T)}) \log(T) + \epsilon^{-6} d^3 \log^5(TD/\sigma) \right). \end{aligned}$$

We conclude the proof here. \square

Remark C.9 (Space usage). *Since the largest matrices that the data structure maintains and updates in each iteration are $\mathbf{B}^{(t)}, \mathbf{N}^{(t)} \in \mathbb{R}^{s^{(t)} \times (d+1)}$, it is straightforward to see that the total space used by the data structure is bounded by $O(s^{(T)} \cdot d)$.*

In Lemma 3.8 we already proved that with probability at least 0.99, we have $s^{(T)} \leq O(\epsilon^{-2}d \cdot \log(TD/\sigma) \log(1/\delta))$. So the total space is bounded by

$$O(\epsilon^{-2}d^2 \cdot \log^2(TD/\sigma)).$$

D MISSING PROOFS FROM SECTION 4

In this section we provide the complete proof of Theorem 4.2. We consider the Word RAM model where the word size $w = O(\log d)$ in all results in Section 4.

Remark D.1 (Limited precision in the Word RAM model). *For simplicity, we present the proofs in Section D.1 and Section D.2 assuming infinite precision arithmetic. These proofs are still correct if we work in the Word RAM model with word size $w = 10 \log(dT) = O(\log d)$. This is because the $2^{-w} = 1/(d^{10}T^{10})$ error caused by the Word RAM model is much smaller than both the $1/d^2$ approximation error of Lemma 4.3, and the $1/(d^8T^2)$ approximation error of Theorem 4.2.*

We prove OMv-hardness for well-conditioned PSD matrix in Section D.1, and reduce OMv to dynamic least squares regression in Section D.2.

D.1 OMV-HARDNESS FOR PSD MATRIX

We prove Lemma 4.3 in this section.

Lemma D.2 (Restatement of Lemma 4.3). *Let $d \in \mathbb{N}, T = \text{poly}(d)$. Let $\gamma > 0$ be any constant. Let $\mathbf{H} \in \mathbb{R}^{d \times d}$ be a symmetric matrix whose eigenvalues satisfy $1 \leq \lambda_d(\mathbf{H}) \leq \dots \leq \lambda_1(\mathbf{H}) \leq 3$. For any $t \in [T]$, $\mathbf{z}^{(t)} \in \mathbb{R}^d$ is revealed at the t -th step, and $\|\mathbf{z}^{(t)}\|_2 \leq 1$. Assuming the OMv conjecture is true, then there is no algorithm with $\text{poly}(d)$ preprocessing time and $O(d^{2-\gamma})$ amortized running time that can return an $O(1/d^2)$ -approximate answer to $\mathbf{H}\mathbf{z}^{(t)}$ for all t , i.e., a vector $\mathbf{y}^{(t)} \in \mathbb{R}^d$ s.t. $\|\mathbf{y}^{(t)} - \mathbf{H}\mathbf{z}^{(t)}\|_2 \leq \epsilon$, and has an error probability $\leq 1/3$.*

Proof. Given a Boolean matrix $\mathbf{B} \in \{0, 1\}^{d \times d}$ in the OMv conjecture (Conjecture 4.1), we construct a PSD matrix

$$\mathbf{H} = \begin{bmatrix} 2\mathbf{I}_d & \frac{1}{d}\mathbf{B} \\ \frac{1}{d}\mathbf{B}^\top & 2\mathbf{I}_d \end{bmatrix} \in \mathbb{R}^{2d \times 2d}.$$

We note that \mathbf{H} is symmetric and $1 \leq \lambda_d(\mathbf{H}) \leq \lambda_1(\mathbf{H}) \leq 3$, since for any $\mathbf{z} = (\mathbf{z}', \mathbf{z}'') \in \mathbb{R}^{2n}$ with $\|\mathbf{z}\|_2^2 = 1$, one has

$$\begin{aligned} \mathbf{z}^\top \mathbf{H} \mathbf{z} &= 2\|\mathbf{z}'\|_2^2 + 2\|\mathbf{z}''\|_2^2 + \frac{1}{d}(\mathbf{z}')^\top \mathbf{B}^\top \mathbf{z}'' + \frac{1}{d}(\mathbf{z}'')^\top \mathbf{B} \mathbf{z}' \\ &\leq 2\|\mathbf{z}'\|_2^2 + 2\|\mathbf{z}''\|_2^2 + 2\|\mathbf{z}'\|_2 \|\mathbf{z}''\|_2 \in (1, 3). \end{aligned}$$

The second step follows from $\mathbf{B} \in \{0, 1\}^d$, and therefore, $\|\mathbf{B}\mathbf{y}\|_2 \leq d\|\mathbf{y}\|_2$ and $\|\mathbf{B}^\top \mathbf{y}\|_2 \leq d\|\mathbf{y}\|_2$ hold for any $\mathbf{y} \in \mathbb{R}^d$.

Given an online query $\mathbf{z} \in \{0, 1\}^d$ for \mathbf{B} , we can assume w.l.o.g. that $\mathbf{z} \neq \mathbf{0}_d$. We construct a query vector for \mathbf{H} as $\bar{\mathbf{z}} = (\mathbf{0}_d, \frac{\mathbf{z}}{\|\mathbf{z}\|_2}) \in \mathbb{R}^{2d}$. Clearly one has $\|\bar{\mathbf{z}}\|_2 = 1$.

We prove that one can recover $\mathbf{B}\mathbf{z}$ from an $O(1/d^2)$ -approximate answer to $\mathbf{H}\bar{\mathbf{z}}$. Let $\hat{\mathbf{y}} = (\mathbf{y}', \mathbf{y}'') \in \mathbb{R}^{2d}$ be such an approximate answer, i.e., $\|\hat{\mathbf{y}} - \mathbf{H}\bar{\mathbf{z}}\|_2 \leq O(1/d^2)$. We construct $\mathbf{y} \in \{0, 1\}^d$ such that $\forall i \in [d]$, $y_i = 1$ if $d\|\mathbf{z}\|_2 \cdot y'_i \geq 0.5$ and $y_i = 0$ otherwise. Next we prove that $\mathbf{y} = \mathbf{B}\mathbf{z}$.

$$\|(\mathbf{B}\mathbf{z} - d\|\mathbf{z}\|_2 \cdot \mathbf{y}')\|_2 = d\|\mathbf{z}\|_2 \cdot \left\| \frac{1}{d}\mathbf{B} \frac{\mathbf{z}}{\|\mathbf{z}\|_2} - \mathbf{y}' \right\|_2 \leq d\|\mathbf{z}\|_2 \cdot \|\mathbf{H}\bar{\mathbf{z}} - \hat{\mathbf{y}}\|_2 \leq d \cdot \sqrt{d} \cdot O(1/d^2) < 0.5.$$

The second step follows from $\mathbf{H}\bar{\mathbf{z}} = (\frac{1}{d}\mathbf{B} \frac{\mathbf{z}}{\|\mathbf{z}\|_2}, \frac{2\mathbf{z}}{\|\mathbf{z}\|_2}) \in \mathbb{R}^{2d}$, and in the third step we use the fact that $\|\hat{\mathbf{y}} - \mathbf{H}\bar{\mathbf{z}}\|_2 \leq O(1/d^2)$, and $\|\mathbf{z}\|_2 \leq \sqrt{d}$ since $\mathbf{z} \in \{0, 1\}^d$.

Since each entry of $\mathbf{B}\mathbf{z}$ is an integer, rounding $d\|\mathbf{z}\|_2 \cdot \mathbf{y}'$ to the closest integer gives the exact solution. Hence, we have $\mathbf{y} = \mathbf{B}\mathbf{z}$, and we conclude the proof here. \square

D.2 REDUCING DYNAMIC LEAST-SQUARES REGRESSION TO OMV

For any two vectors \mathbf{x}, \mathbf{y} and $\delta > 0$, we write $\mathbf{x} = \mathbf{y} \pm \delta$ if $\|\mathbf{x} - \mathbf{y}\|_2 \leq \delta$. With Lemma 4.3 in hand, we can prove Theorem 4.2.

Theorem D.3 (Restatement of Theorem 4.2). *Let $d \in \mathbb{N}$, $T = \text{poly}(d)$, $\epsilon = \frac{1}{d^8 T^2} = 1/\text{poly}(d)$, and let $\gamma > 0$ be any constant. Assuming the OMv conjecture is true, any dynamic algorithm that maintains an ϵ -approximate solution of the least squares regression requires at least $\Omega(d^{2-\gamma})$ amortized time per update.*

Proof. We reduce from the problem of Lemma 4.3. Given a PSD matrix $\mathbf{H} \in \mathbb{R}^{d \times d}$ for the problem of Lemma 4.3, where $1 \leq \lambda_1(\mathbf{H}) \leq \lambda_d(\mathbf{H}) \leq 3$, we compute $\mathbf{A}^\top \mathbf{A} = \mathbf{H}^{-1}$. In the least squares regression problem, we take \mathbf{A} to be the initial matrix, and let $\mathbf{b} = \mathbf{0}_d \in \mathbb{R}^d$ be the initial label vector. The preprocessing step takes $O(d^\omega)$ time.

In the online stage of the problem of Lemma 4.3, let the query at the t -th step be $\mathbf{z}^{(t)} \in \mathbb{R}^d$ where $\|\mathbf{z}^{(t)}\|_2 \leq 1$. We scale the vector to construct

$$\mathbf{a}^{(t)} = \frac{1}{d^2 \sqrt{T}} \cdot \mathbf{z}^{(t)} \in \mathbb{R}^d,$$

and use $(\mathbf{a}^{(t)}, 1) \in \mathbb{R}^d \times \mathbb{R}$ as the update to the regression problem at the t -th step. Let $\mathbf{x}^{(t)} \in \mathbb{R}^d$ be the solution returned by the dynamic algorithm for least squares regression at the t -th step. We prove that one can answer the matrix-vector query with

$$\mathbf{y}^{(t)} = d^2 \sqrt{T} (\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}),$$

and we have the guarantee that

$$\|\mathbf{y}^{(t)} - \mathbf{H}\mathbf{z}^{(t)}\|_2 \leq O(1/d^2). \quad (10)$$

By Lemma 4.3, this is impossible if the OMv conjecture is true.

We first introduce some notations. Let $\mathbf{A}^{(t)} \in \mathbb{R}^{(d+t) \times d}$ be the data matrix after the t -th time step, $\mathbf{b}^{(t)} \in \mathbb{R}^{d+t}$ be the labels, and $\mathbf{H}^{(t)} := ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1} \in \mathbb{R}^{d \times d}$. For simplicity, we also define $\mathbf{A}^{(0)} = \mathbf{A}$, $\mathbf{b}^{(0)} = \mathbf{b}$, $\mathbf{H}^{(0)} = \mathbf{H}$. For any $t \in [T]$, let $\mathbf{x}_*^{(t)}$ be the optimal solution at the t -th step, and it has the closed-form $\mathbf{x}_*^{(t)} = ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1} (\mathbf{A}^{(t)})^\top \cdot \mathbf{b}^{(t)}$. The proof divides into three steps:

- Step 1. $\mathbf{x}^{(t)}$ and $\mathbf{x}_*^{(t)}$ are close, i.e., $\mathbf{x}^{(t)} = \mathbf{x}_*^{(t)} \pm O(\frac{1}{d^4 \sqrt{T}})$.
- Step 2. $\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}$ recovers $\mathbf{H}^{(t-1)} \mathbf{a}^{(t)}$, i.e., $\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)} = \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \pm O(\frac{1}{d^4 \sqrt{T}})$.
- Step 3. $\mathbf{H}^{(t-1)} \mathbf{a}^{(t)}$ is close to $\mathbf{H} \mathbf{a}^{(t)}$, i.e., $\mathbf{H}^{(t-1)} \mathbf{a}^{(t)} = \mathbf{H} \mathbf{a}^{(t)} \pm O(\frac{1}{d^6 \sqrt{T}})$.

In particular, Step 2 and 3 directly implies Eq. (10).

We first prove a useful bound on the singular values of $\mathbf{A}^{(t)}$ and $\mathbf{H}^{(t)}$ for all $t \in [T]$. First note that the assumption $1 \leq \lambda_1(\mathbf{H}) \leq \lambda_d(\mathbf{H}) \leq 3$ implies that $\frac{1}{\sqrt{3}} \leq \sigma_d(\mathbf{A}) \leq \sigma_1(\mathbf{A}) \leq 1$. Since $\|\mathbf{a}^{(t)}\|_2 = \|\frac{1}{d^2 \sqrt{T}} \cdot \mathbf{z}^{(t)}\|_2 \leq \frac{1}{d^2 \sqrt{T}}$ for any $t \in [T]$, we have that for any $\mathbf{x} \in \mathbb{R}^d$ with $\|\mathbf{x}\|_2 = 1$,

$$\begin{aligned} \mathbf{x}^\top (\mathbf{A}^{(T)})^\top \mathbf{A}^{(T)} \mathbf{x} &= \mathbf{x}^\top \mathbf{A}^\top \mathbf{A} \mathbf{x} + \sum_{t=1}^T (\mathbf{x}^\top \mathbf{a}^{(t)})^2 \\ &\leq \|\mathbf{x}\|_2^2 + \sum_{t=1}^T \|\mathbf{x}\|_2^2 \|\mathbf{a}^{(t)}\|_2^2 \leq 1 + \sum_{t=1}^T \frac{1}{d^4 \cdot T} \leq 2. \end{aligned}$$

Thus we have $\frac{1}{2} \leq \sigma_d(\mathbf{A}) \leq \sigma_d(\mathbf{A}^{(t)}) \leq \sigma_1(\mathbf{A}^{(t)}) \leq \sigma_1(\mathbf{A}^{(T)}) \leq 2$.

Therefore, the matrix $\mathbf{H}^{(t)} = ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1}$ satisfies $\frac{1}{4} \leq \lambda_d(\mathbf{H}^{(t)}) \leq \lambda_1(\mathbf{H}^{(t)}) \leq 4$.

Step 1 We prove $\mathbf{x}^{(t)} = \mathbf{x}_\star^{(t)} \pm \frac{1}{d^4\sqrt{T}}$. Since at each step we add a new label 1 to the vector \mathbf{b} , we have $\|\mathbf{b}^{(t)}\|_2 \leq \|\mathbf{b}^{(T)}\|_2 = \sqrt{T}$. Since $\mathbf{x}^{(t)}$ is an ϵ -approximate solution of the least squares regression problem, we have

$$\begin{aligned} (1 + \epsilon)^2 \|\mathbf{A}^{(t)} \mathbf{x}_\star^{(t)} - \mathbf{b}^{(t)}\|_2^2 &\geq \|\mathbf{A}^{(t)} \mathbf{x}^{(t)} - \mathbf{b}^{(t)}\|_2^2 \\ &= \|\mathbf{A}^{(t)} \mathbf{x}_\star^{(t)} - \mathbf{b}^{(t)}\|_2^2 + \|\mathbf{A}^{(t)} (\mathbf{x}^{(t)} - \mathbf{x}_\star^{(t)})\|_2^2 \\ &\geq \|\mathbf{A}^{(t)} \mathbf{x}_\star^{(t)} - \mathbf{b}^{(t)}\|_2^2 + \frac{1}{4} \|\mathbf{x}^{(t)} - \mathbf{x}_\star^{(t)}\|_2^2. \end{aligned}$$

The second step follows from $\mathbf{A}^{(t)} \mathbf{x}_\star^{(t)} - \mathbf{b}^{(t)} = (\mathbf{A}^{(t)} ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1} (\mathbf{A}^{(t)})^\top - \mathbf{I}) \cdot \mathbf{b}^{(t)} \in \ker[(\mathbf{A}^{(t)})^\top]$ is orthogonal to $\mathbf{A}^{(t)} (\mathbf{x}^{(t)} - \mathbf{x}_\star^{(t)}) \in \text{Im}[\mathbf{A}^{(t)}]$, the third step follows from $\sigma_d(\mathbf{A}^{(t)}) \geq \frac{1}{2}$.

Therefore, we conclude

$$\begin{aligned} \|\mathbf{x}^{(t)} - \mathbf{x}_\star^{(t)}\|_2^2 &\leq 4(2\epsilon + \epsilon^2) \|\mathbf{A}^{(t)} \mathbf{x}_\star^{(t)} - \mathbf{b}^{(t)}\|_2^2 \leq 12\epsilon \cdot \|\mathbf{b}^{(t)}\|_2^2 \leq 12\epsilon T \\ &\Rightarrow \|\mathbf{x}^{(t)} - \mathbf{x}_\star^{(t)}\|_2 < 4\sqrt{\epsilon T} = O\left(\frac{1}{d^4\sqrt{T}}\right). \end{aligned} \quad (11)$$

The last step follows from $\epsilon = \frac{1}{d^8 T^2}$ in the theorem statement.

Step 2 We prove $\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)} = \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \pm O(\frac{1}{d^4\sqrt{T}})$. By Woodbury identity, one has

$$\begin{aligned} \mathbf{x}_\star^{(t)} &= ((\mathbf{A}^{(t)})^\top \mathbf{A}^{(t)})^{-1} (\mathbf{A}^{(t)})^\top \mathbf{b}^{(t)} \\ &= ((\mathbf{A}^{(t-1)})^\top \mathbf{A}^{(t-1)} + \mathbf{a}^{(t)} (\mathbf{a}^{(t)})^\top)^{-1} \cdot ((\mathbf{A}^{(t-1)})^\top \mathbf{b}^{(t-1)} + \mathbf{a}^{(t)}) \\ &= \left(\mathbf{H}^{(t-1)} - \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \cdot (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} \cdot (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \right) \cdot ((\mathbf{A}^{(t-1)})^\top \mathbf{b}^{(t-1)} + \mathbf{a}^{(t)}) \\ &= \mathbf{x}_\star^{(t-1)} + \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \left(1 - (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} (\mathbf{a}^{(t)})^\top (\mathbf{x}_\star^{(t-1)} + \mathbf{H}^{(t-1)} \mathbf{a}^{(t)}) \right) \end{aligned} \quad (12)$$

The first step follows from $\mathbf{x}_\star^{(t)}$ is the optimal solution at step t , and the third step follows from the Woodbury identity and $((\mathbf{A}^{(t-1)})^\top \mathbf{A}^{(t-1)})^{-1} = \mathbf{H}^{(t-1)}$. We use $\mathbf{x}_\star^{(t-1)} = \mathbf{H}^{(t-1)} (\mathbf{A}^{(t-1)})^\top \mathbf{b}^{(t-1)}$ in the fourth step.

Consequently, we have

$$\begin{aligned} &\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)} \\ &= \mathbf{x}_\star^{(t)} - \mathbf{x}_\star^{(t-1)} \pm O\left(\frac{1}{d^4\sqrt{T}}\right) \\ &= \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} - \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} (\mathbf{a}^{(t)})^\top (\mathbf{H}^{(t-1)} \mathbf{a}^{(t)} + \mathbf{x}_\star^{(t-1)}) \pm O\left(\frac{1}{d^4\sqrt{T}}\right) \\ &= \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \pm O\left(\frac{1}{d^4\sqrt{T}}\right) \pm O\left(\frac{1}{d^4\sqrt{T}}\right) \\ &= \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \pm O\left(\frac{1}{d^4\sqrt{T}}\right). \end{aligned} \quad (13)$$

The first step comes from Eq. (11), the second step follows from Eq. (12), the third step follows from

$$\begin{aligned} &\|\mathbf{H}^{(t-1)} \mathbf{a}^{(t)} \cdot (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} \cdot (\mathbf{a}^{(t)})^\top \cdot (\mathbf{H}^{(t-1)} \mathbf{a}^{(t)} + \mathbf{x}_\star^{(t-1)})\|_2 \\ &\leq \|\mathbf{H}^{(t-1)}\| \cdot \|\mathbf{a}^{(t)}\|_2 \cdot (1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} \cdot \|(\mathbf{a}^{(t)})^\top\|_2 \cdot (\|\mathbf{H}^{(t-1)} \mathbf{a}^{(t)}\|_2 + \|\mathbf{x}_\star^{(t-1)}\|_2) \\ &\leq 4 \cdot \frac{1}{d^2\sqrt{T}} \cdot 1 \cdot \frac{1}{d^2\sqrt{T}} \cdot \left(\frac{4}{d^2\sqrt{T}} + 8\sqrt{T}\right) = O\left(\frac{1}{d^4\sqrt{T}}\right) \end{aligned}$$

where we use $\|\mathbf{H}^{(t-1)}\| \leq 4$, $\|\mathbf{a}^{(t)}\|_2 \leq \frac{1}{d^2\sqrt{T}}$, $(1 + (\mathbf{a}^{(t)})^\top \mathbf{H}^{(t-1)} \mathbf{a}^{(t)})^{-1} \leq 1$, $\|\mathbf{H}^{(t-1)} \mathbf{a}^{(t)}\|_2 \leq \|\mathbf{H}^{(t-1)}\| \|\mathbf{a}^{(t)}\|_2 \leq \frac{4}{d^2\sqrt{T}}$, $\|(\mathbf{A}^{(t)})^\top\| \leq 2$ and

$$\begin{aligned} \|\mathbf{x}_*^{(t-1)}\|_2 &= \mathbf{H}^{(t-1)} (\mathbf{A}^{(t-1)})^\top \mathbf{b}^{(t-1)} \\ &\leq \|\mathbf{H}^{(t-1)}\| \cdot \|(\mathbf{A}^{(t-1)})^\top\| \cdot \|\mathbf{b}^{(t-1)}\|_2 \leq 4 \cdot 2 \cdot \|\mathbf{b}^{(t-1)}\|_2 \leq 8\sqrt{T}. \end{aligned}$$

Step 3 We prove $\mathbf{H}^{(t-1)} \mathbf{a}^{(t)} = \mathbf{H} \mathbf{a}^{(t)} \pm O(\frac{1}{d^4\sqrt{T}})$. Denote $\mathbf{U} = [\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(t-1)}] \in \mathbb{R}^{d \times t}$, we have $\|\mathbf{U}\| = \|\mathbf{U}^\top\| \leq \|\mathbf{U}\|_F \leq \frac{1}{d^2}$ since $\|\mathbf{a}^{(i)}\|_2 \leq \frac{1}{d^2\sqrt{T}}$ for all $i \in [t-1]$. Then we have that

$$\begin{aligned} \mathbf{H}^{(t-1)} \mathbf{a}^{(t)} &= (\mathbf{A}^\top \mathbf{A} + \mathbf{U} \mathbf{U}^\top)^{-1} \mathbf{a}^{(t)} \\ &= (\mathbf{H} - \mathbf{H} \mathbf{U} (\mathbf{I} + \mathbf{U}^\top \mathbf{H} \mathbf{U})^{-1} \mathbf{U}^\top \mathbf{H}) \cdot \mathbf{a}^{(t)} \\ &= \mathbf{H} \mathbf{a}^{(t)} - \mathbf{H} \mathbf{U} (\mathbf{I} + \mathbf{U}^\top \mathbf{H} \mathbf{U})^{-1} \mathbf{U}^\top \mathbf{H} \mathbf{a}^{(t)} \\ &= \mathbf{H} \mathbf{a}^{(t)} \pm O\left(\frac{1}{d^6\sqrt{T}}\right). \end{aligned} \tag{14}$$

The second step follows from the Woodbury identity and $\mathbf{H} = (\mathbf{A}^\top \mathbf{A})^{-1}$. The fourth step follows from

$$\begin{aligned} \|\mathbf{H} \mathbf{U} (\mathbf{I} + \mathbf{U}^\top \mathbf{H} \mathbf{U})^{-1} \mathbf{U}^\top \mathbf{H} \mathbf{a}^{(t)}\|_2 &\leq \|\mathbf{H}\| \cdot \|\mathbf{U}\| \cdot \|(\mathbf{I} + \mathbf{U}^\top \mathbf{H} \mathbf{U})^{-1}\| \cdot \|\mathbf{U}^\top\| \cdot \|\mathbf{H}\| \cdot \|\mathbf{a}^{(t)}\|_2 \\ &\leq 4 \cdot \frac{1}{d^2} \cdot 1 \cdot \frac{1}{d^2} \cdot 4 \cdot \frac{1}{d^2\sqrt{T}} = \frac{16}{d^6\sqrt{T}}, \end{aligned}$$

as $\|\mathbf{H}\| \leq 4$, $\|\mathbf{U}\| = \|\mathbf{U}^\top\| \leq \|\mathbf{U}\|_F \leq \frac{1}{d^2}$, $\|(\mathbf{I} + \mathbf{U}^\top \mathbf{H} \mathbf{U})^{-1}\| \leq 1$, and $\|\mathbf{a}^{(t)}\|_2 \leq \frac{1}{d^2\sqrt{T}}$.

Combining three steps We conclude that

$$\|d^2\sqrt{T}(\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}) - \mathbf{H} \mathbf{z}^{(t)}\|_2 = d^2\sqrt{T} \|(\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}) - \mathbf{H} \mathbf{a}^{(t)}\|_2 = O(1/d^2)$$

where the first step follows from $\mathbf{a}^{(t)} = \frac{1}{d^2\sqrt{T}} \cdot \mathbf{z}^{(t)}$, and the second step follows from Eq. (13) and Eq. (14). Hence one can recover the matrix-vector query from solutions of dynamic least squares regression. On the other side, the reduction only takes $O(d)$ time per update. Hence, we have shown a $\Omega(d^2)$ lower bound on the amortized running time for dynamic least squares regression problem under the OMv conjecture. \square

E DETAILS OF EXPERIMENTS

We provide more details of the setting of our experiments in Section 5.

Synthetic dataset The scaling parameters $w^{(t)}$ of the elliptical model are chosen as follows: After the initial phase (the first 10% of the data), we randomly choose $d/10$ rows among the next 10% of data to have a large scaler of \sqrt{T} . The rest of the data have a scaler of 1.

Parameters In both our method and the row sampling method, we use an error parameter ϵ . We set the sampling probability to be $p = \min\{\tau\epsilon^{-2}/2, 1\}$ (except when $\epsilon = 1$ we set $p = \min\{\tau, 1\}$ to make it non-trivial), where τ is the approximate online leverage score of different methods. The JL matrix in our algorithm has k number of rows, where we set $k = c_\epsilon \cdot \epsilon^{-2}$ for some constants c_ϵ so that $k \approx 20$.

The raw data of our experiments (Figure 1) are shown in Table 1.

Dataset	Method	Error	Time	Other*
Synthetic	Kalman	1	224.7s	
Synthetic	ours	1.42	3.85s	$\epsilon = 1$
Synthetic	ours	1.22	8.6s	$\epsilon = 0.5$
Synthetic	ours	1.03	25.1s	$\epsilon = 0.2$
Synthetic	ours	1.009	82.0s	$\epsilon = 0.1$
Synthetic	row sampling	1.38	36.39s	$\epsilon = 1$
Synthetic	row sampling	1.23	42.9s	$\epsilon = 0.5$
Synthetic	row sampling	1.039	51.6s	$\epsilon = 0.2$
Synthetic	row sampling	1.009	85.9s	$\epsilon = 0.1$
Synthetic	uniform	56.2	11.1s	$p = 0.05$
Synthetic	uniform	39.8	22.3s	$p = 0.1$
Synthetic	uniform	39.5	44.9s	$p = 0.2$
Synthetic	uniform	20.4	117.3s	$p = 0.5$
VirusShare	Kalman	1	45.3s	
VirusShare	ours	1.24	1.74s	$\epsilon = 1$
VirusShare	ours	1.18	2.50s	$\epsilon = 0.5$
VirusShare	ours	1.03	6.73s	$\epsilon = 0.2$
VirusShare	ours	1.005	17.5s	$\epsilon = 0.1$
VirusShare	row sampling	1.25	7.35s	$\epsilon = 1$
VirusShare	row sampling	1.15	7.94s	$\epsilon = 0.5$
VirusShare	row sampling	1.02	12.0s	$\epsilon = 0.2$
VirusShare	row sampling	1.005	18.1s	$\epsilon = 0.1$
VirusShare	uniform sampling	1.2153e+05	2.62s	$p = 0.05$
VirusShare	uniform sampling	9.7335e+03	5.23s	$p = 0.1$
VirusShare	uniform sampling	2.3363e+03	9.66s	$p = 0.2$
VirusShare	uniform sampling	23.1	23.7s	$p = 0.5$

Table 1: Experiment results.