## Lecture 11 - Applications of Dimension Reduction

Instructor: Alex Andoni
Scribes: Marshall Ball

Today we looked at two applications of dimension reduction for improving the time complexity of two classical problems in P :
(a) Matrix Multiplication.
(b) Least Square Regression

As usual, we will make our lives easier by considering approximate variants (of the optimization versions) of the above problems.

## 1 Matrix Multiplicaiton

Definition 1. (Exact) Matrix Multiplication is the following problem:

- Given $A, B \in \mathbb{R}^{n \times d}$,
- Compute: $C=A^{\top} B \in \mathbb{R}^{d \times d}$.

In general, you may consider the problem for arbitrary fields $\mathcal{K}$, but we will restrict our attention to $\mathbb{R}$. (One may also consider matrices of arbitrary dimension.)

Naively, we can solve the above problem in time $O\left(n d^{2}\right)$. The state of the art for $n \times n$ matrices is time $O\left(n^{\omega}\right)$ for $\omega \approx 2.36 \ldots$ This will yield an algorithm for our problem with complexity $O\left(d^{2} n^{\omega-2}\right)$.

However as usual, we are interested in a near linear time, $\sim O(n d)$, algorithm. To do this exactly is hard, so we will relax the problem to an approximate version.

First, we define the following norm to characterize our approximation guarrantee;
Definition 2. For a matrix $Z \in \mathbb{R}^{m \times n}$, the (squared) frobenius norm is defined as follows:

$$
\|Z\|_{F}^{2}:=\sum_{i, j} Z_{i, j}^{2}
$$

Definition 3. (Approximate) Matrix Multiplication is the following problem:

- Given $A, B \in \mathbb{R}^{n \times d}$,
- Compute: $C^{\prime} \in \mathbb{R}^{d}$ such that the following holds with high probility,

$$
\left\|C^{\prime}-A^{\top} B\right\|_{F} \leq \varepsilon\|A\|_{F} \times\|B\|_{F}
$$

Some notation for what follows:

$$
A=\left[\begin{array}{c}
x_{1}^{\top} \\
\vdots \\
x_{n}^{\top}
\end{array}\right] \quad B=\left[\begin{array}{c}
y_{1}^{\top} \\
\vdots \\
y_{n}^{\top}
\end{array}\right]
$$

### 1.1 A First Algorithm: Sampling via a Horovitz-Thompson Estimator

We begin by noting the following:
Claim 1. $A^{\top} B=\sum_{k=1}^{n} x_{k} y_{k}^{\top}$ ( $x y^{\top}$ is the "outer-product" of vectors $x$ and $y$ ). Proof.

$$
C_{i j}=\left(\sum_{k=1}^{n} x_{k} y_{k}^{t}\right)_{i j}=\sum_{k=1}^{n} x_{k i} y_{k j} .
$$

From this, we derive the following algorithm (we will fix parameters in the analysis):

- Sample $m$ coordinates $k_{t}$ from $[n]$ ( $2 m$ vectors: $x_{k_{t}}, y_{k_{t}}, t \in[m]$ ) where the probability of sampling coordinate $k$ is $p_{k} \propto\|x\|_{k}\|y\|_{k}$.
- Then simply output,

$$
C^{\prime}=\sum_{t=1}^{m} \frac{x_{k_{t}} y_{k_{t}}}{p_{k_{t}}}
$$

## Theorem 2.

$$
\operatorname{Pr}\left[\left\|C^{\prime}-C\right\|_{F}>\varepsilon\|A\|_{F}\|B\|_{F}\right]<\frac{1}{\varepsilon^{2} m} .
$$

Notice that this means we can take $m=\Omega\left(1 / \varepsilon^{2}\right)$.
Proof. • Expectation

$$
\begin{aligned}
\mathbb{E}\left[C^{\prime}\right] & =\frac{1}{m} \mathbb{E}\left[\sum_{t=1}^{m} \frac{x_{k_{t}} y_{k_{k}}^{\top}}{p_{k_{t}}}\right] \\
& =\frac{1}{m} \sum_{t=1}^{m} \sum_{k=1}^{n} \frac{p_{k} x_{k} y_{k}^{\top}}{p_{k}} \\
& =\sum_{k=1}^{n} x_{k} y_{k}^{\top}=C .
\end{aligned}
$$

- Variance

$$
\begin{aligned}
V & =\mathbb{E}\left[\left\|C^{\prime}-C\right\|_{F}^{2}\right] \\
& =\mathbb{E}\left[\sum_{i, j}\left(C_{i j}^{\prime}-C_{i j}\right)^{2}\right] \\
& =\sum_{i, j} \operatorname{Var}\left[C_{i j}^{\prime}\right] \\
& \leq \sum_{i, j} \operatorname{Var}[\frac{1}{m} \sum_{t=1}^{m} \underbrace{\frac{x_{k t i} y_{k t j}}{p_{k_{t}}}}_{\text {id. dist. var. }}] \\
& =\sum_{i, j} \frac{1}{m} \operatorname{Var}\left[\frac{x_{k i} y_{k j}}{p_{k}}\right] \quad \text { (randomness over } k \text { ) } \\
& \leq \frac{1}{m} \sum_{i, j} \mathbb{E}\left[\left(\frac{x_{k i} y_{k j}}{p_{k}}\right)^{2}\right] \\
& =\frac{1}{m} \sum_{i, j} \sum_{k=1}^{n} p_{k}\left(\frac{x_{k i} y_{k j}}{p_{k}}\right)^{2} \\
& =\frac{1}{m} \sum_{k=1}^{n} \frac{1}{p_{k}} \sum_{i, j} x_{k i}^{2} y_{k j}^{2} \\
& =\frac{1}{m} \sum_{k=1}^{n} \frac{1}{p_{k}}\left\|x_{k}\right\|_{F}^{2}\left\|y_{k}\right\|_{F}^{2}
\end{aligned}
$$

So, take

$$
p_{k}:=\frac{\left\|x_{k}\right\|_{F}\left\|y_{k}\right\|_{F}}{\sum_{i=1}^{n}\left\|x_{i}\right\|_{F}\left\|y_{i}\right\|_{F}} .
$$

Then (via Cauchy-Schwartz),

$$
V \leq \frac{\left(\sum_{k=1}^{n}\left\|x_{k}\right\|_{F}\left\|y_{k}\right\|_{F}\right)^{2}}{m} \leq \frac{\left(\sum_{k=1}^{n}\left\|x_{k}\right\|_{F}^{2}\right)\left(\sum_{k=1}^{n}\left\|y_{k}\right\|_{F}^{2}\right)}{m}=\frac{\|A\|_{F}^{2}\|B\|_{F}^{2}}{m}
$$

- So applying Chebyshev to the above,

$$
\operatorname{Pr}\left[\left\|C^{\prime}-C\right\|_{F}^{2}>\varepsilon^{2}\|A\|_{F}^{2}\|B\|_{F}^{2}\right] \leq \frac{\mathbb{E}\left[\left\|C^{\prime}-C\right\|_{F}\right]}{\varepsilon^{2}\|A\|_{F}^{2}\|B\|_{F}^{2}} \leq \frac{1}{m \varepsilon^{2}}
$$

### 1.2 A Second Algorithm: Using Dimension Reduction

Note 1. We can view the above algorithm as the following:

- Choose a random $\Pi \in \mathbb{R}^{m \times n}$ where

$$
\Pi_{i, j}:=\left\{\begin{aligned}
\frac{1}{\sqrt{m p_{k}}} & \text { if }(i, j)=\left(t, k_{t}\right) \\
0 & \text { otherwise }
\end{aligned}\right.
$$

- Compute:

$$
C^{\prime}=(\Pi A)^{\top}(\Pi B) .
$$

Observe that the above algorithm requires two passes over the data, one to sample $\Pi$ (compute the $p_{k}$ 's) and one to compute the "reduced" matrix product (or the sum in our previous formulation).

Given this "randomized-projection/embedding" formulation of our approximation algorithm, it seems an appropriate place to invoke the magic of Johnson-Lindenstrauss. Consider the following definition:

Definition 4. $\Pi \in \mathbb{R}^{m \times n}$ is an $(\varepsilon, \delta)$-dimension reducing matrix, $(\varepsilon, \delta)$-DR, if

$$
\forall x \in \mathbb{R}^{n}, \operatorname{Pr}\left[\left|\|\Pi x\|_{2}^{2}-\|x\|_{2}^{2}\right|>\varepsilon\|x\|_{2}^{2}\right] \leq \delta .
$$

Given some $(\varepsilon, \delta)$-DR matrix $\Pi$, our algorithm is to simply compute:

$$
C^{\prime}=(\Pi A)^{\top}(\Pi B) .
$$

Theorem 3. $\Pi$ is $(\varepsilon, \delta)-D R \Longrightarrow \operatorname{Pr}\left[\left\|C^{\prime}-C\right\|_{F}>3 \varepsilon\|A\|_{F}\|B\|_{F}\right] \leq 3 d^{2} \delta$.
Remark 1. With a more precise version of the Johnson-Lindenstrauss lemma we can remove the $d^{2}$ factor from the above.

Corollary 4. If we choose $m=O\left(1 / \varepsilon^{2} \log (1 / \delta)\right), \delta=\frac{1}{10 d^{2}}$, then (naively) we can compute $C^{\prime}$ in time $O(m n d)+O(d m d)=O\left(\frac{n d+d^{2}}{\varepsilon^{2}} \log d\right)$.

By the above remark, the $\log d$ factor is simply an artifact of our analysis.
To prove the theorem, we will show $C_{i j}^{\prime} \approx C_{i j}$ with probability $\geq 1-3 \delta$ and then take a union bound (hence the $d^{2}$ ).

Proof. First some notation:

$$
\begin{gathered}
A=\left[\begin{array}{lll}
A_{1} & \cdots & A_{d}
\end{array}\right] \quad B=\left[\begin{array}{lll}
B_{1} & \cdots & B_{d}
\end{array}\right] \\
a_{i}:=\frac{A_{i}}{\left\|A_{i}\right\|_{2}}
\end{gathered} b_{i}:=\frac{B_{i}}{\left\|B_{i}\right\|_{2}}, ~ \$
$$

Note:

- $C_{i j}=A_{i}^{\top} B_{j}=\left\|A_{i}\right\|\left\|B_{j}\right\| a_{i}^{\top} b_{j}$.
- With probability $\geq 1-3 \delta$,

$$
\begin{aligned}
C_{i j}^{\prime}=\left(\Pi A_{i}\right)^{\top}\left(\Pi B_{j}\right) & =\left\|A_{i}\right\|\left\|B_{j}\right\|\left(\Pi a_{i}\right)^{\top}\left(\Pi b_{j}\right) \\
& =\left\|A_{i}\right\|\left\|B_{j}\right\|\left[\left\|\Pi a_{i}\right\|^{2}+\left\|\Pi b_{j}\right\|^{2}-\frac{1}{2}\left\|\Pi a_{i}-\Pi b_{j}\right\|^{2}\right] \\
& =\left\|A_{i}\right\|\left\|B_{j}\right\|\left[\left\|a_{i}\right\|^{2}+\left\|b_{j}\right\|^{2}-\left\|a_{i}-b_{j}\right\|^{2} \pm 3 \varepsilon\right] \quad(\Pi \text { is }(\varepsilon, \delta)-\mathrm{DR}) \\
& =\left\|A_{i}\right\|\left\|B_{j}\right\|\left[a_{i} b_{j} \pm 3 \varepsilon\right]
\end{aligned}
$$

So with probability $\geq 1-3 \delta,\left(C_{i j}^{\prime}-C_{i j}\right)^{2} \leq\left\|A_{i}\right\|_{2}^{2}\left\|B_{j}\right\|_{2}^{2}(3 \varepsilon)^{2}$. This implies (via union bound) that with probability $\geq 1-3 \delta d^{2}$,

$$
\left\|C^{\prime}-C\right\|_{F} \leq \sum_{i j} 9 \varepsilon^{2}\left\|A_{i}\right\|^{2}\left\|B_{j}\right\|^{2}=9 \varepsilon^{2}\|A\|_{F}^{2}\|B\|_{F}^{2}
$$

## 2 Least Squares Regression

Definition 5. (Exact) Least Squares Regression is the following problem:

- Given $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^{n}$,
- find $x^{*}=\operatorname{argmin}_{x \in \mathbb{R}^{d}}\|A x-b\|_{2}$.

We can consider least squares regression as a simple learning problem where the $i$-th row of $A, a^{(i)}$, is labeled with $b_{i}$ according to some approximately linear function.

Definition 6. A function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is linear if

$$
\exists y \in \mathbb{R}^{d}: f(x)=\langle x, y\rangle .
$$

If $\exists x: A x=b$ then the problem is easy. In general, we are only assume $\exists x: A x \approx b$.
In general, we can do least squares regression via Singular Value Decomposition (in time $\tilde{O}\left(n d^{\omega-1}\right)$ ), but perhaps we can speed things up by loosening the approximation.

Definition 7. (Approximate) Least Squares Regression is the following problem:

- Given $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^{n}$,
- Let $x^{*}=\operatorname{argmin}_{x \in \mathbb{R}^{d}}\|A x-b\|_{2}$. Find $x \in \mathbb{R}^{d}$ such that

$$
\|A x-b\|_{2} \leq(1+\varepsilon)\left\|A x^{*}-b\right\|_{2} .
$$

To solve this problem we will use dimension reduction, as promised.
First, we define a special kind of dimension reducing matrix:
Definition 8. $\Pi \in \mathbb{R}^{m \times n}$ is a $(d, \varepsilon, \delta)$-subspace embedding, $(d, \varepsilon, \delta)$-SE, if $\forall P \subset \mathbb{R}^{n}$ such that $P$ is a $d$-dimensional subspace,

$$
\operatorname{Pr}[\forall p \in P:|\|\Pi p\|-\|p\|| \leq \varepsilon\|p\|] \geq 1-\delta .
$$

Then given some such SE $\Pi$, our alorithm is simply: find $\operatorname{argmin}_{x}\|\Pi A x-\Pi b\|$ (via SVD).
Naively, the time to reduce dimension is $O(m n d)$. The time to perform SVD on the result is $O\left(m d^{\omega-1}\right.$. So if we take $m=O\left(d / \varepsilon^{2}\right)$, then the resulting algorithm has time complexity

$$
O\left(\frac{n d^{2}}{\varepsilon}+m d^{\omega-1}\right) .
$$

If we use a faster version of Johnson-Lindenstrauss, we can acheive $O_{\varepsilon}\left(\left(n \log n+m^{3}\right) d\right)$ time complexity.

Unfortunately, at this point we ran out of time. We will finish up this application of dimension reduction next lecture.

