# Optimal Algorithms for Linear Algebra in the Current Matrix Multiplication Time 

Yeshwanth Cherapanamjeri* Sandeep Silwal $^{\dagger}$ David P. Woodruff ${ }^{\ddagger}$ Samson Zhou ${ }^{\S}$


#### Abstract

We study fundamental problems in linear algebra, such as finding a maximal linearly independent subset of rows or columns (a basis), solving linear regression, or computing a subspace embedding. For these problems, we consider input matrices $\mathbf{A} \in \mathbb{R}^{n \times d}$ with $n>d$. The input can be read in $n n z(\mathbf{A})$ time, which denotes the number of nonzero entries of $\mathbf{A}$. In this paper, we show that beyond the time required to read the input matrix, these fundamental linear algebra problems can be solved in $d^{\omega}$ time, i.e., where $\omega \approx 2.37$ is the current matrix-multiplication exponent.

To do so, we introduce a constant-factor subspace embedding with the optimal $m=\mathcal{O}(d)$ number of rows, and which can be applied in time $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{poly}(\log d)$ for any trade-off parameter $\alpha>0$, tightening a recent result by Chepurko et. al. [SODA 2022] that achieves an $\exp (\operatorname{poly}(\log \log n))$ distortion with $m=d \cdot \operatorname{poly}(\log \log d)$ rows in $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+d^{2+\alpha+o(1)}\right)$ time. Our subspace embedding uses a recently shown property of stacked Subsampled Randomized Hadamard Transforms (SRHT), which actually increase the input dimension, to "spread" the mass of an input vector among a large number of coordinates, followed by random sampling. To control the effects of random sampling, we use fast semidefinite programming to reweight the rows. We then use our constant-factor subspace embedding to give the first optimal runtime algorithms for finding a maximal linearly independent subset of columns, regression, and leverage score sampling. To do so, we also introduce a novel subroutine that iteratively grows a set of independent rows, which may be of independent interest.


## 1 Introduction

In this paper, we consider fundamental problems in linear algebra, such as finding a maximal linearly independent subset of rows or columns, i.e., a basis, or solving linear regression, or computing a subspace embedding. Surprisingly, we still do not have optimal algorithms for these tasks. The input to these problems is generally a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with $n>d$ that requires nnz( $\mathbf{A}$ ) time to read, where $\mathrm{nnz}(\mathbf{A})$ is the number of non-zero entries of $\mathbf{A}$. Algorithms for these problems frequently use subroutines such as matrix multiplication, inverse computation, or decomposition (singular value, QR, LU, etc.), that use at least $n d^{\omega-1}$ time, where $\omega \approx 2.37$ is the exponent for matrix multiplication [AW21].

[^0]Dimensionality reduction techniques are often utilized to decrease the effective input size, so that a solution to the smaller input is often a good approximation to the optimal solution of the original problem. These approaches transform $\mathbf{A}$ into a matrix $\mathbf{M} \in \mathbb{R}^{m \times d}$ with $m \ll n$ and (approximately) solve the problem on the instance $\mathbf{M}$ with a significantly smaller number of rows. However, existing results could only achieve $m=d$ polylog $(d)$ for dimensionality reduction techniques with inputsparsity runtime, which prevented true matrix-multiplication runtime algorithms, i.e., running times of the form $\mathcal{O}\left(d^{\omega}\right)$.

Here we emphasize that $\omega$ is the parameter between 2 and 3 for the matrix multiplication exponent, possibly depending on the input parameters for matrix multiplication, e.g., matrix multiplication between two $n \times n$ matrices uses $\mathcal{O}\left(n^{\omega}\right)$ time, for some fixed matrix multiplication oracle that we are given. By contrast, many previous works define $\omega$ to be the smallest constant such that matrix multiplication between two $n \times n$ matrices runs in time $\mathcal{O}\left(n^{\omega+\varepsilon}\right)$ for any constant $\varepsilon>0$. In particular, [CW82] showed that given a matrix multiplication algorithm with runtime $\mathcal{O}\left(n^{\omega+\varepsilon_{1}}\right)$, there exists a matrix multiplication algorithm with runtime $\mathcal{O}\left(n^{\omega+\varepsilon_{2}}\right)$ with $\varepsilon_{2} \in\left(0, \varepsilon_{1}\right)$ and this process can continue ad infinitum. However, at some point we require an explicit fixed matrix multiplication algorithm for downstream applications. Thus, we consider access to a fixed matrix multiplication algorithm with runtime $\mathcal{O}\left(n^{\omega}\right)$, so that it is important to track and eliminate additional polylog overheads on top of the matrix multiplication runtime of the fixed algorithm. Indeed, the removal of the last logarithmic factors is related to well-known conjectures on the construction of sparse Johnson-Lindenstrauss transforms [NN13, CCKW22]. In a recent work, [CCKW22] showed that these logarithmic factors were not necessary, achieving algorithms for linear algebra in near matrix-multiplication runtime, up to poly $(\log \log d)$ factors.

### 1.1 Our Contribution

In this work, we give the first algorithms for linear algebra in true matrix-multiplication runtime, removing the last poly $(\log \log d)$ factors in the algorithms of [CCKW22] and thus closing a long line of work. Our results show that beyond the time required to read the input matrix, fundamental linear algebra problems such as finding a maximal linearly independent subset of rows or columns (a basis), linear regression, or computing a subspace embedding can be solved in the current matrixmultiplication runtime.

We first introduce a constant-factor subspace embedding that uses input-sparsity runtime:
Theorem 1.1. For any $\mathbf{A} \in \mathbb{R}^{n \times d}$ and any tradeoff parameter $\alpha>0$, we can compute matrix $\mathbf{G} \in \mathbb{R}^{p \times n}$ such that:

$$
\forall \mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{G} \mathbf{A} \mathbf{x}\|_{2} \leq \xi\|\mathbf{A} \mathbf{x}\|_{2}
$$

with probability at least 0.9 for a fixed constant $\xi>1$. Furthermore, we have $p=\mathcal{O}(d)$ and $\mathbf{G A}$ may be computed in time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{polylog}(d)$.

By comparison, [CCKW22] recently gave a subspace embedding with distortion $\exp (\operatorname{poly}(\log \log n))$ using runtime $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+d^{2+\alpha+o(1)}\right)$, for any tradeoff parameter $\alpha>0$. Theorem 1.1 also extends naturally to the case when $\mathbf{A}$ has rank $k$, in which case it suffices for $\mathbf{G} \in \mathbb{R}^{p \times n}$ to have $p=\mathcal{O}(k)$ rows and the resulting time to compute $\mathbf{G A}$ is $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}\right)+k^{2+\alpha} \operatorname{polylog}(k)$.

Our constant-factor subspace embedding in Theorem 1.1 uses a sampling procedure that leverages a recently observed property of stacked Subsampled Randomized Hadamard Transforms (SRHTs)
to "spread" the mass of an input vector among a large number of coordinates [CN22]. Our constantfactor subspace embedding can be used to improve the efficiency of leverage score sampling, which has applications in a number of important linear algebra problems [Woo14].

In particular, we can further boost our constant-factor subspace embedding to a $(1+\varepsilon)$ approximate subspace embedding through leverage score sampling:
Theorem 1.2. Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, an accuracy parameter $\varepsilon>0$, and any tradeoff parameter $\alpha>0$, there exists an algorithm that computes a matrix SA with $\mathcal{O}\left(\frac{1}{\varepsilon^{2}} d \log d\right)$ rows such that with probability at least $\frac{9}{10}$, for all vectors $\mathbf{x} \in \mathbb{R}^{d}$,

$$
(1-\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{S A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} .
$$

Moreover, SA can be computed in time

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon^{2}} d^{2+\alpha} \text { polylog }(d)
$$

By comparison, recent work by [CCKW22] achieved a $(1+\varepsilon)$-subspace embedding with either $\frac{1}{\varepsilon^{2}}(d \log d) \exp (\operatorname{poly}(\log \log d))$ rows or with runtime $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}\right)+d^{\omega} \operatorname{poly}(\log \log d)+\frac{1}{\varepsilon^{3}} d^{2+o(1)}+$ $\frac{1}{\varepsilon^{2}} n^{\alpha+o(1)} d^{2+o(1)}$. Our result avoids such tradeoffs, which is especially useful in downstream applications, as we soon discuss. Moreover, Theorem 1.2 extends naturally to the case where A has rank $k$, similarly as Theorem 1.1. On the other hand, we remark that unlike our constant-factor subspace embedding, our $(1+\varepsilon)$-subspace embedding does not have the optimal number of rows to perform further tasks downstream. We believe the existence/construction of such a subspace embedding would be an interesting future question. Questions in a similar spirit have also been previously asked for graph theoretic problems, e.g., [LS17, LS18].

We then use our constant-factor subspace embedding and our leverage score sampling framework to find a maximal set of linearly independent rows of an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ :
Theorem 1.3. Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$ and any tradeoff parameter $\alpha>0$, there exists an algorithm that outputs a set of $k$ linearly independent rows of $\mathbf{A}$, using time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+k^{\omega}\right)+$ $k^{2+\alpha}$ polylog $(k)$

By comparison, recent work of [CCKW22] gave an algorithm that finds a set of $k$ linearly independent rows of $\mathbf{A}$ using time $\mathcal{O}\left(\mathrm{nnz}(\mathbf{A})+k^{2+o(1)}\right)+k^{\omega}$ poly $(\log \log k)$. Like the algorithm of [CCKW22], we first reduce the problem to computing a set of $k$ linearly dependent rows of a matrix $\mathbf{B} \in \mathbb{R}^{\mathcal{O}(k \log k) \times \mathcal{O}(k)}$, though due to our more efficient subspace embedding algorithm, we can do this in matrix-multiplication runtime while [CCKW22] cannot. Now to achieve Theorem 1.3, we develop a novel subroutine in Section 4.1, which iteratively grows a set of independent rows of B that may be of independent interest. Crucially, the algorithm avoids additional $k^{\omega}$ poly $(\log \log k)$ dependencies that are incurred by the subroutine of [CCKW22]. We provide a summary of previous work on finding a set of independent rows in Figure 1.

Finally, we use our $(1+\varepsilon)$-approximate subspace embedding to achieve $(1+\varepsilon)$-approximate linear regression:
Theorem 1.4. Given $\mathbf{A} \in \mathbb{R}^{n \times d}, \mathbf{b} \in \mathbb{R}^{n}$, and any tradeoff parameter $\alpha>0$, there exists an algorithm that with probability at least 0.9 , outputs a vector $\mathbf{y}$ such that

$$
\|\mathbf{A} \mathbf{y}-\mathbf{b}\|_{2} \leq(1+\varepsilon) \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A x}-\mathbf{b}\|_{2}
$$

| Reference | Runtime |
| :---: | :---: |
| Gaussian elimination | $\mathcal{O}\left(n d^{\omega-1}\right)$ |
| [CKL13] | $\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \log n+k^{\omega} \log n\right)$ |
| [CCKW22] | $\mathcal{O}\left(\mathrm{nnz}(\mathbf{A})+k^{2+o(1)}\right)+k^{\omega} \operatorname{poly}(\log \log k)$ |
| Theorem 1.3, for any $\alpha>0$ | $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+k^{\omega}\right)+k^{2+\alpha} \operatorname{polylog}(k)$ |

Fig. 1: Summary of previous results for identifying a set of $k$ independent rows from a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$
using time $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon} d^{2+\alpha} \operatorname{polylog}(d)+\frac{1}{\varepsilon} d^{2} \operatorname{poly} \log (d) \log \frac{1}{\varepsilon}$.
By comparison, [CCKW22] output a $(1+\varepsilon)$-approximation to linear regression in time $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+$ $\frac{1}{\varepsilon^{3}} n^{\alpha+o(1)} d^{2+o(1)}+d^{\omega}$ poly $(\log \log d)$ for any tradeoff parameter $\alpha>0$. We remark that our result achieves both the optimal (current) matrix-multiplication runtime and also a better dependence on $\frac{1}{\varepsilon}$, which is a byproduct of our $(1+\varepsilon)$-subspace embedding avoiding the tradeoffs incurred by [CCKW22].

### 1.2 Technical Overview

In this section, we give a brief overview of our technical contributions. A summary of the interplay between our algorithmic contributions can be seen in Figure 3.

To avoid polylogarithmic overhead over matrix-multiplication runtime, we require a dimensionality reduction technique that uses $o(d \log d)$ rows. A folklore result states that a dense matrix $\mathbf{G}$ of $\mathcal{O}(d)$ rows with entries that are independent Sub-Gaussian random variables suffices to achieve a constant-factor subspace embedding [Woo14]. However, computing GA for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ requires time $\mathcal{O}(d \operatorname{nnz}(\mathbf{A}))$ due to the multiplication with the dense matrix $\mathbf{G}$. Although the runtime is unsatisfactory, a subspace embedding with a small number of rows corresponds to an improvement in runtime for downstream tasks. Unfortunately, faster dimensionality reduction techniques such as the sparse Johnson-Lindenstrauss transform use $\Omega(d \log d)$ rows and it is an open question whether these constructions can be improved to only using $\mathcal{O}(d)$ rows [NN13].

Intuition from previous work. [CCKW22] recently sidestepped these issues by first showing that a rescaling of an embedding of [Ind07] roughly maintains the $\ell_{2}$ norm of a $d$-dimensional unit $\ell_{2}$ vector while the $\ell_{1}$ norm becomes $\tilde{\Omega}(\sqrt{d})$. In particular, a constant fraction of the resulting $o(d \log d)$ coordinates have magnitude $\tilde{\Omega}\left(\frac{1}{\sqrt{d}}\right)$. [CCKW22] then used the intuition that under such a "flattening" $\mathbf{y}$ of a unit $\ell_{2}$ vector, a sparse matrix $\mathbf{S}$ of random signs will sample some of the coordinates with "large" magnitude, so that the dot product $\left\langle\mathbf{S}_{i}, \mathbf{y}\right\rangle$ of each row of $\mathbf{S}$ with the flattened vector $\mathbf{y}$ will be at least $\tilde{\Omega}\left(\frac{1}{\sqrt{d}}\right)$, which implies a lower bound on $\|\mathbf{S y}\|_{2}^{2}$. An upper bound on the operator norm $\|\mathbf{S}\|_{2}$ can also be shown, ultimately giving a distortion of $\exp (\operatorname{poly}(\log \log d))$. Unfortunately, Indyk's embedding [Ind07] only governs the sum of the $\ell_{2}$ norms of blocks of coordinates of the resulting embedding and thus must be applied recursively across $\mathcal{O}(\log \log d)$ levels, resulting in a sketching matrix with $d \exp (\log \log (d))$ rows. Therefore, the resulting matrix does not lose polylogarithmic factors over matrix-multiplication runtime, but still cannot quite achieve true matrix-multiplication runtime. Hence, using Indyk's embedding [Ind07] seems to be a major
bottleneck to achieving matrix multiplication runtime in the previous works and thus it seems we need to use significantly new techniques altogether.

Crude subspace embeddings through SRHT. To avoid the extraneous factors over matrixmultiplication runtime, we require a sketching matrix with a smaller number of rows which can also be applied quickly. To that end, we recall that the flattening property of the Subsampled Randomized Hadamard Transform (SRHT) is frequently used to show that it forms a subspace embedding, in the sense that the largest coordinate of the image of a fixed vector is upper bounded by $\tilde{\mathcal{O}}\left(\frac{1}{\sqrt{d}}\right)$. More recently, [CN22] showed that by stacking some number of SRHTs on top of each other to increase dimension, not only is the $\ell_{2}$-norm of any fixed vector preserved exactly, but also for any vector, a constant fraction of the coordinates of the image of the SRHT is lower bounded by $\tilde{\Omega}\left(\frac{1}{\sqrt{d}}\right)$ in absolute value. This property can be seen as a fast embedding of $\ell_{2}$ into $\ell_{1}$ with a small target dimension.

We first left-multiply our input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with OSNAP matrices [NN13] $\mathbf{S}_{1}$ and $\mathbf{S}_{2}$ to obtain a constant-factor subspace embedding. The application of the composition of OSNAP matrices is a standard technique that allows us to first decrease the number of rows to $\ell:=$ $\mathcal{O}(d \log d)$ at the cost of a constant-factor distortion, in $\mathcal{O}\left(\operatorname{nnz}(\mathbf{A})+d^{C}\right)$ time, where $C$ can be made an arbitrarily small constant larger than 2 .

If $\mathbf{M} \in \mathbb{R}^{m \ell \times \ell}$ is the matrix consisting of stacked randomized Hadamard matrices, we define $\mathbf{S} \in \mathbb{R}^{\mathcal{O}(d) \times m \ell}$ to be a matrix that independently and uniformly samples each row of $\mathbf{M} \mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$. It then suffices to bound both the contraction and the dilation of $\mathbf{S M B}:=\mathbf{S M} \mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$. That is, if we could show $\|\mathbf{S M}\|_{2} \leq \mathcal{O}(1)$ and $\|\mathbf{S M B x}\|_{2} \geq \mathcal{O}(1)$ for all unit vectors $\mathbf{B x} \in \mathbb{R}^{\ell}$, then it follows that $\mathbf{S M}$ is a constant factor subspace embedding. To handle contraction, we show that $\|\mathbf{S M B x}\|_{2}^{2} \geq \mathcal{O}(1)$ by first showing concentration for a single vector $\mathbf{x} \in \mathbb{R}^{\ell}$ due to the abundance of "large" coordinates in $\mathbf{M B x}$, followed by taking a union bound over a sufficiently fine net. Unfortunately, it does not seem evident how to bound the dilation by a constant (or whether it is even true). For instance, using either a crude concentration inequality or more sophisticated results bounding the norms of random submatrices, e.g., [Tro08], seems to give an extra logarithmic factor. Thus this approach yields a subspace embedding with a logarithmic distortion, which is not enough for our optimal algorithms for downstream tasks.

Constant-factor subspace embedding. To achieve a constant-factor subspace embedding, we note that the slack in the above analysis is that there are large entries in $\mathbf{M}$ that can be sampled by $\mathbf{S}$, which prevents constant upper bounds on $\|\mathbf{S M}\|_{2}$. On the other hand, due to the abundance of large coordinates in $\mathbf{M B x}$, an accurate estimation can still be acquired without these large entries in $\mathbf{M}$. Indeed, we show that with high probability, there exists a slight reweighting $\mathbf{W}$ of the sampled rows (possibly with weight 0 to remove the large entries) such that the resulting reweighted subsampled matrix has operator bounded by a constant, i.e., $\|\mathbf{W S M}\|_{2} \leq \mathcal{O}(1)$.

Moreover, we show that the reweighting can be efficiently computed by solving a standard packing semidefinite program (SDP). Crucially, the SDP requires a fast projection oracle, which we can implement due to our subspace embedding with a logarithmic distortion discussed above. Finally, we show that under this reweighting, the contraction does not drastically change, so that $\|\mathbf{W S M B x}\|_{2} \geq \mathcal{O}(1)$ for all unit vectors $\mathbf{B x} \in \mathbb{R}^{n}$. Hence, it follows that $\mathbf{W S M B}$ is a constantfactor subspace embedding.

Given input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ :
(1) Apply OSNAP matrix $\mathbf{S}_{2} \in \mathbb{R}^{n^{\prime} \times n}$ with constant factor distortion and tradeoff-parameter $\alpha$ to acquire $\mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{n^{\prime} \times d}$, where $n^{\prime}=\mathcal{O}\left(d^{1+\alpha} \log d\right)$
(2) Apply OSNAP matrix $\mathbf{S}_{1} \in \mathbb{R}^{\ell \times n^{\prime}}$ with constant factor distortion and tradeoff-parameter $\alpha^{\prime}=\frac{1}{\log d}$ to acquire $\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{\ell \times d}$, where $\ell=\mathcal{O}(d \log d)$
(3) Apply SRHT matrix $\mathbf{M} \in \mathbb{R}^{m \ell \times \ell}$ for $m=\operatorname{polylog}(d)$ to acquire $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{m \ell \times d}$
(4) Apply a sampling matrix $\mathbf{S} \in \mathbb{R}^{p \times m \ell}$ that uniformly samples rows with $p=\mathcal{O}(d)$, to acquire $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{d^{\prime} \times d}$, where $d^{\prime}=\mathcal{O}(d)$ with high probability
(5) Solve an SDP to find a set of weights $\mathbf{W} \in \mathbb{R}^{d^{\prime} \times d^{\prime}}$ so that the operator norm of $\mathbf{S M}$ is appropriately bounded and output $\mathbf{W S M S}_{1} \mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{d^{\prime} \times d}$

Fig. 2: High-level summary of our constant-factor subspace embedding

We remark that we require sharper bounds in downstream applications, e.g., basis selection, when $\mathbf{A} \in \mathbb{R}^{n \times d}$ is not full rank. In this case, our algorithms naturally generalize to dimension and runtime dependent on the rank of $\mathbf{A}$ rather than the dimension $d$ of $\mathbf{A}$. We summarize our constant-factor subspace embedding at a high level in Figure 2.

Leverage score sampling. To achieve a $(1+\varepsilon)$-subspace embedding for a matrix $\mathbf{A}$, a standard approach is to perform leverage score sampling, i.e., to sample $\mathcal{O}\left(\frac{1}{\varepsilon^{2}} d \log d\right)$ rows of $\mathbf{A}$ with probabilities proportional to their leverage score. However existing techniques could not be run in matrix-multiplication runtime, and instead ran in at least $\mathcal{O}\left(d^{\omega} \log d\right)$ time.

We instead use our fast constant-factor subspace embedding $\mathbf{S A}$ into an optimal target dimension. We can also efficiently compute its QR decomposition so that $\mathbf{S A}=\mathbf{Q R} \mathbf{R}^{-1}$ for a matrix $\mathbf{Q}$ with orthonormal columns. In particular, since $\mathbf{Q}$ has orthonormal columns, then the leverage scores of $\mathbf{S A}$ are precisely the squared row norms of $\mathbf{Q}$. It follows that the squared row norm of $\mathbf{a}_{i} \mathbf{R}$ is a constant-factor approximation to the leverage score of row $\mathbf{a}_{i}$ for each $i \in[n]$. We can then apply the standard leverage score sampling approach by sampling $\mathcal{O}\left(\frac{1}{\varepsilon^{2}} d \log d\right)$ rows of $\mathbf{A}$ with probabilities proportional to their leverage score to achieve a $(1+\varepsilon)$-subspace embedding in matrix-multiplication runtime. In particular, we first compute ARG for a Johnson-Lindenstrauss matrix $\mathbf{G}$ and then perform rejection sampling to achieve leverage score sampling (see Theorem 3.2 and surrounding discussion for more details). We also remark that we only require the orthogonality of $\mathbf{Q}$ in the QR decomposition, so other methods such as the SVD decomposition that yield a matrix with orthonormal columns would also suffice.

We again emphasize that unlike our constant-factor subspace embedding, our $(1+\varepsilon)$-subspace embedding does not have the optimal number of rows to perform further tasks downstream. That is, our $(1+\varepsilon)$-subspace embedding uses $\mathcal{O}\left(\frac{1}{\varepsilon^{2}} d \log d\right)$ rows. By comparison, our constant-factor subspace embedding uses $\mathcal{O}(d)$ rows, which is better for $\varepsilon=\mathcal{O}(1)$. Thus our result should be interpreted as the ability to perform leverage score sampling in matrix-multiplication runtime, with one such application being a $(1+\varepsilon)$-subspace embedding and another such application being
the selection of an independent basis (see below). An interesting open question is whether our techniques can be further refined to achieve a $(1+\varepsilon)$-subspace embedding with $\mathcal{O}\left(\frac{d}{\varepsilon^{2}}\right)$ rows in matrix-multiplication runtime.

Basis selection. To find a set of $k$ independent rows for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$, we first use our efficient leverage score sampling framework in conjunction with existing techniques to reduce the effective input to size $\mathbb{R}^{\mathcal{O}(k \log (k)) \times \mathcal{O}(k)}$. Namely, we first note that there exists a distribution of matrices that form a rank-preserving sketch, so that $\operatorname{rank}(\mathbf{A})=\operatorname{rank}(\mathbf{A S})$, where $\mathbf{S} \in \mathbb{R}^{d \times c k}$ for some constant $c>0$. Moreover the rank-preserving sketch has the property that any set of independent rows of $\mathbf{A}$ is also a set of independent rows of $\mathbf{A S}$ and vice versa. Thus it would suffice to select a basis of rows from AS and take the corresponding rows in $\mathbf{A}$.

However, we cannot explicitly compute AS. We also cannot use our constant-factor subspace embedding, because multiplication by a Hadamard matrix distorts the mapping of the indices of independent rows in the original matrix and the indices of independent rows in the smaller matrix. Instead, we use our constant-factor subspace embedding, which selects $\mathcal{O}(k \log k)$ reweighted rows from AS. Hence, it remains to select a basis of rows from a matrix $\mathbf{B} \in \mathbb{R}^{\mathcal{O}(k \log (k)) \times \mathcal{O}(k)}$ with rank $k$.

For this sub-problem, there exist a number of previous techniques, such as an approach by [CCKW22] that iteratively removes redundant rows from B. We remark that these techniques are generally not optimized to run in time $\mathcal{O}\left(k^{\omega}\right)$ since other components are usually a larger bottleneck, and thus it seems we require a new set of techniques.

Iteratively growing a basis. To select a basis of rows from a matrix $\mathbf{B} \in \mathbb{R}^{\mathcal{O}(k \log (k)) \times \mathcal{O}(k)}$ with rank $k$, we develop a new algorithm that iteratively grows a set $S$ of independent rows of $\mathbf{B}$. Namely, we use leverage score sampling to sample $\mathcal{O}(k)$ rows of $\mathbf{B}$. Observe that this is not enough to cover the entire row span of $\mathbf{B}$, but for $c=\frac{1}{10}$, we show using approximate matrix product on a well-conditioned verion of our input, that we can get rank $(1-c) k=\frac{9}{10} k$ with probability at least $\frac{2}{3}$. We can add an independent subset of these rows to our growing set $S$ and then compute a basis $\mathbf{Z}_{1}$ for the orthogonal complement of $S$. So far, these procedures, i.e., leverage score sampling, independent subset selection, and orthogonal complement basis computation, all use at most $\gamma k^{\omega}$ runtime for an absolute constant $\gamma>0$.

We now repeat these procedures on $\mathbf{B} \mathbf{Z}_{1}^{\top}$, first using leverage score sampling to sample $\mathcal{O}(k)$ rows of $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$. An observation is that the rows of $\mathbf{B}$ that are spanned by $S$ will all be zero in $\mathbf{B Z}{ }_{1}^{\top}$, since $\mathbf{Z}_{1}$ is the orthogonal complement of $S$. Thus again with probability $\frac{2}{3}$, we can sample a set of rows with rank at least a $\frac{9}{10}$ fraction of the rank of $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$. We can again add an independent subset of these rows to our growing set $S$ and then compute a basis $\mathbf{Z}_{2}$ for the orthogonal complement of $S$. However since conditioned on the success of the previous iteration, the rank of $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$ is at most a $c=\frac{1}{10}$-fraction of the rank of $\mathbf{B}$, these procedures will now take at most $\gamma(c k)^{\omega}$ runtime, which is a constant fraction smaller.

We can thus proceed by iteratively adding rows of B to $S$ until $S$ has rank $k$. We can then output the corresponding rows of $\mathbf{A}$. The runtime in each iteration, conditioned on successful samplings in each previous iteration, follows a geometric series and thus the overall runtime is $\mathcal{O}\left(k^{\omega}\right)$. The runtime analysis is also robust to failures in each iteration because a failure in an iteration means that at worst, no additional rows are added to $S$. Therefore, the algorithm will always terminate with a set of independent rows and we can simply compute the expected runtime
of this procedure, which still follows a geometric series.


Fig. 3: Flowchart of dependencies for our algorithmic contributions.

Linear regression. For linear regression, we would like to find a vector y such that

$$
\|\mathbf{A y}-\mathbf{b}\|_{2} \leq(1+\varepsilon) \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A x}-\mathbf{b}\|_{2} .
$$

We can first compute a $(1+\varepsilon)$-approximate subspace embedding ${ }^{1}[\mathbf{S A} ; \mathbf{S b}]$ of the matrix $[\mathbf{A} ; \mathbf{b}]$. However, the subspace embedding $[\mathbf{S A} ; \mathbf{S b}]$ has $\frac{1}{\varepsilon^{2}} d \operatorname{polylog}(d)$ rows and so we cannot directly solve for the optimal solution on the smaller space, since computing the closed-form solution would not be true matrix-multiplication runtime. On the other hand, we only require finding an approximately optimal solution on the smaller space, i.e., we only want a vector $\mathbf{w} \in \mathbb{R}^{d}$ such that $\|\mathbf{S A w}-\mathbf{S b}\|_{2} \leq$ $(1+\mathcal{O}(\varepsilon)) \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A x}-\mathbf{S b}\|_{2}$. Thus we instead use gradient descent to find such a vector $\mathbf{w}$.

For efficient runtime, gradient descent requires a small condition number and a "good" initial solution. While SA is a $(1+\varepsilon)$-approximate subspace embedding, it may not necessarily have small condition number. To decrease the condition number to $\mathcal{O}(1)$, we instead consider $\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A R x}-\mathbf{S b}\|_{2}$, where $\mathbf{G A}=\mathbf{Q R} \mathbf{R}^{-1}$ is a QR decomposition for a constant-factor subspace embedding GA. Thus in this setting, $\mathbf{R}$ can be considered as a preconditioner. To find a good initial solution, we first find the closed-form solution to $\mathbf{w}^{(0)}=\operatorname{argmin}_{\mathbf{x}_{\in} \in \mathbb{R}^{d}}\|\mathbf{G A x}-\mathbf{G b}\|_{2}$, since GA is a constant-factor subspace embedding. It then remains to account for the preconditioning by computing $\mathbf{w}^{(1)}=\mathbf{R}^{-1} \mathbf{w}^{(0)}$, which is a good initial solution for gradient descent, since it provides a constant-factor approximation to the optimal solution due to the properties of GA.

Instead, we use the SRHT to compute a matrix $\mathbf{G}$ that is a constant-factor subspace embedding, with $\mathcal{O}(d)$ rows. By computing $\mathbf{G A}=\mathbf{Q R}^{-1}$, we can compute a matrix $\mathbf{R}$ such that $\kappa(\mathbf{A R})=$ $\mathcal{O}(1)$, since $\mathbf{G}$ is a constant-factor subspace embedding and $\kappa(\mathbf{G A R})=1$ due to the orthogonality of $\mathbf{Q}$. This implies that an approximate minimizer to $\|\mathbf{G A x}-\mathbf{G b}\|_{2}$ is also a constant-factor approximation to the minimizer of $\|\mathbf{S A R y}-\mathbf{S b}\|_{2}$ and due to the preconditioner $\mathbf{R}, \kappa(\mathbf{S A R})=$ $\mathcal{O}(1)$ since $\mathbf{S}$ is also a $(1+\varepsilon)$-approximate subspace embedding of $\mathbf{A}$. Thus, due to the bounded

[^1]condition number of SAR, it suffices to run a small number of steps of gradient descent (GD) to obtain a $(1+\mathcal{O}(\varepsilon))$-approximation to the minimizer of $\|\mathbf{S A y}-\mathbf{S b}\|_{2}$ and thus a $(1+\varepsilon)$-approximation to the regression problem $\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A x}-\mathbf{b}\|_{2}$.

### 1.3 Preliminaries

In this paper, we use $[n]$ to denote the set $\{1, \ldots, n\}$ for a positive integer $n$. We use $\operatorname{poly}(n)$ to denote a fixed degree polynomial in $n$ that can depend on fixed constants in instantiations of variables throughout the algorithm. When a random event occurs with probability at least $1-\frac{1}{\operatorname{poly}(n)}$, we say the event occurs with high probability. Similarly, we use polylog $(n)$ to denote $\operatorname{poly}(\log n)$. We use $\tilde{\Omega}(f)$ to denote $\Omega(f \operatorname{poly} \log (f))$ or $\Omega\left(\frac{f}{\operatorname{polylog}(f)}\right)$. We use $\mathcal{N}\left(\mu, \sigma^{2}\right)$ to denote the normal distribution with mean $\mu$ and variance $\sigma^{2}$ and $\mathcal{N}(\mu, \boldsymbol{\Sigma})$ to denote the multivariate normal distribution with mean $\mu$ and variance $\boldsymbol{\Sigma}$.

We use the following formulation of the bounded differences inequality:
Definition 1.5 (Bounded differences). For a domain $X$, let $f: X^{n} \rightarrow \mathbb{R}$. Then $f$ satisfies the bounded difference assumption if there exist $c_{1}, \ldots, c_{n} \geq 0$ such that for all $i \in[n]$,

$$
\sup _{x_{1}, \ldots, x_{n}, x_{i}^{\prime} \in X}\left|f\left(x_{1}, \ldots, x_{i}, \ldots, x_{n}\right)-f\left(x_{1}, \ldots, x_{i}^{\prime}, \ldots, x_{n}\right)\right| \leq c_{i}
$$

Theorem 1.6 (Bounded differences inequality, McDiarmid's inequality). [ $M^{+} 89$ ] Let $X_{1}, \ldots, X_{n} \in$ $X$ be independent random variables and suppose $f: x^{n} \rightarrow \mathbb{R}$ satisfies the bounded difference assumption with respect to constants $c_{1}, \ldots, c_{n}$. Then for all $t>0$,

$$
\operatorname{Pr}\left[f\left(X_{1}, \ldots, x_{n}\right)-\mathbb{E}\left[f\left(X_{1}, \ldots, X_{n}\right)\right]\right] \leq 2 \exp \left(-\frac{2 t^{2}}{\sum_{i=1}^{n} c_{i}^{2}}\right)
$$

Given a function $\phi: \mathbb{R} \rightarrow \mathbb{R}$, we say the function is $L$-Lipschitz for a parameter $L>0$ if $|f(x)-f(y)| \leq L \cdot|x-y|$ for all $x, y \in \mathbb{R}$. We use the following formulation of Talagrand's contraction principle:
Theorem 1.7 (Ledoux-Talagrand contraction). [LT91] Let $X_{1}, \ldots, X_{n} \in \mathbb{R}^{d}$ be i.i.d. random vectors, $\mathcal{F}$ be a class of real valued functions on $\mathbb{R}^{d}$ and $\sigma_{1}, \ldots, \sigma_{n}$ be independent Rademacher random variables. If $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is an L-Lipschitz function with $\phi(0)=0$, then:

$$
\mathbb{E}[\text { sup }]_{f \in \mathcal{F}} \sum_{i=1}^{n} \sigma_{i} \phi\left(f\left(X_{i}\right)\right) \leq L \cdot \mathbb{E}[\text { sup }]_{f \in \mathcal{F}} \sum_{i=1}^{n} \sigma_{i} f\left(X_{i}\right) .
$$

We also use the following bound on the sum of independent mean-zero random variables:
Theorem 1.8 (Symmetrization, e.g., Lemma 6.4.2 in [Ver18]). Let $x_{1}, \ldots, x_{n} \in \mathbb{R}$ be independent mean-zero random variables. Then

$$
\mathbb{E}\left[\left\|\sum_{i=1}^{n} x_{i}\right\|_{2}\right] \leq 2 \mathbb{E}\left[\left\|\sum_{i=1}^{n} \sigma_{i} x_{i}\right\|_{2}\right]
$$

where the second expectation is taken over realizations of the random variables $x_{i}$ and independent Rademacher variables $\sigma_{i} \in\{-1,+1\}$ for all $i \in[n]$.

We use bold font variables to represent vectors and matrices. For a vector $\mathbf{v} \in \mathbb{R}^{n}$, we use $\|\mathbf{v}\|_{2}$ to denote its Euclidean norm, so that $\|\mathbf{v}\|_{2}^{2}=\sum_{i=0}^{n} v_{i}^{2}$. We use nnz $(\mathbf{A})$ to denote the number of nonzero entries in a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and we use $\mathbf{A}^{-1}$ to denote the pseudo-inverse of $\mathbf{A}$. For a $\operatorname{matrix} \mathbf{A} \in \mathbb{R}^{n \times d}$, we use

$$
\|\mathbf{A}\|_{2}=\max _{\mathbf{x} \in \mathbb{R}^{d},\|\mathbf{x}\|_{2}=1}\|\mathbf{A x}\|_{2}
$$

to denote its operator norm and we use $\kappa(\mathbf{A})$ to denote its condition number, so that

$$
\kappa(\mathbf{A})=\|\mathbf{A}\|_{2}\left\|\mathbf{A}^{-1}\right\|_{2} .
$$

We use $\|\mathbf{A}\|_{F}$ to denote the Frobenius norm of a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, so that

$$
\|\mathbf{A}\|_{F}^{2}:=\sum_{i=1}^{n} \sum_{j=1}^{d} A_{i, j}^{2} .
$$

For a square matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, we use $\operatorname{Tr}(\mathbf{M})$ to denote its trace, so that $\operatorname{Tr}(\mathbf{M})=\sum_{i=1}^{n} M_{i, i}$. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a matrix $\mathbf{B} \in \mathbb{R}^{n \times p}$, we use $[\mathbf{A} ; \mathbf{B}]$ to denote the $n \times(d+p)$ dimensional matrix $\left[\begin{array}{ll}\mathbf{A} & \mathbf{B}\end{array}\right]$. We use $\mathbf{A} \succeq 0$ to denote that a matrix $\mathbf{A}$ is positive semidefinite (PSD).

We use the following formulation of the Matrix Bernstein inequality:
Theorem 1.9 (Matrix Bernstein inequality, e.g., Theorem 1.6 in [Tro12]). Let $\mathbf{Z}_{1}, \ldots, \mathbf{Z}_{n}$ be $a$ sequence of matrices with dimension $n \times d$ such that $\mathbb{E}\left[\mathbf{Z}_{i}\right]=0^{n \times d}$ and $\left\|\mathbf{Z}_{i}\right\| \leq R$ with high probability, for each $i \in[n]$. Let $\sigma^{2}=\max \left(\left\|\sum_{i \in[n]} \mathbb{E}\left[\mathbf{Z}_{i} \mathbf{Z}_{i}^{\top}\right]\right\|_{2},\left\|\sum_{i \in[n]} \mathbb{E}\left[\mathbf{Z}_{i}^{\top} \mathbf{Z}_{i}\right]\right\|_{2}\right)$. Then for all $t \geq 0$,

$$
\operatorname{Pr}\left[\left\|\sum_{i \in[n]} \mathbf{Z}_{i}\right\|_{2} \geq t\right] \leq(n+d) \exp \left(-\frac{t^{2} / 2}{\sigma^{2}+R t / 3}\right)
$$

Subspace embeddings. For an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, a $(1+\varepsilon)$-subspace embedding for $\mathbf{A}$ is a matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ such that for all $\mathbf{x} \in \mathbb{R}^{d}$,

$$
(1-\varepsilon)\|\mathbf{M A x}\|_{2} \leq\|\mathbf{A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{M A x}\|_{2},
$$

for some accuracy parameter $\varepsilon \in(0,1)$. The subspace embedding is oblivious if the matrix $\mathbf{M}$ is generated from a random distribution that is independent of $\mathbf{A}$; otherwise, the subspace embedding is non-oblivious.

A construction of an oblivious subspace embedding that can be computed in input-sparsity time has the following guarantees:

Theorem 1.10 (OSNAP Matrix). [NN13, CW13, Coh16] Given an accuracy parameter $\varepsilon>0$, for any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$, there exists a matrix $\mathbf{S} \in \mathbb{R}^{m \times n}$ with $m=\mathcal{O}\left(\frac{1}{\varepsilon^{2}} k^{1+\alpha} \log k\right)$ such that for all $\mathbf{x} \in \mathbb{R}^{d}$,

$$
\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{S A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2}
$$

with probability at least 0.99. Moreover, SA can be computed in time $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha \varepsilon}\right)$.
Another construction of oblivious subspace embeddings uses Randomized Hadamard Transforms, which are a family of structured randomized transformations defined as follows:

Definition 1.11 (Randomized Hadamard Transform).

$$
\begin{gathered}
\mathbf{H}_{1}=[1] \quad \mathbf{H}_{d}=\left[\begin{array}{cc}
\mathbf{H}_{d / 2} & \mathbf{H}_{d / 2} \\
\mathbf{H}_{d / 2} & -\mathbf{H}_{d / 2}
\end{array}\right] \\
\forall i \in[m]: \mathbf{D}^{(i)} \in \mathbb{R}^{d \times d} \text { and } D_{j, k}^{(i)} \sim\left\{\begin{array}{ll}
\mathcal{N}(0,1) & \text { if } j=k \\
0 & \text { otherwise }
\end{array}, \quad h(\mathbf{z})=\left[\begin{array}{c}
\mathbf{H D}^{(1)} \\
\mathbf{H D}^{(2)} \\
\vdots \\
\mathbf{H D}^{(m)}
\end{array}\right] \cdot \mathbf{z}\right.
\end{gathered}
$$

We require the following properties of Randomized Hadamard Transforms:
Theorem 1.12. [CN22, Theorem 1.1] Let $d \in \mathbb{N}, \delta, \varepsilon \in(0,1 / 2)$ and $f: \mathbb{R} \rightarrow \mathbb{R}$ be a 1-Lipschitz function. Then for the function $h$ defined in Definition 1.11, we have with probability at least $1-\delta$ :

$$
\forall z \in \mathbb{R}^{d} \text { s.t }\|z\| \leq 1:\left|\frac{1}{m d} \cdot \sum_{i=1}^{m d} f\left(h(z)_{i}\right)-\mathbb{E}_{Z \sim \mathcal{N}\left(0,\|z\|^{2}\right)}[f(z)]\right| \leq \varepsilon
$$

as long as $m \geq C \varepsilon^{-2} \log ^{5}(d / \varepsilon) \log (1 / \delta)$ for some absolute constant $C>0$.
Lemma 1.13. [CN22, Lemma B.6] For any $d \in \mathbb{N}, \varepsilon, \delta \in(0,1 / 2)$, we have that for $m \geq 4$. $\frac{\log d+\log (2 / \delta)}{\varepsilon^{2}}$ and the function $h$ defined in Definition 1.11

$$
\forall \mathbf{x} \in \mathbb{R}^{d}:(1-\varepsilon) \cdot\|\mathbf{x}\|_{2} \leq \frac{1}{\sqrt{m d}} \cdot\|h(\mathbf{x})\|_{2} \leq(1+\varepsilon) \cdot\|\mathbf{x}\|_{2}
$$

with probability at least $1-\delta$.
Corollary 1.14 (Subspace embedding via Randomized Hadamard Transform). Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$, there exist absolute constants $c, C>0$ and an explicit matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ with $m=k \operatorname{polylog}(n)$ rows, such that with probability at least 0.99 , for any vector $\mathbf{x} \in \mathbb{R}^{d}$,

$$
\frac{1}{2}\|\mathbf{A x}\|_{2} \leq\|\mathbf{M A x}\|_{2} \leq \frac{3}{2}\|\mathbf{A} \mathbf{x}\|_{2}
$$

and at least Cm of the coordinates of the vector MAx have magnitude at least $\frac{c}{\sqrt{m}}$. Moreover, MA can be computed in time $n d \operatorname{polylog}(n)$.

Leverage scores. A non-oblivious construction of a subspace embedding uses the notion of leverage score sampling. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, the leverage score of row $\mathbf{a}_{i}$ with $i \in[n]$ is defined as $\mathbf{a}_{i}\left(\mathbf{A}^{\top} \mathbf{A}\right)^{-1} \mathbf{a}_{i}^{\top}$. Equivalently, for the singular value decomposition $\mathbf{A}=\mathbf{U} \mathbf{\Sigma V}$, the leverage score of row $\mathbf{a}_{i}$ is also the squared row norm of $\mathbf{u}_{i}$. Thus, it is apparent that the sum of the leverage scores of $\mathbf{A}$ is at most the rank of $\mathbf{A}$, since the columns of $\mathbf{U}$ are orthogonal.

Theorem 1.15 (Generalization of Foster's Theorem, [Fos53]). For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, the sum of its leverage scores is $\operatorname{rank}(\mathbf{A})$.

It is well-known that leverage score sampling can generate a non-oblivious subspace embedding:

Theorem 1.16 (Leverage score sampling). [DMM06a, DMM06b, Mag10, Woo14] Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, let $\tau_{i}$ be the leverage score of the $i$-th row of $\mathbf{A}$. Let $C>1$ be a universal constant and $\alpha>1$ be a parameter and suppose that $p_{i} \in\left[\min \left(1, \frac{C \tau_{i} \log k}{\varepsilon^{2}}\right), \min \left(1, \frac{C \alpha \tau_{i} \log k}{\varepsilon^{2}}\right)\right]$ for each $i \in[n]$. Let $\mathbf{S}$ be a random diagonal matrix so that the $i$-th diagonal entry of $\mathbf{S}$ is $\frac{1}{\sqrt{p_{i}}}$ with probability $p_{i}$ and 0 with probability $1-p_{i}$. Then with probability at least 0.99 , for all vectors $\mathbf{x} \in \mathbb{R}^{d}$,

$$
(1-\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{S A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} .
$$

Moreover, $\mathbf{S}$ has at most $\mathcal{O}\left(\frac{\alpha}{\varepsilon^{2}} d \log d\right)$ nonzero entries with probability at least $1-e^{-\Theta(d)}$.
Leverage scores are particularly useful because Theorem 1.15 upper bounds the sum of the leverage scores by the rank of the matrix, which is at most $d$ for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with $n \geq d$. Thus Theorem 1.16 implies that only $\mathcal{O}(d \log d)$ rows of $\mathbf{A}$ need to be sampled for a constant factor subspace embedding of A, given constant-factor approximations to the leverage scores of $\mathbf{A}$.

Gradient descent for linear regression. Gradient descent is a well-known iterative method for finding a local minimum of a differentiable function $f$. Given a learning rate $\eta>0$ and a point $\mathbf{x}_{n} \in \mathbb{R}^{n}$ for iteration $n$, the point $\mathbf{x}_{n+1}$ for iteration $n+1$ is defined by

$$
\mathbf{x}_{n+1}=\mathbf{x}_{n}+\eta \nabla f\left(\mathbf{x}_{n}\right),
$$

where $\nabla f\left(\mathbf{x}_{n}\right)$ is the gradient of $f$ at $\mathbf{x}_{n}$. It is known that the learning rate $\eta$ can be explicitly chosen so that gradient descent achieves the following convergence rate guarantees:

Theorem 1.17 (Convergence of gradient descent, e.g., Theorem 3 in [Sin16]). For a convex set $S \subseteq \mathbb{R}^{d}$, let $f: S \rightarrow \mathbb{R}$ be strongly convex on $S$, so that there exist $M>m>0$ such that $m I_{d} \preceq \nabla^{2} f \preceq M I_{d}$. Then for any $\zeta>0$, we have $f\left(x^{(k)}\right)-\min _{x \in S} f(x) \leq \varepsilon$, i.e., $k$ iterations of gradient descent suffice to obtain an additive $\zeta$-approximation, for

$$
k \geq \frac{\log \frac{f\left(x^{(0)}\right)-y^{*}}{\zeta}}{\log \frac{M}{M-m}},
$$

where $y^{*}=\min _{x \in S} f(x)$.
In our context, we use gradient descent to approximately solve linear regression, i.e., $\min _{\mathbf{x} \in \mathbb{R}^{d}} f(\mathbf{x}):=$ $\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A x}-\mathbf{b}\|_{2}$ for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a vector $\mathbf{b} \in \mathbb{R}^{n}$. This is equivalent to minimizing the squared objective value $\min _{\mathbf{x} \in \mathbb{R}^{d}} f(\mathbf{x}):=\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A} \mathbf{x}-\mathbf{b}\|_{2}^{2}$. In this case, $\nabla f(\mathbf{x})=2 \mathbf{A}^{\top} \mathbf{A} \mathbf{x}-2 \mathbf{A}^{\top} \mathbf{b}$, which can be explicitly computed from $\mathbf{A}$ and $\mathbf{b}$. However, if the dimensions of $\mathbf{A}$ are prohibitively large, it is often desirable to first apply dimensionality reduction techniques to decrease the input size.

## 2 Constant-Factor Subspace Embedding

In this section, we describe our constant-factor subspace embedding. We first require a crude polylogarithmic approximate subspace embedding, which we describe in Section 2.1. Our polylogarithmic subspace embedding employs a recently observed property of the SRHT to "spread" the
mass of an input vector among a large number of coordinates. Since a large number of coordinates have a "large" amount of mass, it suffices by standard concentration inequalities to uniformly sample rows after the SRHT is applied.

We then show how to utilize the polylogarithmic subspace embedding to boost the approximation guarantees into a constant-factor subspace embedding in Section 2.2. Namely, some of the sampled rows in the SRHT could be too large. Thus we use fast semidefinite programming to reweight sampled rows of the SRHT to achieve a constant factor approximation. We provide more details of this high-level approach in the individual sections.

### 2.1 Polylogarithmic Subspace Embedding

In this section, we describe our polylogarithmic distortion subspace embedding. We first show that after applying the Hadamard Transform, a constant fraction of the resulting vector has "large" coordinates.

Lemma 2.1. There exist constants $c, C>0$ such that for the function $h$ defined in Definition 1.11:

$$
\forall \mathbf{x} \in \mathbb{R}^{d} \text { s.t }\|\mathbf{x}\|_{2}=1: \frac{1}{m d} \cdot \sum_{i=1}^{m d} 1\left\{\left|h(\mathbf{x})_{i}\right| \geq c\right\} \geq 0.999 \cdot m d
$$

with probability at least $1-\frac{1}{d^{10}}$ as long as $m \geq C \log ^{6}(d)$.
Proof. Consider the following approximation to an indicator function:

$$
f_{c}(x)= \begin{cases}0 & \text { if }|x| \leq c \\ \frac{|x|-c}{c} & \text { if } c \leq|x| \leq 2 c \\ 1 & \text { when }|x| \geq 2 c\end{cases}
$$

and let $c$ be small enough such that $\Phi(2 c)-\Phi(-2 c) \leq 10^{-5}$, where $\Phi$ is the CDF of a standard normal distribution. Note that such a $c$ exists as the pdf of a standard normal random variable is upper bounded by 1 . Now, for large enough $C$ and noting that $f_{c}$ is a $(1 / c)$-Lipschitz function that $\forall \mathbf{x} \in \mathbb{R}^{d}$ s.t $\|\mathbf{x}\|_{2}=1$, we have by Theorem 1.12:

$$
\frac{1}{m d} \cdot \sum_{i=1}^{m d} f\left(h(\mathbf{x})_{i}\right) \geq \mathbb{E}_{Z \sim \mathcal{N}(0,1)}[f(Z)]-10^{-5} \geq(1-(\Phi(c)-\Phi(-c)))-10^{-5} \geq 1-2 \cdot 10^{-5}
$$

with probability at least $1-d^{-10}$. The lemma now follows from the fact that $f(x) \leq \mathbf{1}\{|x| \geq c\}$.
We now describe our polylogarithmic distortion subspace embedding. Given an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we first apply OSNAP matrices $\mathbf{S}_{1}$ and $\mathbf{S}_{2}$ to obtain a constant-factor subspace embedding. The OSNAP matrix $\mathbf{S}_{2}$ will have sparsity $\alpha=\mathcal{O}(1)$ and thus dimension $\mathcal{O}\left(d^{1.1} \log d\right) \times n$ for $\alpha=0.1$, for example. This OSNAP matrix will allow us to achieve polylog( $d$ ) dependencies rather than $\operatorname{poly} \log (n)$ dependencies. Here sparsity $\alpha \in(0,1)$ means that a column of the OSNAP matrix will have $\frac{1}{\alpha \varepsilon}$ nonzero entries, so that an application of the OSNAP matrix incurs runtime proportional to $\frac{1}{\alpha \varepsilon}$. Next, the OSNAP matrix $\mathbf{S}_{1}$ will have sparsity $\alpha^{\prime}=\frac{1}{\log d}$ and thus dimension $\mathcal{O}(d \log d) \times \mathcal{O}\left(d^{1.1} \log d\right)$. The purpose of this OSNAP matrix is to slightly decrease the matrix
multiplication time in our analysis, though we remark that without $\mathbf{S}_{1}$, our analysis can still be performed but simply requiring a smaller value of $\alpha$ for $\mathbf{S}_{2}$.

Our polylogarithmic subspace embedding is simple. We apply an SRHT matrix $\mathbf{M}$ with $d$ polylog $(d)$ rows to $\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$. From standard results on Hadamard Transforms, $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A}$ is actually a constant-factor subspace embedding for $\mathbf{A}$. In light of Lemma 2.1, $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A x}$ has a "large" number of coordinates that are "large", for any unit vector $\mathbf{A x} \in \mathbb{R}^{d}$. Thus we can uniformly sample rows of $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A}$ and achieve a "good" approximation to $\|\mathbf{A x}\|_{2}$. Hence, our polylogarithmic subspace embedding is simply the matrix $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$, where $\mathbf{S}$ is a matrix that uniformly samples $\mathcal{O}(d)$ rows independently.
Theorem 2.2. For any $\mathbf{A} \in \mathbb{R}^{n \times d}$ and a tradeoff parameter $\alpha>0$, we may compute matrix $\mathbf{G} \in \mathbb{R}^{p \times n}$ such that:

$$
\forall \mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{G} \mathbf{A} \mathbf{x}\|_{2} \leq \operatorname{polylog}(d)\|\mathbf{A} \mathbf{x}\|_{2}
$$

with probability at least 0.9 for some constant $\xi>0$. Furthermore, we have $p=\mathcal{O}(d)$ and $\mathbf{G A}$ may be computed in time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{polylog}(d)$.
Proof. Let $\mathbf{S}_{1}$ and $\mathbf{S}_{2}$ be OSNAP matrices (Theorem 1.10) that induce a 2-approximate subspace embedding, such that $\alpha=0.1$ for $\mathbf{S}_{2}$ and $\alpha=\frac{1}{\log d}$ for $\mathbf{S}_{1}$. Consider two successive applications of OSNAP matrices $\mathbf{S}_{1}$ and $\mathbf{S}_{2}$ to $\mathbf{A}$ to obtain $\mathbf{B}=\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A} \in \mathbb{R}^{\ell \times d}$, where $\ell=\mathcal{O}(d \log d)$ by Theorem 1.10. Note that by Theorem 1.10, $\mathbf{S}_{2} \mathbf{A}$ can be computed in time $\mathcal{O}(\mathrm{nnz}(\mathbf{A}))$ and thus subsequently, $\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$ can be computed in time $\mathcal{O}\left(d^{2.1} \log d\right)$. Moreover, the exponent 2.1 is due to the choice $\alpha=0.1$ for $\mathbf{S}_{2}$ and can be made any arbitrary constant greater than 2 .

Next, consider an SRHT, characterized by matrix $\mathbf{M}$ with $m=\operatorname{polylog}(\ell)=\operatorname{polylog}(d)$ rows. For the corresponding linear mapping $h$, we have:

$$
\begin{gather*}
\forall \mathbf{x} \in \mathbb{R}^{\ell}:(1-\varepsilon)\|\mathbf{x}\|_{2} \leq \frac{1}{\sqrt{m \ell}} \cdot\|h(\mathbf{x})\|_{2} \leq(1+\varepsilon) \cdot\|\mathbf{x}\|_{2} \\
\forall \mathbf{x} \in \mathbb{R}^{\ell} \text { s.t }\|\mathbf{x}\|_{2}=1: \frac{1}{m \ell} \cdot \sum_{i \in[m \ell]} \mathbf{1}\left\{\left|h(\mathbf{x})_{i}\right| \geq c\right\} \geq 0.999 \\
\forall i \in[m \ell]:\left\|\mathbf{m}_{i}\right\|_{2} \leq 2 \sqrt{\ell} \tag{SRHT-COND}
\end{gather*}
$$

with probability at least $1-1 / \ell^{10}$ by Lemma 1.13 , Lemma 2.1, and the fact that the lengths of the rows of $\mathbf{M}$ correspond to the length of one of $m$ independently distributed standard normal random vectors. Here, we use $\mathbf{1}$ to denote an indicator variable and $\mathbf{m}_{i}$ to denote the $i$-th row of M.

Next, consider a subsampling matrix, $\mathbf{S} \in \mathbb{R}^{p \times m \ell}$ where each row of $\mathbf{S}$ is uniformly sampled from the set of elementary vectors $\left\{\mathbf{e}_{i}\right\}_{i \in[m \ell]}$ for $p=C d$ for some suitably large constant $C$. We now have by an application of the matrix Bernstein inequality, i.e., Theorem 1.9:

$$
\frac{1}{\sqrt{C d}} \cdot\|\mathbf{S M}\|_{2} \leq \operatorname{poly} \log (d)
$$

with probability at least $1-1 / \ell^{10}$. Now letting $T$ be the multiset of indices selected in the construction of $\mathbf{S}$, consider the random variable:

$$
Z=\sup _{\mathbf{u} \in \operatorname{Span}(\mathbf{B}) \text { s.t }\|\mathbf{u}\|_{2}=1}\left|\frac{1}{p} \cdot \sum_{i \in T} \mathbf{1}\left\{\left\langle\mathbf{m}_{i}, \mathbf{u}\right\rangle \geq c\right\}-\mathbb{E}_{i \in[m \ell]}\left[\mathbf{1}\left\{\left\langle\mathbf{m}_{i}, \mathbf{u}\right\rangle \geq c\right\}\right]\right|
$$

Since $Z$ satisfies the bounded differences assumption with respect to the elements of $T$ and the fact that $Z$ corresponds to the empirical concentration of indicator functions of halfspaces of dimension $d$, we have by standard VC Theory and McDiarmid's inequality, i.e., Theorem 1.6, that $Z \leq 0.001$ with probability at least $1-1 / \ell^{10}$. From this, we get that for a suitably large $C>0$ :

$$
\forall \mathbf{u} \in \operatorname{Span}(\mathbf{B}): \frac{1}{\sqrt{p}} \cdot\|\mathbf{S M u}\|_{2} \geq c \cdot\|\mathbf{u}\|_{2} \Longrightarrow\|\mathbf{u}\|_{2} \leq \frac{C}{\sqrt{p}}\|\mathbf{S M u}\|_{2} \leq \operatorname{poly} \log (d) \cdot\|\mathbf{u}\|_{2}
$$

Hence, for the matrix $\mathbf{G A}=\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$, we have that $\mathbf{G} \mathbf{A}$ is a subspace embedding with polylog $(d)$ distortion that can be computed in time $\mathcal{O}\left(\operatorname{nnz}(\mathbf{A})+d^{2.1} \log d\right)$. However, as we previously remarked, the exponent 2.1 is due to the choice $\alpha=0.1$ for $\mathbf{S}_{2}$ and can be made any arbitrary constant greater than 2 with the tradeoff that the $\mathrm{nnz}(\mathbf{A})$ term becomes $\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}$ in the application of the OSNAP matrix $\mathbf{S}_{2}$ in Theorem 1.10. Therefore, the overall runtime is $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{polylog}(d)$.

### 2.2 Constant-Factor Subspace Embedding

In this section, we describe how to improve our polylogarithmic factor subspace embedding into a constant-factor subspace embedding. We first require the following guarantees for (approximately) solving semidefinite programs (SDPs).

Theorem 2.3 (Theorem 1.1 in [PTZ12]). For a primal positive SDP with $m \times m$ matrices and $n$ constraints and an accuracy parameter $\varepsilon>0$, there exists an algorithm that produces a $(1+\varepsilon)$ approximation in $\mathcal{O}\left(\frac{1}{\varepsilon^{3}} \log ^{3} n\right)$ iterations, where each iteration consists of computing matrix sums and a special primitive that computes $\exp (\Phi) \bullet \mathbf{A}$ for positive semidefinite matrices (PSD) $\Phi$ and A.

Here, $\exp (\Phi) \bullet \mathbf{A}$ denotes the pointwise dot product between matrices $\exp (\Phi)$ and $\mathbf{A}$, i.e., the Hadamard product.

Theorem 2.4 (Theorem 4.1 in [PTZ12]). There exists an algorithm that takes input an $m \times m$ matrix $\Phi$ with $p$ nonzero entries, $\kappa \geq \max \left(1,\|\Phi\|_{2}\right)$, and PSD $m \times m$ matrices $\mathbf{A}_{i}$ in factorized form $\mathbf{A}_{i}=\mathbf{Q}_{i} \mathbf{Q}_{i}^{\top}$, where the total number of nonzero entries across all matrices $\mathbf{Q}_{i}$ is $q$, and outputs a $(1+\varepsilon)$-approximation to all $\exp (\Phi) \bullet \mathbf{A}_{i}$. The algorithm uses $\mathcal{O}\left(\frac{1}{\varepsilon^{2}}\left(p \kappa \log \frac{1}{\varepsilon}+q\right) \log m\right)$ total work.

We now describe our constant-factor subspace embedding. Recall that our polylogarithmic distortion subspace embedding for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is a matrix $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$, where $\mathbf{S}$ is a matrix that randomly and independently samples $\mathcal{O}(d)$ rows, $\mathbf{M}$ is an SRHT matrix with $d$ polylog $(d)$ rows and $\mathbf{S}_{1}$ and $\mathbf{S}_{2}$ are OSNAP matrices. Recall furthermore that from standard results for randomized Hadamard Transforms, $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A}$ is already a constant-factor subspace embedding. The reason $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$ is not a constant-factor subspace embedding is because there are certain rows of $\mathbf{M S}_{1} \mathbf{S}_{2} \mathbf{A}$ that are too large.

We first show that with high probability, there exists a reweighting $\mathbf{W}$ of the sampled rows so that $\mathbf{W S M S}_{1} \mathbf{S}_{2} \mathbf{A}$ is a constant-factor subspace embedding of $\mathbf{A}$. We can thus use semidefinite programming to efficiently compute such a set of weights and quickly output $\mathbf{W S M S}_{1} \mathbf{S}_{2} \mathbf{A}$. A high-level description of our constant-factor subspace embedding is summarized in Figure 2.
Theorem 1.1. For any $\mathbf{A} \in \mathbb{R}^{n \times d}$ and any tradeoff parameter $\alpha>0$, we can compute matrix $\mathbf{G} \in \mathbb{R}^{p \times n}$ such that:

$$
\forall \mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{G} \mathbf{A} \mathbf{x}\|_{2} \leq \xi\|\mathbf{A} \mathbf{x}\|_{2}
$$

with probability at least 0.9 for a fixed constant $\xi>1$. Furthermore, we have $p=\mathcal{O}(d)$ and $\mathbf{G A}$ may be computed in time $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{polylog}(d)$.

Proof. Let $\mathbf{S}_{1} \in \mathbb{R}^{\mathcal{O}(\ell) \times \tilde{\mathcal{O}}\left(d^{1+\alpha}\right)}$ and $\mathbf{S}_{2} \in \mathbb{R}^{\tilde{\mathcal{O}}\left(d^{1+\alpha}\right) \times n}$ be the OSNAP matrices defined in Theorem 2.2 , so that each matrix is a constant-factor subspace embedding and $\alpha=0.1$ for $\mathbf{S}_{2}$ and $\alpha=\frac{1}{\log d}$ for $\mathbf{S}_{1}$. In particular, we have $\ell=\mathcal{O}(d \log d)$. Let $\mathbf{M}$ be an $\operatorname{SRHT}$ with $m=\operatorname{polylog}(\ell)=\operatorname{polylog}(d)$ rows, as in Theorem 2.2. Finally, we let $\mathbf{S} \in \mathbb{R}^{p \times m \ell}$ be a subsampling matrix, where each row of $\mathbf{S}$ is uniformly sampled from the set of elementary vectors $\left\{\mathbf{e}_{i}\right\}_{i \in[m \ell]}$ for $p=C d$ for some suitably large constant $C$, as in Theorem 2.2. Note that as in the proof of Theorem 2.2, we may set $\alpha$ for $\mathbf{S}_{2}$ to any constant $\alpha>0$ to obtain the final result.

Let $\mathbf{x}_{1}, \ldots, \mathbf{x}_{p}$ denote the rows of $\mathbf{S M}, \mathbf{B}=\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$ and $\mathbf{U}$ be an orthonormal basis for $\operatorname{Span}(\mathbf{B})$. We will now find a set of weights $w \in \mathcal{W}:=\left\{w: \sum_{i \in[p]} w_{i}=1\right.$ and $\left.\forall i \in[p], 0 \leq w_{i} \leq \frac{2}{p}\right\} \subset \mathbb{R}^{p}$ such that $\left\|\mathbf{U U}^{\top} \cdot \boldsymbol{\Sigma}_{w} \cdot \mathbf{U} \mathbf{U}^{\top}\right\|_{2}$ is minimized where $\boldsymbol{\Sigma}_{w}=\sum_{i \in[p]} w_{i} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$. We start by showing that this quantity is small with high probability.

Lemma 2.5. We have for some suitably large constant $C>0$ :

$$
\min _{w \in \mathcal{W}}\left\|\mathbf{U U}^{\top} \cdot \boldsymbol{\Sigma}_{w} \cdot \mathbf{U} \mathbf{U}^{\top}\right\|_{2} \leq C
$$

With probability at least $1-1 / p^{10}$.
Proof. We start by analyzing the quantity via the approach from [LM19]:

$$
Z_{r}=\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i=1}^{p} \mathbf{1}\left\{\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right| \geq r\right\}
$$

Note that $Z_{r}$ satisfies the bounded differences inequality with respect to the rows $x_{i}$ drawn i.i.d. from the rows of $\mathbf{M}$. Furthermore, we have:

$$
\begin{aligned}
\mathbb{E}\left[Z_{r}\right] & \leq \frac{1}{r} \cdot \mathbb{E}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]}\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|\right] \\
& \leq \frac{1}{r}\left(\mathbb{E}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]}\left(\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|-\mathbb{E}_{\mathbf{x} \in \mathbf{M}}|\langle\mathbf{u}, \mathbf{x}\rangle|\right)\right]+\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \mathbb{E}_{\mathbf{x} \in \mathbf{U}}[|\langle\mathbf{u}, \mathbf{x}\rangle|]\right)
\end{aligned}
$$

Since $\mathbf{u}$ and $\mathbf{x}$ are both unit vectors, then $\mathbb{E}_{\mathbf{x} \in \mathbf{U}}[|\langle\mathbf{u}, \mathbf{x}\rangle|] \leq 1$. Thus we have

$$
\mathbb{E}\left[Z_{r}\right] \leq \frac{1}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \mathbf{x}_{i}^{\prime}}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]}\left(\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|-\left|\left\langle\mathbf{u}, \mathbf{x}_{i}^{\prime}\right\rangle\right|\right)\right]+1\right)
$$

Note that $\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|-\left|\left\langle\mathbf{u}, \mathbf{x}_{i}^{\prime}\right\rangle\right|$ is a zero-mean random variable. Thus by using standard symmetrization arguments, i.e., Theorem 1.8, we can insert Radamacher variables $\sigma_{i}$ as follows where $x_{i}^{\prime}$ represent independent copies of $x_{i}$. Therefore,

$$
\mathbb{E}\left[Z_{r}\right] \leq \frac{2}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \mathbf{x}_{i}^{\prime}, \sigma_{i}}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i}\left(\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|-\left|\left\langle\mathbf{u}, \mathbf{x}_{i}^{\prime}\right\rangle\right|\right)\right]+1\right)
$$

$$
\begin{aligned}
& \leq \frac{2}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \mathbf{x}_{i}^{\prime}, \sigma_{i}}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i}\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|+\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]}-\sigma_{i}\left|\left\langle\mathbf{u}, \mathbf{x}_{i}^{\prime}\right\rangle\right|\right]+1\right) \\
& =\frac{4}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \sigma_{i}}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i}\left|\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right|\right]+1\right) .
\end{aligned}
$$

Since the Rademacher random variables $\sigma_{i} \in\{ \pm 1\}$ are chosen uniformly at random and independent of the $x_{i}$, we can remove the absolute values around $\left\langle u, x_{i}\right\rangle$ using Talagrand's contraction principle, i.e., Theorem 1.7. Hence, by the definition of the operator norm,

$$
\begin{aligned}
\mathbb{E}\left[Z_{r}\right] & \leq \frac{4}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \sigma_{i}}\left[\sup _{\mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1} \frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i}\left\langle\mathbf{u}, \mathbf{x}_{i}\right\rangle\right]+1\right) \\
& =\frac{4}{r}\left(\mathbb{E}_{\mathbf{x}_{i}, \sigma_{i}}\left[\left\|\mathbf{U}^{\top}\left(\frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i} \mathbf{x}_{i}\right)\right\|_{2}\right]+1\right) .
\end{aligned}
$$

By convexity and Jensen's inequality, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{r}\right] & \leq \frac{4}{r}\left(\left(\mathbb{E}_{\mathbf{x}_{i}, \sigma_{i}}\left[\left\|\mathbf{U}^{\top}\left(\frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i} \mathbf{x}_{i}\right)\right\|_{2}^{2}\right]\right)^{1 / 2}+1\right) \\
& \leq \frac{4}{r}\left(\left(\mathbb{E}_{\mathbf{x}_{i}, \sigma_{i}}\left[\operatorname{Tr}\left(\mathbf{U}^{\top}\left(\left(\frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i} \mathbf{x}_{i}\right)\right)\left(\left(\frac{1}{p} \cdot \sum_{i \in[p]} \sigma_{i} \mathbf{x}_{i}\right)^{\top}\right) \mathbf{U}\right)\right]\right)^{1 / 2}+1\right) \\
& \leq \frac{4}{r} \cdot\left(\left(\operatorname{Tr}\left(\mathbf{U}^{\top}\left(\frac{1}{p} \cdot \mathbb{E}_{\mathbf{x} \in \mathbf{M}}\left[\mathbf{x} \mathbf{x}^{\top}\right]\right) \mathbf{U}\right)\right)^{1 / 2}+1\right),
\end{aligned}
$$

where the last inequality follows from the linearity of the trace operator. Since $\mathbf{U}$ is an orthonormal basis for $\operatorname{Span}(\mathbf{B})$, a subspace of dimension $d$ and $\mathbf{M}$ satisfies $\mathbb{E}_{x \in \mathbf{M}}\left[x x^{\top}\right] \preccurlyeq 2 \cdot I$ (SRHT-COND):

$$
\mathbb{E}\left[Z_{r}\right] \leq \frac{4}{r} \cdot\left(\left(\frac{2 d}{p}\right)^{1 / 2}+1\right)
$$

Hence, we get for large enough $r$ that $\mathbb{E}\left[Z_{r}\right] \leq 0.0001$. The bounded differences inequality, i.e., Theorem 1.6, now yields that $Z_{r} \leq 0.0002$ with probability at least $1-1 / \ell^{10}$. Now, we analyze the random variable:

$$
\min _{\mathbf{w} \in \mathcal{W}}\left\|\mathbf{U} \mathbf{U}^{\top} \boldsymbol{\Sigma}_{w} \mathbf{U} \mathbf{U}^{\top}\right\|_{2}=\min _{\mathbf{w} \in \mathcal{W}} \max _{\mathbf{Y} \succcurlyeq 0, \operatorname{Tr}(\mathbf{Y})=1}\left\langle\mathbf{U} \mathbf{U}^{\top} \boldsymbol{\Sigma}_{w} \mathbf{U U}^{\top}, \mathbf{Y}\right\rangle=\max _{\mathbf{Y} \succcurlyeq 0, \operatorname{Tr}(\mathbf{Y})=1} \min _{\mathbf{w} \in \mathcal{W}}\left\langle\mathbf{U} \mathbf{U}^{\top} \boldsymbol{\Sigma}_{w} \mathbf{U U}^{\top}, \mathbf{Y}\right\rangle
$$

where the exchange of the min and max follows from von Neumann's equality. We will now show that for all $\mathbf{Y} \succcurlyeq 0, \operatorname{Tr}(\mathbf{Y})=1$, we have $T_{\mathbf{Y}}:=\left|\left\{i: \mathbf{x}_{i}^{\top} \mathbf{U U}^{\top} \mathbf{Y U U}^{\top} \mathbf{x}_{i} \geq 32768 r^{2}\right\}\right|<0.5 p$ using an analysis similar to [DL22]. Suppose for the sake of contradiction, there exists such an $\mathbf{Y}$ satisfying this. Then, consider a Gaussian random variable $g \sim \mathcal{N}(0, \mathbf{Y})$. We now have for
all $i \in T_{\mathbf{Y}}$ by noting that $\mathbf{g}^{\top} \mathbf{U U}^{\top} \mathbf{x}_{i}$ is a zero mean gaussian random variable with variance $\mathbb{E}\left[\left(\mathbf{g}^{\top} \mathbf{U} \mathbf{U}^{\top} \mathbf{x}_{i}\right)^{2}\right]=\mathbb{E}\left[\mathbf{x}_{i}^{\top} \mathbf{U U}^{\top} \mathbf{g g}^{\top} \mathbf{U U}^{\top} \mathbf{x}_{i}\right]=\left\langle\mathbf{Y}, \mathbf{U U}^{\top} \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \mathbf{U U}^{\top}\right\rangle \geq 32768 r^{2}$ and standard upper bounds on the pdf of a gaussian random variable:

$$
\begin{gathered}
\mathbb{P}\left\{\left|\mathbf{g}^{\top} \mathbf{U} \mathbf{U}^{\top} \mathbf{x}_{i}\right| \geq 8 r\right\} \geq \frac{9}{10} \\
\mathbb{P}\left\{\|\mathbf{g}\|_{2} \leq 4\right\} \geq \frac{9}{10}
\end{gathered}
$$

Hence, we get by a union bound on the above two events:

$$
\mathbb{P}\left\{\frac{1}{\|\mathbf{g}\|_{2}} \cdot\left|\mathbf{g}^{\top} \mathbf{U U}^{\top} \mathbf{x}_{i}\right| \geq 2 r\right\} \geq \frac{8}{10}
$$

And we get:

$$
p \cdot Z_{r} \geq \mathbb{E}_{\mathbf{g}}\left[\sum_{i \in T_{\mathbf{Y}}} \mathbf{1}\left\{\frac{1}{\|\mathbf{g}\|_{2}} \cdot\left|\mathbf{g}^{\top} \mathbf{U} \mathbf{U}^{\top} \mathbf{x}_{i}\right| \geq 2 r\right\}\right] \geq \frac{8}{10} \cdot\left|T_{\mathbf{Y}}\right| \geq 0.4 p
$$

which yields a contradiction and establishes the lemma.
The task of finding a suitable set of weights can be formulated as the following packing semidefinite program:

$$
\begin{array}{ll}
\max & 1^{\top} \mathbf{w} \\
\text { s.t } & \sum_{i \in[p]} w_{i} \mathbf{A}_{i} \preccurlyeq C \cdot \mathbb{I} \\
& w_{i} \geq 0 \\
& \text { where } \mathbf{A}_{i}=\mathbf{C}_{i} \mathbf{C}_{i}^{\top} \text { with } \mathbf{C}_{i}=\left[\begin{array}{cc}
\mathbf{U U}^{\top} \cdot \mathbf{x}_{i} & 0 \\
0 & \sqrt{\frac{2}{p}} \cdot \mathbf{e}_{i}
\end{array}\right] .
\end{array}
$$

These families of SDPs may be solved to constant accuracy in time $d^{2}$ polylog $(d)$ from Theorem 2.3 by noting that the matrix exponentials computed in Theorem 2.4 may be implemented with the fast projection oracle onto the span of $\mathbf{U}$ via gradient descent to inverse polynomial accuracy.

Towards concluding the proof, let $\mathbf{W}$ denote the diagonal matrix with $W_{i, i}=\sqrt{w_{i}}$. We now have $\left\|\mathbf{W S U U}{ }^{\top}\right\|_{2} \leq C$ for some $C>0$ from the constraints of the program.
Lemma 2.6. We have for some absolute $c>0$ :

$$
\forall \mathbf{w} \in \mathcal{W}, \mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1: \mathbf{u}^{\top}\left(\sum_{i \in[p]} w_{i} \cdot \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right) \mathbf{u} \geq c
$$

with probability at least $1-1 / p^{10}$.
Proof. We have from SRHT-COND, that for any $\mathbf{u} \in \mathbf{U}$ with $\|\mathbf{u}\|_{2}=1$ for some $c>0$ :

$$
Q(\mathbf{u}):=\frac{1}{m l} \sum_{i \in[m l]} \mathbf{1}\left\{\left|h(\mathbf{u})_{i}\right| \geq c\right\} \geq 0.999
$$

Noting that the VC-dimension of halfspaces of dimension $d$ is $d+1$, we have by [Ver18, Theorem 8.3.23]:

$$
\mathbb{E}[\underbrace{\sup _{\mathbf{u} \in \mathbf{U},\|\mathbf{u}\|_{2}=1}\left|\frac{1}{p} \cdot \sum_{i \in[p]} \mathbf{1}\left\{\left|\left\langle\mathbf{x}_{i}, \mathbf{u}\right\rangle\right| \geq c\right\}-Q(\mathbf{u})\right|}_{Q}] \leq C \cdot \sqrt{\frac{d}{p}}
$$

for some absolute constant $C>0$. Furthermore, noting that the random variable $Q$ satisfies the bounded differences inequality as the rows $\mathbf{x}_{i}$ are drawn i.i.d. from the rows of $\mathbf{M}$. Hence, we have by an application of the bounded differences inequality and the previous display that with probability at least $1-1 / p^{10}$ :

$$
\forall \mathbf{u} \in \mathbf{U} \text { s.t }\|\mathbf{u}\|_{2}=1: \frac{1}{p} \cdot \sum_{i \in[p]} \mathbf{1}\left\{\left|\left\langle\mathbf{x}_{i}, \mathbf{u}\right\rangle\right| \geq c\right\} \geq 0.99
$$

The lemma now follows from the fact that for any $\mathbf{w} \in \mathcal{W}$, from the fact that $w_{i} \leq 2 / p$ :

$$
\|\mathbf{W S M u}\|_{2}=\sqrt{\sum_{i \in[p]} w_{i}\left\langle\mathbf{x}_{i}, \mathbf{u}\right\rangle^{2}} \geq c \cdot \sqrt{\sum_{i \in[p]} w_{i} \mathbf{1}\left\{\left|\left\langle\mathbf{x}_{i}, \mathbf{u}\right\rangle\right| \geq c\right\}} \geq \frac{c}{2} .
$$

To conclude the proof, note that $\mathbf{S}_{1} \mathbf{S}_{2}$ is a valid constant-factor subspace embedding for $\mathbf{A}$; that is, there exist constants $C^{\prime}, \xi^{\prime}>0$ such that:

$$
\forall \mathbf{x} \in \mathbb{R}^{d}:\|\mathbf{A} \mathbf{x}\|_{2} \leq C^{\prime}\left\|\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A} \mathbf{x}\right\|_{2} \leq \xi^{\prime}\|\mathbf{A} \mathbf{x}\|_{2} .
$$

Furthermore, we have as a consequence of Lemmas 2.5 and 2.6 that there exist constants $C^{\prime \prime}, \xi^{\prime \prime}$ :

$$
\begin{aligned}
\forall \mathbf{u} \in \mathbf{U}:\|\mathbf{u}\|_{2} \leq C^{\prime \prime}\|\mathbf{W S M u}\|_{2} & =\sqrt{\mathbf{u}^{\top} \cdot\left(\sum_{i \in[p]} w_{i} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right) \cdot \mathbf{u}} \\
& =\sqrt{\mathbf{u}^{\top} \mathbf{U} \mathbf{U}^{\top} \Sigma_{\mathbf{w}} \mathbf{U} \mathbf{U}^{\top} \mathbf{u}} \leq \sqrt{\left\|\mathbf{U U}^{\top} \Sigma_{\mathbf{w}} \mathbf{U U}^{\top}\right\|_{2}} \cdot\|\mathbf{u}\|_{2} \leq \xi^{\prime \prime}\|\mathbf{u}\|_{2}
\end{aligned}
$$

Recalling that $\mathbf{U}$ is the span of $\mathbf{S}_{1} \mathbf{S}_{2} \mathbf{A}$, the previous two displays yield:

$$
\forall\|\mathbf{x}\|_{2}=1:\|\mathbf{A} \mathbf{x}\|_{2} \leq C\left\|\mathbf{W S M S}_{1} \mathbf{S}_{2} \mathbf{A} \mathbf{x}\right\|_{2} \leq \xi\|\mathbf{A} \mathbf{x}\|_{2}
$$

for some constants $C, \xi>0$ with probability at least $1-1 / d^{10}$.
Runtime analysis. We start by showing that we may approximately project onto $\mathbf{U}$ by computing a pre-conditioner, $\mathbf{R}$, of $\mathbf{B}$ such that $\mathbf{B R}$ has condition number polylog $(d)$ :
Lemma 2.7. We may compute in time $\mathcal{O}\left(d^{\omega}\right)$ a matrix $\mathbf{R}$ such that $\mathbf{B R}$ has condition number polylog(d).

Proof. We start by computing a pre-conditioner of $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$, $\mathbf{R}$, in time $\mathcal{O}\left(d^{\omega}\right)$ as $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A}$ is of dimension $\mathcal{O}(d) \times d$. Consequently, $\mathbf{S M S}_{1} \mathbf{S}_{2} \mathbf{A R}$ has condition number polylog $(d)$. The lemma will now follow by showing that $\mathbf{S M}$ is a polylog $(d)$-subspace embedding for $\mathbf{U}$. The lower bound follows from Lemma 2.6. The upper bound now follows from applying matrix Bernstein (i.e., Theorem 1.9) to the random matrix $\sum_{i \in[p]} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$ by noting that $\left\|\mathbf{x}_{i}\right\|_{2} \leq \sqrt{d} \operatorname{polylog}(d)$ and $\mathbb{E}\left[\left(\mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right)^{2}\right] \preccurlyeq 4 \ell \cdot \mathbb{I} \preccurlyeq 4 d$ polylog $(d) \cdot \mathbb{I}$ by SRHT-COND.

As a consequence of the above lemma, we may compute an approximate projection onto $\mathbf{U}$ in time $d^{2} \operatorname{polylog}(d) \log (1 / \gamma)$ with accuracy $\gamma$; i.e., for any $\|\mathbf{u}\|_{2}=1$, we can compute a vector $\hat{\mathbf{u}} \in \mathbf{U}$ such that $\left\|\hat{\mathbf{u}}-\mathbf{u}^{*}\right\| \leq \gamma$ where $\mathbf{u}^{*}=\operatorname{argmin}_{\mathbf{z} \in \mathbf{U}}\|\mathbf{u}-\mathbf{z}\|$ via gradient descent. Now, to determine the runtime, we first note that by Theorem $1.10, \mathbf{S}_{1} \mathbf{A}$ can be computed in time $\mathcal{O}(\mathrm{nnz}(\mathbf{A}))$ and has dimension $d^{1.1} \log d \times d$. Similarly by Theorem $1.10, \mathbf{S}_{2} \mathbf{S}_{1} \mathbf{A}$ can be subsequently computed in time $\mathcal{O}\left(d^{2.1} \log ^{2}(d)\right)$ and has dimension $d \log d \times d$. By Corollary 1.14, $\mathbf{M S}_{2} \mathbf{S}_{1} \mathbf{A}$ can then subsequently be computed in time $d^{2}$ polylog $(d)$ and has dimensions $d$ polylog $(d) \times d$. Since $\mathbf{S}$ is a sampling matrix that samples $\mathcal{O}(d)$ rows, then $\mathbf{S M S}_{2} \mathbf{S}_{1} \mathbf{A}$ can be subsequently computed in time $\mathcal{O}\left(d^{2}\right)$ and has dimensions $\mathcal{O}(d) \times d$. Since the SDP can be solved to constant accuracy in time $d^{2}$ polylog $(d)$, then $\mathbf{W}$ can be computed in time $d^{2} \operatorname{polylog}(d)$. Since $\mathbf{W}$ is simply a reweighting matrix, then $\mathbf{W S M S}_{2} \mathbf{S}_{1} \mathbf{A}$ can subsequently be computed in time $\mathcal{O}\left(d^{2}\right)$. Therefore, the total time to compute the subspace embedding $\mathbf{W S M S}_{2} \mathbf{S}_{1} \mathbf{A}$ is

$$
\mathcal{O}\left(\mathrm{nnz}(\mathbf{A})+d^{2.1} \log ^{2}(d)\right)+d^{2} \operatorname{polylog}(d) .
$$

More generally, we can use an arbitrary $\alpha$ instead of setting $\alpha=0.1$ to achieve the total runtime

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{2+\alpha} \log ^{2}(d)\right)+d^{2} \operatorname{polylog}(d)=\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{poly} \log (d) .
$$

## 3 Subspace Embedding through Leverage Score Sampling

In this section, we show that our constant factor approximation can be used to achieve leverage score sampling in the current matrix-multiplication runtime. Leverage score sampling is an important tool that will allow us to achieve a $(1+\varepsilon)$-subspace embedding in this section, approximate linear regression in Section 5, and independent row selection in Section 4.

We first recall the following standard result, which states that a constant-factor subspace embedding can be used to achieve constant-factor approximations to the leverage scores.

Lemma 3.1. Suppose $\mathbf{S}$ is a subspace embedding for $\mathbf{A} \in \mathbb{R}^{n \times d}$ so that for any $\mathbf{x} \in \mathbb{R}^{d}$,

$$
\|\mathbf{A x}\|_{2} \leq\|\mathbf{S A x}\|_{2} \leq \alpha\|\mathbf{A} \mathbf{x}\|_{2} .
$$

Then for all $i \in[n]$,

$$
\frac{\tau}{\alpha^{2}} \leq\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2} \leq \tau
$$

where $\mathbf{S A}=\mathbf{Q R}^{-1}$ for an orthonormal matrix $\mathbf{Q}$, i.e., $\mathbf{Q R}^{-1}$ is the $Q R$ decomposition of $\mathbf{S A}$, and $\tau_{i}$ is the leverage score of the $i$-th row of $\mathbf{A}$.

Lemma 3.1 follows from the fact that $\mathbf{Q}$ has orthonormal columns and that the leverage score of each row of $\mathbf{A}$ is just the squared row norm of $\mathbf{U}$ in the singular value decomposition $\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}$, see, e.g., Section 1.3, while $\mathbf{S}$ is an $\alpha$-distortion subspace embedding. More detailed proofs of Lemma 3.1 appear in [DMMW12, Woo14, CCKW22].

To quickly obtain a constant factor approximation to $\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}$ for all $i \in[n]$, a standard approach is to use a Gaussian matrix $\mathbf{G}$ with $\mathcal{O}(\log n)$ columns and then compute $\left\|\mathbf{a}_{i} \mathbf{R G}\right\|_{2}^{2}$ [DMMW12, Woo14]. However, this multiplication by a dense matrix incurs a high runtime. Instead, [CCKW22] showed that a two-stage sampling process can be performed by first using a Gaussian matrix $\mathbf{G}^{\prime}$ with only $\mathcal{O}(1 / \gamma)$ columns, so that $\left\|\mathbf{a}_{i} \mathbf{R G}\right\|_{2}^{2}$ is an $\mathcal{O}\left(n^{\gamma} \log n\right)$-approximation to $\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}$ for each $i \in[n]$. Then after sampling a large number of rows, i.e., oversampling by an $\mathcal{O}\left(n^{\gamma} \log n\right)$ factor, we can compute constant-factor approximations to the sampled rows and then perform rejection sampling to reduce the overall number of rows. Formally, the guarantees are as follows:

Theorem 3.2 (Lemma 7.3 in [CCKW22]). Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, suppose $\mathbf{R} \in \mathbb{R}^{d \times d}$ is a matrix such that for any vector $\mathbf{x} \in \mathbb{R}^{d}$, ARx can be computed in time $T_{1}$ and $\mathbf{R x}$ can be computed in time $T_{2}$. Given parameters $\alpha, s>0$, there exists an algorithm that with probability at least 0.95 , samples a random subset $S \subseteq[n]$ such that each $i \in S$ with probability $f_{i}$, where

$$
\min \left(1, \frac{s}{16} \frac{\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}}{\|\mathbf{A R}\|_{F}^{2}}\right) \leq f_{i} \leq \min \left(1, s \frac{\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}}{\|\mathbf{A R}\|_{F}^{2}}\right)
$$

The algorithm outputs $S$ along with $f_{i}$ for each $i \in S$ in time $\mathcal{O}\left(\frac{T_{1}}{\alpha}+T_{2} \log n+s d n^{\alpha} \log ^{2}(n)\right)$.
Given these previous results, our algorithm is simple. To perform leverage score sampling on an input matrix $\mathbf{A}$ to obtain a $(1+\varepsilon)$-subspace embedding, we first obtain a constant-factor subspace embedding $\mathbf{S}_{1} \mathbf{A}$ through Theorem 1.1. We then compute a QR decomposition of $\mathbf{S}_{1} \mathbf{A}$ so that $\mathbf{Q R} \mathbf{R}^{-1}=\mathbf{S}_{1} \mathbf{A}$ and apply Theorem 3.2 to $\mathbf{A}$ and $\mathbf{R}$. The algorithm in full appears in Algorithm 1.

```
Algorithm 1 Subspace embedding through leverage score sampling
Input: \(\mathbf{A} \in \mathbb{R}^{n \times d}, \varepsilon, \alpha>0\)
Output: Subspace embedding SA
    Let \(\mathbf{S}_{1} \mathbf{A}\) be a fast embedding of \(\mathbf{A} \quad \triangleright\) Theorem 1.1
    Let \(\mathbf{Q R}{ }^{-1}\) be a QR decomposition of \(\mathbf{S}_{1} \mathbf{A}\)
    \(s \leftarrow \frac{1}{\varepsilon^{2}} d \operatorname{polylog}(d)\)
    Let \(S\) and \(\left\{f_{i}\right\}\) be the output of Theorem 3.2 with inputs \(\mathbf{A}, \mathbf{R}, s, \alpha\)
    Set the \(i\)-th diagonal entry of \(\mathbf{S}=\frac{1}{\sqrt{J_{i}}}\) for each \(i \in S\)
    return SA
```

We now show that Algorithm 1 can be used to obtain a $(1+\varepsilon)$-subspace embedding in the current matrix-multiplication runtime.

Theorem 1.2. Given $\mathbf{A} \in \mathbb{R}^{n \times d}$, an accuracy parameter $\varepsilon>0$, and any tradeoff parameter $\alpha>0$, there exists an algorithm that computes a matrix $\mathbf{S A}$ with $\mathcal{O}\left(\frac{1}{\varepsilon^{2}} d \log d\right)$ rows such that with probability at least $\frac{9}{10}$, for all vectors $\mathbf{x} \in \mathbb{R}^{d}$,

$$
(1-\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{S A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} .
$$

Moreover, SA can be computed in time

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon^{2}} d^{2+\alpha} \operatorname{poly} \log (d) .
$$

Proof. By Theorem 1.1, there exists a matrix $\mathbf{S}_{1}$ such that with probability at least $0.99, \mathbf{S}_{1}$ has $\mathcal{O}(d)$ rows and

$$
\|\mathbf{A} \mathbf{x}\|_{2} \leq\left\|\mathbf{S}_{1} \mathbf{A} \mathbf{x}\right\|_{2} \leq \xi\|\mathbf{A} \mathbf{x}\|_{2}
$$

for $\xi=\mathcal{O}(1)$ and for all $\mathbf{x} \in \mathbb{R}^{d}$. Moreover, $\mathbf{S}_{1} \mathbf{A}$ can be computed in $\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}\right)+d^{2+\alpha} \operatorname{polylog}(d)$ time for any $\alpha>0$. Thus the QR decomposition of $\mathbf{S}_{1} \mathbf{A}$ can be computed in time $\mathcal{O}\left(d^{\omega}\right)$ to output matrices $\mathbf{Q}$ and $\mathbf{R}^{-1}$ such that $\mathbf{Q}$ has orthonormal columns and $\mathbf{Q} \mathbf{R}^{-1}=\mathbf{S}_{1} \mathbf{A}$. By Lemma 3.1,

$$
\frac{\tau_{i}}{\xi^{2}} \leq\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2} \leq \tau_{i}
$$

for all $i \in[n]$, so that

$$
\frac{\tau_{i}}{d \xi^{2}} \leq \frac{\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}}{\|\mathbf{A R}\|_{F}^{2}}
$$

since $\|\mathbf{A R}\|_{F}^{2}=\sum_{i \in[n]}\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2} \leq \sum_{i \in[n]} \tau_{i} \leq d$ by Theorem 1.15. By Theorem 3.2, there exists an algorithm that with probability at least 0.95 , will output a set $S$ along with corresponding sampling probabilities $f_{i}$, for each $i \in S$, such that

$$
f_{i} \geq \min \left(1, \frac{s}{16} \frac{\left\|\mathbf{a}_{i} \mathbf{R}\right\|_{2}^{2}}{\|\mathbf{A R}\|_{F}^{2}}\right) \geq \min \left(1, \frac{s}{16} \frac{\tau_{i}}{d \xi^{2}}\right)
$$

Setting $s=\frac{1}{\varepsilon^{2}} d \xi^{2}$, it follows that $f_{i} \geq \min \left(1, \frac{C \tau_{i} \log d}{\varepsilon^{2}}\right)$ for some constant $C>0$. Thus, by Theorem 1.16, with probability at least 0.9 for the matrix $\mathbf{S}$ of Algorithm 1,

$$
(1-\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} \leq\|\mathbf{S A} \mathbf{x}\|_{2} \leq(1+\varepsilon)\|\mathbf{A} \mathbf{x}\|_{2} .
$$

By setting $T_{1}=\operatorname{nnz}(\mathbf{A})+d^{2}$ and $T_{2}=d^{2}$ in Theorem 3.2, it follows that the total runtime is

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon^{2}} n^{\alpha} d^{2} \operatorname{polylog}(d) .
$$

Now we note that either $d>n^{0.1}$, in which case $n^{\alpha}$ can be replaced with $d^{\alpha}$ after a reparametrization of $\alpha$ or $d<n^{0.1}$, in which case the $n^{\alpha} d^{2} \operatorname{polylog}(k)$ term is lower-order since nnz $(\mathbf{A})$ can be assumed to be at least $n$ by throwing out zero rows. Therefore, the final runtime is

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon^{2}} d^{2+\alpha} \operatorname{poly} \log (d) .
$$

## 4 Independent Row Selection

In this section, we show how our leverage score sampling framework and our constant-factor subspace embedding can be used to select a maximal set of independent rows of an input matrix $\mathbf{A}$ in the current matrix-multiplication runtime. We first require the following definition of rankpreserving sketches:

Definition 4.1 (Rank-preserving sketch). A distribution $\mathcal{S}$ on matrices $\mathbf{S} \in \mathbb{R}^{m \times n}$ is a rank preserving sketch if there exists a constant $c>0$ such that for $\mathbf{S} \sim \mathcal{S}$, with high probability, for an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we have $\min \left(\operatorname{rank}(\mathbf{S A}), \frac{m}{c}\right)=\min \left(\operatorname{rank}(\mathbf{A}), \frac{m}{c}\right)$.
[CKL13] gave a sparse construction of a rank-preserving sketch with the following properties:
Theorem 4.2. [CKL13] There exists a rank-preserving sketch distribution with $c=11$ such that (1) $\mathbf{S A}$ can be computed in $\mathcal{O}(\mathrm{nnz}(\mathbf{A}))$ time, (2) $\mathbf{S}$ has at most two nonzero entries in a column, and (3) $\mathbf{S}$ has at most $\frac{2 n}{m}$ nonzeros in a row.

We also require the following guarantees from the approximate matrix multiplication algorithm. More precisely, the approximate matrix multiplication algorithm samples a fixed number of rows with replacement from an input matrix $\mathbf{M}$, where each row $\mathbf{m}_{i}$ is sampled with probability proportional to $\left\|\mathbf{m}_{i}\right\|_{2}^{2}$. The result is a matrix $\mathbf{S M}$ such that $\mathbf{M}^{\top} \mathbf{S}^{\top} \mathbf{S M}$ is a "good" approximation to $\mathbf{M}^{\top} \mathbf{M}$.

Theorem 4.3. [DKM06] Suppose $\mathbf{S}$ is a sampling matrix with $r$ rows randomly generated from the approximate matrix multiplication algorithm on input $\mathbf{M} \in \mathbb{R}^{n \times d}$. Then with probability at least $\frac{2}{3}$,

$$
\left\|\mathbf{M}^{\top} \mathbf{S}^{\top} \mathbf{S M}-\mathbf{M}^{\top} \mathbf{M}\right\|_{F}^{2} \leq \frac{10}{\sqrt{r}}\|\mathbf{M}\|_{F}^{4}
$$

The approximate matrix multiplication algorithm is simply squared row norm sampling, which is equivalent to leverage score sampling after preconditioning, since the leverage scores of each row of $\mathbf{A}$ are just the squared row norms of $\mathbf{U}$ in the singular value decomposition $\mathbf{A}=\mathbf{U} \boldsymbol{\Sigma V}$, e.g., see Section 1.3.

### 4.1 Independent Row Selection for a Reduced Matrix

We now describe Algorithm 2, our algorithm for independent row selection from a matrix $\mathbf{B} \in$ $\mathbb{R}^{\mathcal{O}(k \log (k)) \times \mathcal{O}(k)}$ with rank $k$ - we shall ultimately reduce the input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$ down to this case. Our algorithm will iteratively grow a set $S$ of independent rows of $\mathbf{B}$.

We first use leverage score sampling to sample $\mathcal{O}(k)$ rows of $\mathbf{B}$. Although we require $\mathcal{O}(k \log k)$ samples to cover the entire row span of $\mathbf{B}$, we can show using the guarantees of approximate matrix multiplication by Theorem 4.3 that $\mathcal{O}(k)$ samples suffice to cover a span with rank $(1-c) k=\frac{9}{10} k$ with probability $\frac{2}{3}$, for $c=\frac{1}{10}$. We can efficiently compute both an independent subset of these rows and a basis $\mathbf{Z}_{1}$ for their orthogonal complement. We can then add the independent subset to $S$. These procedures in total use at most $\gamma k^{\omega}$ runtime for some fixed constant $\gamma>0$.

For the next iteration, we repeat these procedures on $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$. Namely, we use leverage score sampling to sample $\mathcal{O}(k)$ rows of $\mathbf{B Z}_{1}^{\top}$. We can again efficiently compute both an independent subset of these rows and a basis $\mathbf{Z}_{2}$ for their orthogonal complement. We can again add an independent subset of these rows to our growing set $S$ and then compute a basis $\mathbf{Z}_{2}$ for the orthogonal
complement of $S$. The main idea is that the rows of $\mathbf{B}$ that are spanned by $S$ will have leverage score zero in $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$ because they must be orthogonal to $\mathbf{Z}_{1}$, the orthogonal complement of $S$. Hence, the sampled rows will cover a constant fraction of the remaining subspace orthogonal to $S$ and in particular with probability $\frac{2}{3}$, we can sample a set of rows with rank at least a $\frac{9}{10}$ fraction of the rank of $\mathbf{B Z} \mathbf{Z}_{1}^{\top}$.

Moreover, since the rank of $\mathbf{B Z}_{1}^{\top}$ is at most a $c=\frac{1}{10}$-fraction of the rank of $\mathbf{B}$ conditioned on the success of the first iteration, then the runtime of the second iteration is at most $\gamma(c k)^{\omega}$, which is a constant fraction smaller than the runtime of the first iteration. We can now repeatedly apply this approach by iteratively adding a set of rows to $S$ that cover a constant fraction of the remaining subspace orthogonal to $S$, while using runtime a constant fraction of that in the previous iterations. Since the runtime follows a geometric series in expectation, the overall runtime is $\mathcal{O}\left(k^{\omega}\right)$ in expectation, and so by Markov's inequality, the overall runtime is $\mathcal{O}\left(k^{\omega}\right)$ with constant probability. We remark that the correctness and runtime analysis is robust to failures in each iteration because a failure in an iteration means that at worst, no additional rows are added to $S$, which does not affect the correctness of the algorithm and only slightly increases the runtime, which is absorbed into the computation of the expected runtime.

```
Algorithm 2 Independent Row Selection
Input: \(\mathbf{A} \in \mathbb{R}^{m \times d}\) for \(m=d \operatorname{polylog}(d)\)
Output: A set of \(\operatorname{rank}(\mathbf{A})\) independent rows of \(\mathbf{A}\)
    Use Theorem 1.16 so that \(S\) is a set of \(\mathcal{O}(\operatorname{rank}(\mathbf{A}))\) rows of \(\mathbf{A}\)
    Let SA be the submatrix of \(\mathcal{O}(\operatorname{rank}(\mathbf{A}))\) rows of \(\mathbf{A}\) in \(S\)
    Let \(\mathbf{Z}^{(1)}\) be a basis for the orthogonal complement of \(\mathbf{S A}\)
    Reduce \(S\) to a set of independent rows
    \(t \leftarrow 1\)
    while \(\mathbf{Z}^{(t)}\) is non-empty do
        Use Theorem 1.16 so that \(S^{\prime}\) is a set of \(\mathcal{O}(r)\) rows of \(\mathbf{A Z}{ }^{(1)} \ldots \mathbf{Z}^{(t)}\), where \(r=\)
    \(\operatorname{rank}\left(\mathbf{A} \mathbf{Z}^{(1)} \ldots \mathbf{Z}^{(t)}\right)\)
        \(S \leftarrow S \cup S^{\prime}\)
        Reduce \(S\) to a set of independent rows
        Let SA be the submatrix of \(\mathcal{O}(r)\) rows of \(\mathbf{A}\) in \(S\)
        Let \(\mathbf{Z}^{(t+1)}\) be a basis for the orthogonal complement of SA
    return \(S\)
```

We show that the subroutine Algorithm 2 can be used to find a set of $\operatorname{rank}(\mathbf{A})$ independent rows of $\mathbf{A}$ in matrix-multiplication runtime. We shall ultimately apply Algorithm 2 to a matrix $\mathbf{A} \in \mathbb{R}^{\mathcal{O}(k \log k) \times \mathcal{O}(k)}$.

Lemma 4.4. Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ of rank $k$, there exists an algorithm that with probability at least $\frac{2}{3}$, outputs a set $S$ of $k$ independent rows of $\mathbf{A}$ in time

$$
\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \log d+k^{\omega}\right)
$$

Proof. Suppose A has rank $k$ and let $S$ be a set of $n_{1} \geq \frac{9}{10} k$ rows of A, with rank $r_{1}$. Let $\mathbf{S}$ be the corresponding sampling matrix so that $S$ consists of the rows of $\mathbf{S A} \in \mathbb{R}^{n_{1} \times d}$. Let $\mathbf{Z}^{(1)}$ be a basis for the orthogonal complement of $\mathbf{S A}$, so that $\mathbf{Z}^{(1)} \in \mathbb{R}^{\left(k-n_{1}\right) \times d}$.

Observe that for any row $\mathbf{a}_{i}$ in the span of $S$, we have $\mathbf{a}_{i}\left(\mathbf{Z}^{(1)}\right)^{\top}=0^{\left(k-n_{1}\right)}$, where the right-hand side denotes the all zeros vector of length $k-n_{1}$. Thus the only nonzero rows of $\mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$ are the rows that are independent of $S$ and so we would like to sample rows of A proportional to their leverage score sample in $\mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$. However, due to our desired runtime, we cannot afford to explicitly compute $\mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$. Instead we apply the techniques of Theorem 1.2 to perform leverage score sampling on $\mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$. Namely, we first generate a matrix $\mathbf{G}^{(1)} \in \mathbb{R}^{m_{1} \times n_{1}}$ via Theorem 1.1 and suppose that

$$
\left\|\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top} \mathbf{x}\right\|_{2} \leq\left\|\mathbf{G}^{(1)} \mathbf{S} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top} \mathbf{x}\right\|_{2} \leq \gamma\left\|\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top} \mathbf{x}\right\|_{2},
$$

for an absolute constant $\gamma>1$ and for all $\mathbf{x} \in \mathbb{R}^{d}$. Crucially, $\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top} \in \mathbb{R}^{m_{1} \times d}$, where $m_{1} \leq C\left(k-n_{1}\right)$ and can be computed in time $C\left(\operatorname{nnz}(\mathbf{A})+m_{1} k^{\omega-1}\right)$ for an absolute constant $C>0$. As in Theorem 1.2, we then use $C\left(\mathrm{nnz}(\mathbf{A})+m_{1} k^{\omega-1}\right)$ time to compute a QR decomposition of $\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$ to output matrices $\mathbf{Q}$ and $\mathbf{R}$ such that $\mathbf{Q}$ has orthonormal columns and $\mathbf{Q R}=$ $\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$.

By Lemma 3.1, we have that $\left\|\mathbf{a}_{i}\left(\mathbf{Z}^{(1)}\right)^{\top} \mathbf{R}^{-1}\right\|_{2}^{2}$ is a $\xi^{2}$-approximation to the leverage score of the $i$-th row of $\mathbf{a}_{i}\left(\mathbf{Z}^{(1)}\right)^{\top}$. Thus by Theorem 3.2, we can sample $\mathcal{O}\left(d-n_{1}\right)$ rows of $\mathbf{A}$ with probabilities proportional to the leverage scores of $\mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top}$.

By setting $\mathbf{M}$ in the context of Theorem 4.3 to $\mathbf{M}=\mathbf{G}^{(1)} \mathbf{A}\left(\mathbf{Z}^{(1)}\right)^{\top} \mathbf{R}^{-1}$, we have that $\mathbf{M}^{\top} \mathbf{M}$ is the diagonal matrix consisting of $r_{1}:=d-n_{1}$ ones and zeros elsewhere, since $\mathbf{Q}$ has orthonormal columns. Thus, $\|\mathbf{M}\|_{F}^{2}=r_{1}$ and so by Theorem 4.3, we have that for $r=\mathcal{O}\left(r_{1}\right)$ with probability at least $\frac{2}{3}$,

$$
\left\|\mathbf{M}^{\top} \mathbf{S}^{\top} \mathbf{S M}-\mathbf{M}^{\top} \mathbf{M}\right\|_{F}^{2} \leq \frac{r_{1}}{100}
$$

where $\mathbf{S}$ is the sampling matrix induced by approximate matrix multiplication, which is equivalent to leverage score sampling in this case. On the other hand, we have $\left\|\mathbf{M}^{\top} \mathbf{S}^{\top} \mathbf{S M}-\mathbf{M}^{\top} \mathbf{M}\right\|_{F}^{2} \geq$ $\operatorname{rank}(\mathbf{S M})$, since $\mathbf{M}^{\top} \mathbf{M}$ is a diagonal matrix consisting of only ones and zeros. Therefore, it follows that SM, i.e., the set of rows from leverage score sampling, has found at least a $\frac{9}{10}$ fraction of the remaining independent rows.

By arguing inductively, the algorithm outputs a set of $d$ independent rows. Namely, we use $\mathbf{S}$ to compute the matrix $\mathbf{Z}^{(2)}$ for the orthogonal complement of the sampled rows. Then given a sequence of matrices $\mathbf{Z}^{(1)}, \ldots, \mathbf{Z}^{(i)}$, we generate a matrix $\mathbf{G}^{(i)}$ via Theorem 1.1 and its QR decomposition to iteratively perform leverage score sampling.

Call an iteration successful if the number of remaining independent rows of $\mathbf{A}$ has decreased by at least a $\frac{9}{10}$ fraction. We define a round to be a number of iterations such that the number of remaining independent rows of $\mathbf{A}$ has decreased by at least a $\frac{9}{10}$ fraction. Since each iteration in round $t$ runs in time

$$
C\left(\operatorname{nnz}(\mathbf{A})+\left(\frac{1}{10}\right)^{i-1} k \cdot k^{\omega-1}\right)
$$

and succeeds with probability at least $\frac{2}{3}$, then the expected runtime $R_{t}$ of round $t$ is at most

$$
\mathbb{E}\left[R_{t}\right] \leq C\left(\mathrm{nnz}(\mathbf{A})+\left(\frac{1}{10}\right)^{i-1} k \cdot k^{\omega-1}\right)+\frac{1}{3} \mathbb{E}\left[R_{t}\right] \leq 2 C\left(\mathrm{nnz}(\mathbf{A})+\left(\frac{1}{10}\right)^{i-1} k \cdot k^{\omega-1}\right) .
$$

Thus the total expected runtime is at most

$$
\sum_{i=1}^{\log d}=2 C\left(\mathrm{nnz}(\mathbf{A})+\left(\frac{1}{10}\right)^{i-1} k \cdot k^{\omega-1}\right)=\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \log d+k^{\omega}\right) .
$$

Hence we have by Markov's inequality that with probability at least $\frac{2}{3}$, the algorithm uses total time

$$
\mathcal{O}\left(\mathrm{nnz}(\mathbf{A}) \log d+k^{\omega}\right) .
$$

### 4.2 Input Matrix Reduction

We now show there exists an algorithm for independent row selection that uses the current matrixmultiplication runtime. We would like to reduce from an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$ to a matrix $\mathbf{B} \in \mathbb{R}^{\mathcal{O}(k \log k) \times k}$, which would allow us to apply Algorithm 2 and therefore, Lemma 4.4. To that end, we apply a rank-preserving sketch $\mathbf{S}$ to $\mathbf{A}$, so that $\operatorname{rank}(\mathbf{A})=\operatorname{rank}(\mathbf{A S})$, where $\mathbf{S} \in \mathbb{R}^{d \times c k}$ for some constant $c>0$ and any set of independent rows of $\mathbf{A}$ is also a set of independent rows of AS. We then use our constant-factor subspace embedding to select $\mathcal{O}(k \log k)$ reweighted rows from AS. These reweighted rows form the input matrix $\mathbf{B}$ to Algorithm 2.
Theorem 1.3. Given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank $k$ and any tradeoff parameter $\alpha>0$, there exists an algorithm that outputs a set of $k$ linearly independent rows of $\mathbf{A}$, using time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+k^{\omega}\right)+$ $k^{2+\alpha}$ polylog $(k)$
Proof. Let $\mathbf{S}_{1} \in \mathbb{R}^{c k \times d}$ be a rank-preserving sketch and let $\mathcal{E}$ be the event that $\operatorname{rank}\left(\mathbf{A} \mathbf{S}_{1}^{\top}\right)=$ $\operatorname{rank}(\mathbf{A})=k$. Then by Theorem 4.2, we have $\operatorname{Pr}[\mathcal{E}] \geq 1-\frac{1}{\text { poly }(n)}$.

Conditioned on $\mathcal{E}$, let $I \subseteq[n]$ be a subset of independent rows of $\mathbf{A S} \mathbf{S}_{1}^{\top}$, so that $\operatorname{rank}\left(\mathbf{A}_{I}\right)=k$. Thus to find $k$ linearly independent rows of $\mathbf{A}$, it suffices to find $k$ linearly independent rows of $\mathbf{A S}_{1}^{\top}$. Let $\mathbf{B}=\mathbf{A} \mathbf{S}_{1}^{\top}$, so that $\mathrm{nnz}(\mathbf{B})=\mathcal{O}(\mathrm{nnz}(\mathbf{A}))$ by Theorem 4.2.

By Theorem 1.2, there exists an algorithm that samples $\mathcal{O}(k \log k)$ rows of $\mathbf{B}$ to form a matrix $\mathbf{B}^{\prime}$ such that $\operatorname{rank}(\mathbf{B})=\operatorname{rank}\left(\mathbf{B}^{\prime}\right)=k$, using time

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+k^{\omega}\right)+k^{2+\alpha} \operatorname{polylog}(k),
$$

for any trade-off parameter $\alpha>0$. Thus, we can then apply Lemma 4.4 to $\mathbf{B}^{\prime}$ to compute $k$ linearly independent rows in time $k^{2}$ polylog $(k)+\mathcal{O}\left(k^{\omega}\right)$. Thus, the overall runtime is

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+k^{\omega}\right)+k^{2+\alpha} \operatorname{polylog}(k) .
$$

## 5 Linear Regression

In this section, we show how our $(1+\varepsilon)$-subspace embedding can be used to solve approximate linear regression in the current matrix-multiplication runtime.

We first recall the following statement that shows how a $(1+\sqrt{\varepsilon})$-subspace embedding suffices to achieve a $(1+\mathcal{O}(\varepsilon))$-approximate solution to linear regression.

Lemma 5.1 (Theorem 14 in [BDN15]). Let SA be $a(1+\sqrt{\varepsilon})$-subspace embedding of an input matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$. Then

$$
\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A}-\mathbf{S b}\|_{2} \leq(1+\mathcal{O}(\varepsilon)) \cdot \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A}-\mathbf{b}\|_{2} .
$$

Thus Lemma 5.1 implies that we should simply find an approximate solution to linear regression after applying a $(1+\sqrt{\varepsilon})$-subspace embedding, e.g., Theorem 1.2. However, this is not straightforward because the resulting dimension after applying Theorem 1.2 would be $\mathcal{O}\left(\frac{1}{\varepsilon} d \log d\right)$, which is not small enough to compute the closed-form solution for linear regression in the current matrixmultiplication runtime, due to the extra logarithmic factor. Instead, we consider gradient descent to find an approximately optimal solution for linear regression, recalling the guarantees on the convergence rate of gradient descent in Theorem 1.17.

We now show that there exists an algorithm for approximate linear regression in the current matrix-multiplication runtime. The main approach is similar to that of [CCKW22], but we have a better runtime due to our constant-factor subspace embedding, and also we have a better dependence on $\varepsilon$ due to invoking Lemma 5.1.

The main idea is that Lemma 5.1 states that it suffices to solve approximate linear regression after applying a $(1+\mathcal{O}(\sqrt{\varepsilon}))$-subspace embedding $\mathbf{S}$, which by Theorem 1.2 results in $\mathcal{O}\left(\frac{1}{\varepsilon} d \log d\right)$ rows. Unfortunately, the dimension of $\mathbf{S}$ is too high to find a closed form solution to $\min _{\mathbf{x} \in \mathbb{R}^{d}} \| \mathbf{S A x}-$ $\mathbf{S b} \|_{2}$. On the other hand, since we only require finding a vector $\mathbf{w} \in \mathbb{R}^{d}$ such that $\|\mathbf{S A w}-\mathbf{S b}\|_{2} \leq$ $(1+\mathcal{O}(\varepsilon)) \min _{\mathbf{x}_{\boldsymbol{R}} \mathbb{R}^{d}}\|\mathbf{S A x}-\mathbf{S b}\|_{2}$, we instead use gradient descent to find such a vector $\mathbf{w}$. However, gradient descent requires a small condition number and a "good" initial solution. To decrease the condition number to $\mathcal{O}(1)$, we instead consider $\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A R x}-\mathbf{S b}\|_{2}$, where $\mathbf{G} \mathbf{A}=\mathbf{Q R}^{-1}$ is a QR decomposition for a constant-factor subspace embedding GA. To find a good initial solution, we first find the closed-form solution to $\mathbf{w}^{(0)}=\operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{G A x}-\mathbf{G b}\|_{2}$, since $\mathbf{G A}$ is a constant-factor subspace embedding. We then account for the preconditioning by computing $\mathbf{w}^{(1)}=\mathbf{R}^{-1} \mathbf{w}^{(0)}$, which is a good starting point for our gradient descent because it provides a constant-factor approximation to the optimal solution due to the properties of GA.

Theorem 1.4. Given $\mathbf{A} \in \mathbb{R}^{n \times d}, \mathbf{b} \in \mathbb{R}^{n}$, and any tradeoff parameter $\alpha>0$, there exists an algorithm that with probability at least 0.9 , outputs a vector $\mathbf{y}$ such that

$$
\|\mathbf{A y}-\mathbf{b}\|_{2} \leq(1+\varepsilon) \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{A} \mathbf{x}-\mathbf{b}\|_{2},
$$

using time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon} d^{2+\alpha} \operatorname{poly} \log (d)+\frac{1}{\varepsilon} d^{2} \operatorname{polylog}(d) \log \frac{1}{\varepsilon}$.
Proof. Let $\alpha>0$ be a fixed constant. By Theorem 1.2, we can in time $\mathcal{O}\left(\frac{\operatorname{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+$ $\frac{1}{\varepsilon} d^{2+\alpha} \operatorname{polylog}(d)$, compute a $(1+\mathcal{O}(\sqrt{\varepsilon}))$ subspace embedding $[\mathbf{S A} ; \mathbf{S b}]$ of the matrix $[\mathbf{A} ; \mathbf{b}]$ with probability at least 0.9 . By Theorem 1.1, we can compute a matrix $\mathbf{G}$ such that with probability at least $0.9, \mathbf{G}$ has $\mathcal{O}(d)$ rows and $\mathbf{G A}$ is a $\xi$-distortion subspace embedding with $\xi=\mathcal{O}(1)$. Moreover, since $\mathbf{G A} \in \mathbb{R}^{\mathcal{O}(d) \times d}$, then we can compute its QR decomposition $\mathbf{G A}=\mathbf{Q R}^{-1}$ in $\mathcal{O}\left(d^{\omega}\right)$ time. Because $\mathbf{Q}$ has orthonormal columns, then the condition number of $\mathbf{G A R}$ is $\kappa(\mathbf{G A R})=1$. Since GA is a $\xi$-distortion subspace embedding with $\xi=\mathcal{O}(1)$, it follows that $\kappa(\mathbf{A R})=1$. Similarly, we have $\kappa(\mathbf{S A R})=\mathcal{O}(1)$ since $\mathbf{S}$ is also a $(1+\mathcal{O}(\sqrt{\varepsilon}))$ subspace embedding of $\mathbf{A}$. Intuitively, $\mathbf{R}$ serves as a good preconditioner to the matrix $\mathbf{A}$.

More precisely, let w be an approximate minimizer of the resulting matrix, so that

$$
\|\mathbf{S A R w}-\mathbf{S b}\|_{2} \leq(1+\varepsilon) \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A R x}-\mathbf{S b}\|_{2} .
$$

Then by Lemma 5.1, $\mathbf{R w}$ is a $(1+\mathcal{O}(\varepsilon))$-approximate solution to the linear regression problem, so it remains to compute $\mathbf{R w}$.

Unfortunately, since $\mathbf{S A R}$ has $\mathcal{O}\left(\frac{1}{\varepsilon} d \log d\right)$ rows, we cannot afford to immediately use the closed-form solution to compute $\min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A R x}-\mathbf{S b}\|_{2}$. On the other hand, since $\mathbf{R w}$ is a $(1+\mathcal{O}(\varepsilon))$-approximate solution to the linear regression problem, we can use gradient descent to compute $\mathbf{R w}$ after finding a "good" initial point.

To that end, we first find the closed-form solution to $\mathbf{w}^{(0)}=\operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{G A x}-\mathbf{G b}\|_{2}$, since GA is a constant-factor subspace embedding. We then account for the preconditioning by computing $\mathbf{w}^{(1)}=\mathbf{R}^{-1} \mathbf{w}^{(0)}$, which will serve as a starting point for our gradient descent.

More specifically, let $\mathbf{w}^{(0)}=(\mathbf{G A})^{+}(\mathbf{G b})$, where $(\mathbf{G A})^{+}$is the pseudo-inverse of $\mathbf{G A}$. Since $\mathbf{G A} \in \mathbb{R}^{\mathcal{O}(d) \times d}$, then we can compute its pseudo-inverse in $\mathcal{O}\left(d^{\omega}\right)$ time. Moreover, since the construction of $\mathbf{G}$ in Theorem 1.1 consists of a reweighted subsampled Hadamard transform, then we can compute $\mathbf{G b}$ in $\mathcal{O}(n \log n)=\mathcal{O}(\mathrm{nnz}(\mathbf{A}))$ time. Thus we can compute $\mathbf{w}^{(0)}$ in total time $\mathcal{O}\left(\mathrm{nnz}(\mathbf{A})+d^{\omega}\right)$ after computing $\mathbf{G}$.

We can now compute $\mathbf{w}^{(1)}=\mathbf{R}^{-1} \mathbf{w}^{(0)}$ in $\mathcal{O}\left(d^{2}\right)$ time and furthermore,

$$
\left\|\mathbf{S A R w}{ }^{(1)}-\mathbf{S b}\right\|_{2} \leq(1+\varepsilon)\left\|\mathbf{A R} \mathbf{w}^{(1)}-\mathbf{b}\right\|_{2}=(1+\varepsilon)\left\|\mathbf{A} \mathbf{w}^{(0)}-\mathbf{b}\right\|_{2} \leq(1+\varepsilon)\left\|\mathbf{G A} \mathbf{w}^{(0)}-\mathbf{G b}\right\|_{2} .
$$

Let $\mathbf{z}=\operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A x}-\mathbf{S b}\|_{2}$. Since $\mathbf{w}^{(0)}=(\mathbf{G A})^{+}(\mathbf{G b})$, then $\mathbf{w}^{(0)}=\operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^{d}} \| \mathbf{G A} \mathbf{x}-$ $\mathbf{G b} \|_{2}$. Therefore,

$$
\begin{aligned}
\| \mathbf{S A R w} & (1)-\mathbf{S b} \|_{2}
\end{aligned} \leq(1+\varepsilon)\left\|\mathbf{G A} \mathbf{w}^{(0)}-\mathbf{G b}\right\|_{2} .
$$

In other words, $\mathbf{w}^{(1)}$ is an $\mathcal{O}(1)$-approximation to the optimizer of the linear regression problem for the input matrix $\mathbf{S A R}$ and the measurement vector $\mathbf{S b}$, since $\gamma=\mathcal{O}(1)$. Since $\mathbf{Q}$ is orthonormal, then (squared) linear regression for $[\mathbf{S A R} ; \mathbf{S b}]$ is 1 -strongly convex. Moreover, since $\kappa(\mathbf{S A R})=$ $\mathcal{O}(1)$, then we can set $m=1$ and $M=\mathcal{O}(1)$ in Theorem 1.17. Further, setting the parameters $f\left(x^{(0)}\right) \leq \gamma \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A x}-\mathbf{S b}\|_{2}$ and the gradient descent accuracy $\zeta=\varepsilon \min _{\mathbf{x} \in \mathbb{R}^{d}}\|\mathbf{S A x}-\mathbf{S b}\|_{2}$ in Theorem 1.17, we obtain a $(1+\varepsilon)$-approximation by using $\mathcal{O}\left(\log \frac{1}{\varepsilon}\right)$ iterations of gradient descent with the initial solution as $\mathbf{w}^{(1)}$. Since $\mathbf{S}$ has $\mathcal{O}\left(\frac{1}{\varepsilon} d \log (d)\right)$ rows from Theorem 1.2, each iteration of gradient descent can be performed in time $\frac{1}{\varepsilon} d^{2}$ polylog $d$. Hence, the overall runtime to compute a $(1+\varepsilon)$-approximate solution to the linear regression problem on input matrix $\mathbf{A}$ and measurement vector $\mathbf{b}$ is

$$
\mathcal{O}\left(\frac{\mathrm{nnz}(\mathbf{A})}{\alpha}+d^{\omega}\right)+\frac{1}{\varepsilon} d^{2+\alpha} \text { polylog}(d)+\frac{1}{\varepsilon} d^{2} \text { polylog }(d) \log \frac{1}{\varepsilon} .
$$

Acknowledgments. We thank Jelani Nelson for helpful discussions over the course of the project. David P. Woodruff and Samson Zhou were supported by a Simons Investigator Award and by the National Science Foundation under Grant No. CCF-1815840. Sandeep Silwal is supported by an NSF Graduate Research Fellowship under Grant No. 1745302, and NSF TRIPODS program (award DMS-2022448), NSF award CCF-2006798, and Simons Investigator Award (via Piotr Indyk).

## References

[AW21] Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 522-539, 2021. 1
[BDN15] Jean Bourgain, Sjoerd Dirksen, and Jelani Nelson. Toward a unified theory of sparse dimensionality reduction in euclidean space. In Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC, pages 499-508. ACM, 2015. 27
[CCKW22] Nadiia Chepurko, Kenneth L. Clarkson, Praneeth Kacham, and David P. Woodruff. Near-optimal algorithms for linear algebra in the current matrix multiplication time. In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 3043-3068, 2022. 2, 3, 4, 7, 21, 27
[CKL13] Ho Yee Cheung, Tsz Chiu Kwok, and Lap Chi Lau. Fast matrix rank algorithms and applications. J. ACM, 60(5):31:1-31:25, 2013. 4, 23
[CN22] Yeshwanth Cherapanamjeri and Jelani Nelson. Uniform approximations for randomized hadamard transforms with applications. In STOC: 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 659-671, 2022. 3, 5, 11
[Coh16] Michael B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 278-287, 2016. 10
[CW82] Don Coppersmith and Shmuel Winograd. On the asymptotic complexity of matrix multiplication. SIAM J. Comput., 11(3):472-492, 1982. 2
[CW13] Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In Symposium on Theory of Computing Conference, STOC, pages 81-90, 2013. 10
[DKM06] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast monte carlo algorithms for matrices I: approximating matrix multiplication. SIAM J. Comput., 36(1):132-157, 2006. 23
[DL22] Jules Depersin and Guillaume Lecué. Robust sub-Gaussian estimation of a mean vector in nearly linear time. Ann. Statist., 50(1):511-536, 2022. 17
[DMM06a] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-based methods. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX and 10th International Workshop on Randomization and Computation, RANDOM, Proceedings, pages 316-326, 2006. 12
[DMM06b] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Subspace sampling and relative-error matrix approximation: Column-row-based methods. In Algorithms ESA 2006, 14th Annual European Symposium, Proceedings, pages 304-314, 2006. 12
[DMMW12] Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. J. Mach. Learn. Res., 13:3475-3506, 2012. 21
[Fos53] Frederic G Foster. On the stochastic matrices associated with certain queuing processes. The Annals of Mathematical Statistics, 24(3):355-360, 1953. 11
[Ind07] Piotr Indyk. Uncertainty principles, extractors, and explicit embeddings of 12 into 11. In Proceedings of the 39th Annual ACM Symposium on Theory of Computing, pages 615-620, 2007. 4
[LM19] Gábor Lugosi and Shahar Mendelson. Sub-Gaussian estimators of the mean of a random vector. Ann. Statist., 47(2):783-794, 2019. 16
[LS17] Yin Tat Lee and He Sun. An sdp-based algorithm for linear-sized spectral sparsification. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 678-687, 2017. 3
[LS18] Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almostlinear time. SIAM J. Comput., 47(6):2315-2336, 2018. 3
[LT91] Michel Ledoux and Michel Talagrand. Probability in Banach spaces, volume 23 of Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1991. Isoperimetry and processes. 9
[ $\left.\mathrm{M}^{+} 89\right]$ Colin McDiarmid et al. On the method of bounded differences. Surveys in combinatorics, 141(1):148-188, 1989. 9
[Mag10] Malik Magdon-Ismail. Row sampling for matrix algorithms via a non-commutative bernstein bound. CoRR, abs/1008.0587, 2010. 12
[NN13] Jelani Nelson and Huy L. Nguyen. OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. In 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS, pages 117-126, 2013. 2, 4, 5, 10
[PTZ12] Richard Peng, Kanat Tangwongsan, and Peng Zhang. Faster and simpl semidefinite programming. In 24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, pages 101-108, 2012. 15
[Sin16] Yaron Singer. Lecture notes. http://people.seas.harvard.edu/~yaron/AM221S16/lecture_notes/AM221_lecture9.pdf, 2016. 12
[Tro08] Joel A Tropp. Norms of random submatrices and sparse approximation. Comptes Rendus Mathematique, 346(23-24):1271-1274, 2008. 5
[Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. Found. Comput. Math., 12(4):389-434, 2012. 10
[Ver18] Roman Vershynin. High-dimensional probability: An introduction with applications in data science, volume 47. Cambridge university press, 2018. 9, 19
[Woo14] David P. Woodruff. Sketching as a tool for numerical linear algebra. Found. Trends Theor. Comput. Sci., 10(1-2):1-157, 2014. 3, 4, 12, 21


[^0]:    *UC Berkeley. E-mail: yeshwanth@berkeley.edu
    ${ }^{\dagger}$ MIT. E-mail: silwal@mit.edu
    ${ }^{\ddagger}$ Carnegie Mellon University. E-mail: dwoodruf@cs.cmu.edu
    ${ }^{\S}$ UC Berkeley and Rice University. Work done in part while at Carnegie Mellon University. E-mail: samsonzhou@gmail.com

[^1]:    ${ }^{1}$ It is known that even a $(1+\mathcal{O}(\sqrt{\varepsilon}))$-approximate subspace embedding suffices, see Lemma 5.1. To facilitate the intuition for our algorithm, we defer this discussion to Section 5.

