## A Proof of Lemma 3.1

Proof of Lemma 3.1. The algorithm is straightforward: choose a random point in $S$, and check if strictly more than $n / 2$ points lie within a ball of radius $2 r$ around this point. If so, include all points with distance at most $4 r$ from this point. Note that we cannot calculate the $\ell_{2}$ distance directly, instead, using JL-lemma, we project all points onto $\mathbb{R}^{O(\log d)}$. If not, repeat, and run for $O(\log 1 / \delta)$ iterations.
Similar to proof of Lemma E.5, let $J \in \mathbb{R}^{r \times d^{2}}$ matrix whose each entries are i.i.d. entries from $\mathcal{N}(0,1 / r)$ where $r=O(\log d)$. Note that $Z v=\left(Y \operatorname{diag}(v) Y^{\top}\right)^{b}$, then we can compute $J \cdot Z$ using fact rectangular matrix multiplication by multiply each row of $J$ to $Z$. Then, this takes $\widetilde{O}(T(N, d))$ time. Note that for each iteration, we compute $O(N)$ many $\ell_{2}$ distance, which takes time $\widetilde{O}(N)$. Thus, the total time is $\widetilde{O}(T(N, d)+N \log (1 / \delta))=\widetilde{O}(T(N, d) \log (1 / \delta))$.
By the triangle inequality, if we ever randomly select a point from $S^{\prime}$, then we terminate, and in this case it is easy to see that the output satisfies the desired property. Thus, it is easy to see that the probability we have not terminated after $t$ iterations is at most $2^{-t}$. Suppose we have terminated. Then in that iteration, we selected a point $X \in S$ that has distance at most $2 r$ to more than $n / 2$ other points in $S$. This implies that it has distance at most $2 r$ to some point in $S^{\prime}$. By triangle inequality, this implies that all points in $S^{\prime}$ are at distance at most $4 r$ from $X$, and so the output in this iteration must satisfy the claims of the Lemma.

## B Proof of Lemma 4.3

Before we prove Lemma 4.3, we need the following preliminaries.
Lemma B. 1 (Hanson-Wright). Let $X_{1}, X_{2}, \ldots X_{N}$ be i.i.d. random vectors in $\mathbb{R}^{d}$ where $X_{i} \sim$ $\mathcal{N}(0, \Sigma)$ and $\Sigma \preceq I$. Let $U \in \mathbb{R}^{d \times d}$ and $U \succeq 0$ and $\|U\|_{F}=1$. Then, there exists a universal constant $C$ so that for all $T>0$, we have

$$
\operatorname{Pr}\left[\left|\frac{1}{N} \sum_{i=1}^{n} \operatorname{tr}\left(X_{i} X_{i}^{\top} U\right)-\operatorname{tr}(U)\right|>T\right] \leq 2 \exp \left(-C N \min \left(T, T^{2}\right)\right)
$$

Corollary B.2. Under the same setting as Lemma B.1, let $Z_{i}=X_{i} \otimes X_{i}$ and $v \in \mathbb{R}^{d^{2}}$ be an arbitrary unit vector. Then, there exists a universal constant $C>0$ so that for all $T>0$, we have

$$
\operatorname{Pr}\left[\left|\frac{1}{N} \sum_{i=1}^{n}\left\langle v, Z_{i}\right\rangle-\mathbb{E}[\langle v, Z\rangle]\right|>T\right] \leq 2 \exp \left(-C N \min \left(T, T^{2}\right)\right)
$$

Proof. This follows by letting the $v$ in the statement be the flattening of $U$ in Lemma B.1.
Using a standard $\varepsilon$-net argument (see e.g. [Ver10]), we get the following concentration bounds for the empirical mean.
Lemma B.3. Under the same setting as Lemma B.1, let $Z_{i}=X_{i} \otimes X_{i}$ there exist universal constants $A, C>0$ so that for all $T>0$, we have

$$
\operatorname{Pr}\left[\left\|\frac{1}{N} \sum_{i=1}^{n} Z_{i}-\mu_{Z}\right\|_{2}>T\right] \leq 2 \exp \left(A d^{2}-C N \min \left(T, T^{2}\right)\right)
$$

Now, we are ready to prove Lemma 4.3.
Proof of Lemma 4.3. The parameter of $\gamma_{1}$ directly follows from Lemma B. 3 by choosing $T=$ $\varepsilon \sqrt{\log (1 / \varepsilon)}$ and $N=\Omega\left(\frac{d^{2}}{\varepsilon^{2} \log (1 / \varepsilon)}\right)$.
For $\beta_{1}$, we apply Lemma B. 3 for any fixed set $S \subset[N]$ of size $2 \varepsilon N$, we have

$$
\operatorname{Pr}\left[\left\|\frac{1}{|S|} \sum_{i \in S} Z_{i}-\mu_{Z}\right\|_{2}>T\right] \leq 2 \exp \left(A d^{2}-C \varepsilon N \min \left(T^{2}, T\right)\right)
$$

Taking the union bound over all subsets of size $2 \varepsilon N$, we get

$$
\begin{aligned}
& \operatorname{Pr}\left[\exists S:|S|=2 \varepsilon N \text { and }\left\|\frac{1}{|S|} \sum_{i \in S} Z_{i}-\mu_{Z}\right\|_{2}>T\right] \\
\leq & 2 \exp \left(A d^{2}+\log \binom{2 \varepsilon N}{N}-C \varepsilon N \min \left(T^{2}, T\right)\right) \\
\leq & 2 \exp \left(A d^{2}+O(N \cdot \varepsilon \log 1 / \varepsilon)-C \varepsilon N \min \left(T^{2}, T\right)\right)
\end{aligned}
$$

By choose $T=O(\log (1 / \varepsilon))$ and $N=\Omega\left(\frac{d^{2}}{\varepsilon^{2} \log (1 / \varepsilon)}\right)$, we have the probability above is less than $O\left(d^{-3}\right)$.
For proof of $\gamma_{2}$ and $\beta_{2}$, see Theorem 4.13 of $\left[\mathrm{DKK}^{+} 16\right]$ and Proposition A. 28 of [DKK ${ }^{+}$17].

## C Proof of Lemma 4.4

Before we prove this lemma, we require the following pair of technical lemmata. The first is a standard fact about the covariance of $X \otimes X$ for $X$ Gaussian.
Fact C. 1 (see e.g. [CDGW19]). Let $X \sim \mathcal{N}(0, \Sigma)$. Then the covariance of $X \otimes X$ is $2 \Sigma \otimes \Sigma$.
This implies:
Lemma C.2. Let $X \sim \mathcal{N}(\mathbf{0}, \Sigma)$ and $Z=X \otimes X$. Let $\Sigma_{Z} \in \mathbb{R}^{d^{2} \times d^{2}}$ be the covariance matrix of $Z$. We have:

1. If $\Sigma \preceq I$, then $\Sigma_{Z} \preceq 2 I$.
2. If $\|\bar{\Sigma}-I\| \leq \zeta$ for $0 \leq \zeta<1$, then $\left\|\Sigma_{Z}-2 I\right\| \leq 6 \zeta$.

Proof. The first claim follows directly from Fact C.1, as $\left\|\Sigma_{Z}\right\|=2\|\Sigma\| \leq 2$. To prove the second statement, note that if $\lambda_{1}, \ldots, \lambda_{d}$ are the eigenvalues of $\Sigma$, the assumption implies that $\left|\lambda_{i}-1\right| \leq \zeta$ for all $i$, and also that $\lambda_{i}<2$ for all $i$. But all the eigenvalues of $\Sigma_{Z}$ are given by $2 \lambda_{i} \lambda_{j}$ for $i, j \in[d]$, and $\left|2 \lambda_{i} \lambda_{j}-2\right| \leq 2\left(\zeta^{2}+\left|\lambda_{i}\right| \zeta+\left|\lambda_{j}\right| \zeta\right) \leq 6 \zeta$. This proves the claim.

Proof of Lemma 4.4. The first claim follows because $\left(\gamma_{1}, \gamma_{2}\right)$-goodness is affine invariant. We now turn our attention to the second claim. First, we show that the $\gamma_{1}$ and $\beta_{1}$ parameters are changed by at most a constant multiplicative factor. Since $X_{i}=\Sigma^{-1 / 2} \bar{X}_{i}$, then we have

$$
Z_{i}=\left(\Sigma^{1 / 2} \otimes \Sigma^{1 / 2}\right)\left(\bar{X}_{i} \otimes \bar{X}_{i}\right)=\left(\Sigma^{1 / 2} \otimes \Sigma^{1 / 2}\right) \bar{Z}_{i}
$$

Then, we have

$$
\gamma_{1}\left(Z_{i}\right)=\left\|\mu\left(Z_{i}\right)-\mu_{Z_{i}}\right\|_{2}=\left\|\left(\Sigma^{1 / 2} \otimes \Sigma^{1 / 2}\right)\left(\mu\left(\bar{Z}_{i}\right)-\mu_{\bar{Z}_{i}}\right)\right\|_{2} \leq O(\varepsilon \log 1 / \varepsilon)
$$

where the last step follows by $\|A \otimes B\| \leq\|A\|\|B\|$ and $\|\Sigma\| \leq 2$. Similarly, we have that the $\beta_{1}$ parameter increases by at most a constant multiplicative factor.
Now, we consider the second moment parameters $\gamma_{2}$ and $\beta_{2}$. Note that we have $Z_{i} Z_{i}^{\top}=(\Sigma \otimes$ $\Sigma)^{1 / 2} \bar{Z}_{i} \bar{Z}_{i}^{\top}(\Sigma \otimes \Sigma)^{1 / 2}$. By the goodness of $\bar{Z}_{i}$, then we have

$$
\left\|\frac{1}{|S|} \sum_{i \in S}\left(Z_{i}-\mu(S)\right)\left(Z_{i}-\mu(S)\right)^{\top}-2(\Sigma \otimes \Sigma)\right\| \leq\|\Sigma\|^{2} \cdot O(\varepsilon \sqrt{\log 1 / \varepsilon})=O(\varepsilon \sqrt{\log 1 / \varepsilon})
$$

Then, by Lemma C.2, we have $\|\Sigma \otimes \Sigma-2 I\| \leq 6 \zeta$, so by, we have

$$
\left\|\frac{1}{|S|} \sum_{i \in S}\left(Z_{i}-\mu(S)\right)\left(Z_{i}-\mu(S)\right)^{\top}-2 I\right\| \leq O(\varepsilon \sqrt{\log 1 / \varepsilon})+6 \zeta,
$$

as claimed. The bound on the $\beta_{2}$ parameter is identical, and omitted.

## D Proof of Theorem 4.6 and Theorem 4.7

## D. 1 Approximate Score Oracles for Tensor Inputs

A key algorithmic ingredient to implementing both Theorem 4.6 and Theorem 4.7 will be the following. We will be given access to a set of points $X_{1}, \ldots, X_{N}$, and we will need approximate augmented score oracles for the tensored set of points $X_{1} \otimes X_{1}, \ldots, X_{N} \otimes X_{N}$. Note that we cannot even afford to write down the tensored versions of the $X_{i}$ in the desired runtime. Despite this, we show that we can construct these approximate augmented score oracles very efficiently:
Theorem D.1. Let $\delta>0$. Let $X_{1}, \ldots, X_{N} \in \mathbb{R}^{d}$, and let $Z_{i}=X_{i} \otimes X_{i}$ for all $i=1, \ldots, N$. Let $t>0$, and let $w_{1}, \ldots, w_{t} \in \Gamma_{N}$. Let $\alpha$ be such that $\alpha$ satisfies Equation (3). Then, there is an algorithm APPROXIMATESCORE which takes as input $\alpha, \delta,\left\{X_{1}, \ldots, X_{n}\right\}$, and $w_{1}, \ldots, w_{t}$, which runs in time $\widetilde{O}\left(t^{2} \cdot T(N, d) \log 1 / \delta\right)$, and with probability $1-\delta$, is an approximate augmented score oracle for $\left\{Z_{1}, \ldots, Z_{N}\right\}$ with weights $w_{1}, \ldots, w_{t}$.

We defer the proof to Section E.

## D. 2 Getting $O(\sqrt{\varepsilon})$ error

In this section, we describe and analyze the routine FIRSTPHASE, which achieves a coarse estimate of the true covariance. We restate the theorem below for convenience.

Theorem 4.6 (First Phase). Let $S$ be a set of points satisfying Assumption 4.2. Moreover, let $\Sigma_{t} \in \mathbb{R}^{d \times d}$ be such that $\Sigma \preceq \Sigma_{t}$. Then there is an algorithm FIRSTPHASE, which given $S$ and $\Sigma_{t}$, runs in time $\widetilde{O}(T(N, d) \log \log \kappa)$ and outputs a new upper bound matrix $\Sigma_{t+1}$ and a approximate covariance matrix $\widehat{\Sigma}$ such that, with probability $1-1 / \operatorname{poly}(d, \log \kappa)$,

$$
\Sigma \preceq \Sigma_{t+1} \preceq \Sigma+O(\sqrt{\varepsilon}) \Sigma_{t}, \quad \text { and } \quad\|\widehat{\Sigma}-\Sigma\|_{F} \leq O(\sqrt{\varepsilon})\left\|\Sigma_{t}\right\|
$$

We give the pseudocode for FIRSTPHASE in Algorithm 2. Our algorithm is simple: we simply run a naive pruning step on the tensored inputs, then apply Lemma 3.3 to the remaining tensored inputs.

```
Algorithm 2 Robust Covariance Estimation With Bounded Covariance
    procedure \(\operatorname{FirstPhASE}\left(S=\left\{X_{1}, X_{2}, \ldots, X_{n}\right\}, \varepsilon, \Sigma_{t}\right) \quad \triangleright\) Theorem 4.6
        for \(i=1, \ldots, N\) do
            \(Y_{i} \leftarrow \Sigma_{t}^{-1 / 2} X_{i}\)
            \(Z_{i} \leftarrow Y_{i} \otimes Y_{i}\)
        end for
        \(S^{\prime} \leftarrow \operatorname{NaivePrune}\left(Z_{1}, \ldots, Z_{N}, 4 d^{2} N^{2}, 1 /(d N \log \kappa)\right) \quad \triangleright\) Lemma 3.1
        Let \(\mathscr{O}_{\text {aug }}=\) APPROXIMATESCORE with \(\delta=1 / \operatorname{poly}(d, \log \kappa)\)
        Let \(\widetilde{\Sigma}\) be the estimate of \(S^{\prime}\) computed by Lemma 3.3 with score oracle \(\mathscr{O}_{\text {aug }}\).
        \(\widehat{\Sigma} \leftarrow \Sigma_{t}^{1 / 2} \widetilde{\Sigma} \Sigma_{t}^{1 / 2}\).
        \(\Sigma_{t+1} \leftarrow \widehat{\Sigma}+O(\sqrt{\varepsilon}) \Sigma_{t}\)
        return \(\widehat{\Sigma}, \Sigma_{t+1}\)
    end procedure
```

Proof of Theorem 4.6. By Assumption 4.2 and Lemma 4.4, the points $Z_{1}, \ldots, Z_{N}$ are $O(\varepsilon, O(\sqrt{\varepsilon}), O(1))$-corrupted good with respect to $\Sigma_{t}^{-1 / 2} \Sigma \Sigma_{t}^{-1 / 2}$. Let $\sigma_{t}^{2}=\left\|\Sigma_{t}^{-1 / 2} \Sigma \Sigma_{t}^{-1 / 2}\right\|$. Let $S^{\prime}$ be the output of applying naive pruning to the $Z_{i}$. Observe that for all uncorrupted $i \in[N]$, we have that $\left\|Z_{i}\right\| \leq O\left(\sigma_{t}^{2} d \log N\right)$. Thus, by the guarantee of Lemma 3.1, we have $\left\|Z_{i}\right\|_{2} \leq O\left(\sigma_{t}^{2} d \log N\right)$ for all $Z_{i} \in S^{\prime}$, and $S^{\prime}$ contains all remaining uncorrupted points in $[N]$.

Now, we have $S^{\prime}$ satisfy all the conditions of Lemma 3.3 with $\sigma^{2}=\sigma_{t}^{2}$ and $R \leq O\left(d \sigma^{2} \log N\right)$. Let $\widetilde{\Sigma}$ be the estimation of $Z_{i}$ computed by Lemma 3.3 reshaped into $d \times d$ matrix. Condition on the event that the output of APPROXIMATESCORE is a valid output of an approximate augmented score
oracle in every iteration it is called in, which by our choice of parameters, occurs with probability $1-1 / \operatorname{poly}(d, \log \kappa)$. In this event, by the guarantee on $\widetilde{\Sigma}$, we have

$$
\left\|\widetilde{\Sigma}-\Sigma_{Y}\right\|_{F}=\left\|\widetilde{\Sigma}-\Sigma_{t}^{-1 / 2} \Sigma \Sigma_{t}^{-1 / 2}\right\|_{F} \leq O(\sqrt{\varepsilon}) .
$$

This immediately implies $\widehat{\Sigma}-O(\sqrt{\varepsilon}) \Sigma_{t} \preceq \Sigma \preceq \widehat{\Sigma}+O(\sqrt{\varepsilon}) \Sigma_{t}$. Moreover, using the fact that $\|A B\|_{F} \leq\|A\|\|B\|_{F}$, we have

$$
\|\widehat{\Sigma}-\Sigma\|_{F}=\left\|\Sigma_{t}^{1 / 2} \widetilde{\Sigma} \Sigma_{t}^{-1 / 2}-\Sigma\right\|_{F} \leq O(\sqrt{\varepsilon})\left\|\Sigma_{t}\right\|,
$$

This proves the correctness guarantee. The runtime guarantee follows by combining Lemma 3.1, Lemma 3.3, and Theorem D.1.

## D. 3 Getting $O(\varepsilon \log 1 / \varepsilon)$ error

We now turn our attention to SECONDPHASE, which allows us to refine a coarse estimate down to error $O(\varepsilon \log 1 / \varepsilon)$. We restate the theorem below for convenience.
Theorem 4.7 (Second Phase). Let $S$ be a set of points satisfying Assumption 4.2. Let $0<\zeta<\zeta_{0}$ for some universal constant $\zeta_{0}$. Given $\zeta_{t}$ and $\Sigma_{t}$ where $\Sigma \preceq \Sigma_{t} \preceq\left(1+\zeta_{t}\right) \Sigma$ as input, Algorithm 3 runs in time $\widetilde{O}(T(N, d))$ and outputs a new upper bound matrix $\Sigma_{t+1}$ and a approximate covariance matrix $\widehat{\Sigma}$ such that, with probability $1-1 / N d$, for $\zeta_{t+1}=O\left(\sqrt{\varepsilon \zeta_{t}}+\varepsilon \log 1 / \varepsilon\right)$, we have

$$
\Sigma \preceq \Sigma_{t+1} \preceq \Sigma+\zeta_{t+1} \Sigma_{t}, \quad \text { and } \quad\left\|\Sigma^{-1 / 2} \widehat{\Sigma} \Sigma^{-1 / 2}-I\right\|_{F} \leq \zeta_{t+1}
$$

```
Algorithm 3 Robust Covariance Estimation With Approximately Known Covariance
    procedure \(\operatorname{SecondPhaSE}\left(S=\left\{X_{1}, X_{2}, \ldots, X_{n}\right\}, \varepsilon, \Sigma_{t}, \zeta_{t}\right) \quad \triangleright\) Theorem 4.7
        for \(i=1, \ldots, N\) do
            \(Y_{i} \leftarrow \Sigma_{t}^{-1 / 2} X_{i}\)
            \(Z_{i} \leftarrow Y_{i} \otimes Y_{i}\)
        end for
        \(S^{\prime} \leftarrow \operatorname{NaivePrune}\left(Z_{1}, \ldots, Z_{N}, O(d \log N), 1 / d N\right) \quad \triangleright\) Lemma 3.1
        Let \(\mathscr{O}_{\text {aug }}=\) APPROXIMATESCORE with \(\delta=1 / \operatorname{poly}(d, N)\)
        Let \(\widetilde{\Sigma}\) be the estimation of \(S^{\prime}\) computed by Lemma 3.4.
        \(\widehat{\Sigma} \leftarrow \Sigma_{t}^{1 / 2} \widetilde{\Sigma} \Sigma_{t}^{1 / 2}\).
        return \(\widehat{\Sigma}\)
    end procedure
```

Proof of Theorem 4.7. The proof is very similar to that of Theorem 4.6. By the definition of $Y_{i}$ and condition on $\Sigma_{t}$, we know that

$$
\left\|\Sigma_{Y}-I\right\|=\left\|\Sigma_{t}^{-1 / 2} \Sigma \Sigma_{t}^{-1 / 2}-I\right\| \leq \zeta_{t}
$$

Since $\left\|\Sigma_{Y}-I\right\| \leq \zeta$, by Lemma 4.4, the set $\left\{Z_{1}, \ldots, Z_{N}\right\}$ is $O(\varepsilon, O(\varepsilon \sqrt{\log 1 / \varepsilon}), O(\varepsilon \sqrt{\log 1 / \varepsilon}+$ $\zeta), O(\log 1 / \varepsilon), O\left(\log ^{2} 1 / \varepsilon+\zeta\right)$-corrupted good with respect to $D_{I}$.
Therefore, $S^{\prime}$ satisfies all condition of Lemma 3.4. Condition on the event that the output of APPROXIMATESCORE is a valid output of an approximate augmented score oracle in every iteration it is called. By Lemma 3.4 and our choice of parameters, this occurs with probability $1-1 / \operatorname{poly}(d, N)$. Then, Lemma 3.4 guarantees that we can find a $\widetilde{\Sigma}$ where

$$
\left\|\widetilde{\Sigma}^{b}-\mathbb{E}\left[Z_{i}\right]\right\|_{2}=\left\|\widetilde{\Sigma}-\Sigma_{Y}\right\|_{F} \leq O\left(\sqrt{\zeta_{t} \varepsilon}+\varepsilon \log 1 / \varepsilon\right)
$$

Then, we have

$$
\left\|\Sigma^{-1 / 2} \widehat{\Sigma} \Sigma^{-1 / 2}-I\right\|_{F}=\left\|\Sigma^{-1 / 2} \Sigma_{t}^{1 / 2} \widetilde{\Sigma} \Sigma_{t}^{1 / 2} \Sigma^{-1 / 2}-I\right\|_{F} \leq O\left(\sqrt{\zeta_{t} \varepsilon}+\varepsilon \log 1 / \varepsilon\right),
$$

where the last inequality follows by $\|A B\|_{F} \leq\|A\|\|B\|_{F}$ and $\left\|\Sigma-\Sigma_{t}\right\| \leq \zeta_{t}=O(1)$. This proves the correctness guarantee. Finally, the runtime guarantee follows by combining Lemma 3.1, Lemma 3.4, and Theorem D.1.

## E Fast Implementations for Tensor Inputs

In this section we describe how to implement the approximate score oracles described in Section 3.3.2 fast when the input is given as tensor products. Recall that an approximate score oracle takes as input a set of points $S$ of size $n$, and a sequence of weights $w_{0}, \ldots, w_{t-1}, w_{t}$, and computes $\tilde{\lambda}$ so that $\tilde{\lambda} \approx_{0.1}\left\|M\left(w_{0}\right)-I\right\|$, and $\tilde{\tau}_{t} \in \mathbb{R}^{n}$ where $\tilde{\tau}_{t, i} \approx_{0.1} \tau_{t, i}$ for all $i$ and $\tilde{q}_{t}$ where

$$
\begin{align*}
\tau_{t, i} & =\left(Z_{i}-\mu\left(w_{t}\right)\right)^{\top} U_{t}\left(Z_{i}-\mu\left(w_{t}\right)\right)  \tag{6}\\
\left|\tilde{q}_{t}-q_{t}\right| & \leq 0.1 q_{t}+0.05\left\|M\left(w_{t}\right)-I\right\|, \text { where } q_{t}=\left\langle M\left(w_{t}\right)-I, U_{t}\right\rangle, \tag{7}
\end{align*}
$$

and

$$
U_{t}=\exp \left(c I_{d}+\frac{1}{1.1 \alpha} \sum_{i=0}^{t-1} M\left(w_{i}\right)\right)=\frac{\exp \left(\alpha \sum_{i=0}^{t-1} M\left(w_{i}\right)\right)}{\operatorname{tr} \exp \left(\alpha \sum_{i=0}^{t-1} M\left(w_{i}\right)\right)}
$$

where $\alpha>0$ is a parameter, and $c$ is chosen so that $\operatorname{tr}\left(U_{t}\right)=1$.

## E. 1 Several Ingredients

To implement the approximate score oracle efficient, we need several ingredients. The first one is Taylor series approximation for the matrix exponential:
Lemma E. 1 ([AK16]). Let $0<\varepsilon<1$, let $X$ be a PSD matrix where $\|A\|<M$, there is a degree- $\ell$ polynomial $P_{\ell}$, where $\ell=O(\max (M, \log 1 / \varepsilon))$, such that

$$
(1-\varepsilon) \exp (X) \preceq P_{\ell}(x) \preceq(1+\varepsilon) \exp (X)
$$

Another difficulty is that writing down the matrix $U_{t}$ takes the time $\Omega\left(d^{4}\right)$. Then, we need the Johnson-Lindenstrauss Lemma [JL84] to construct a matrix in much lower dimension:
Lemma E. 2 (Johnson-Lindenstrauss Lemma (JL Lemma)). Let $J \in \mathbb{R}^{r \times d}$ be a matrix whose each entries are i.i.d. samples from $\mathcal{N}(0,1 / r)$. For every vector $v \in \mathbb{R}^{d}$ and every $\varepsilon \in(0,1)$,

$$
\operatorname{Pr}\left[\|J v\|_{2} \approx_{\varepsilon}\|v\|_{2}\right]>1-\exp \left(-\Omega\left(\varepsilon^{2} r\right)\right)
$$

Lemma $\mathbf{E} .3$ (Tellegen's Theorem, [BCS97]). Fix a matrix $A \in \mathbb{R}^{r \times c}$, if we can compute matrixvector product $A x$ for any $x \in \mathbb{R}^{c}$ in time $t$. Then, we can compute $A^{\top} y$ for any $y \in \mathbb{R}^{r}$ in time $O(t)$.
Lemma E. 4 (Power method). For any matrix $A \in \mathbb{R}^{m \times m}$, there exists an randomized algorithm, with probability $1-\delta$, outputs its $1 \pm \varepsilon$ approximation using $O(\log m \log (1 / \delta) / \varepsilon)$ many matrix-vector multiplications.

## E. 2 Fast approximate score oracle

Observe that $\mathscr{O}_{\text {aug }}$ is strictly more powerful than $\mathscr{O}_{\text {approx }}$. In this section, we describe how to implement $\mathscr{O}_{\text {aug }}$ fast.
Lemma E.5. Assuming $\alpha$ is chosen such that it always satisfies Equation (3), then $\operatorname{APProximateScore}\left(S, w_{1}, \ldots, w_{t}\right)$ runs in time $\widetilde{O}\left(t^{2} \cdot T(N, d) \log 1 / \delta\right)$.

First, we show how to compute $M\left(w_{i}\right) v$ for any $v$ by utilizing the fast rectangular matrix multiplication.
Lemma E.6. For any vector $v \in \mathbb{R}^{d^{2}}$ and $w_{i} \in \mathbb{R}^{N}$. We can compute the matrix-vector product $M\left(w_{i}\right) \cdot v$ in $O(T(N, d))$ time.

Proof. Let $Z \in \mathbb{R}^{d^{2} \times N}$ be the matrix whose $i$-th column is $Z_{i}$ and $Y \in \mathbb{R}^{d \times N}$ be the matrix whose $i$-th column is $Y_{i}$. Note that $M\left(w_{i}\right)=\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right) \operatorname{diag}\left(w_{i} /\left|w_{i}\right|\right)\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right)^{\top}$.
By Lemma E.3, $\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right)^{\top} v$ has the same running time as $\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right) v$. We have $\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right) v=Z v-\left(\mathbf{1}^{\top} v\right) \mu\left(w_{i}\right)$. We observe that $Z v=\left(Y \operatorname{diag}(v) Y^{\top}\right)^{b}$, then we can compute $Z v$ in $O(T(N, d))$ time. Then we can compute $\operatorname{diag}\left(w_{i} /\left|w_{i}\right|\right)\left(Z-\mu\left(w_{i}\right) \mathbf{1}^{\top}\right)^{\top} v$ in time $O(T(N, d))$ since multiply a diagonal matrix by a vector can be done in $O(N)$. Thus, we can compute $M\left(w_{i}\right) v$ for any $v$ in time $O(T(N, d))$.

```
Algorithm 4 Approximate Score Oracle
    procedure \(\operatorname{APPROXIMATESCORE}\left(\delta, \alpha, S=\left\{Y_{1}, \ldots, Y_{N}\right\}, w_{1}, w_{2}, \ldots, w_{t}\right)\)
        \(r \leftarrow O(\log 1 / \delta), \ell \leftarrow t\)
        Let \(J \in \mathbb{R}^{r \times d^{2}}\) matrix whose each entries are i.i.d. samples from \(\mathcal{N}(0,1 / r)\).
        Let \(Z_{i}=Y_{i} \otimes Y_{i}\) and \(\mu\left(w_{i}\right)=\frac{1}{\left|w_{i}\right|} \sum_{j=1}^{N} w_{i, j} Z_{j}\),
        Let \(M\left(w_{i}\right)=\frac{1}{\left|w_{i}\right|} \sum_{j=1}^{N} w_{i, j}\left(Z_{j}-\mu\left(w_{i}\right)\right)\left(Z_{j}-\mu\left(w_{i}\right)\right)\) and \(A=\frac{\alpha}{2} \sum_{i=0}^{t-1} M\left(w_{i}\right)\).
        Compute \(S_{r, \ell}\) where \(S_{r, \ell}=J \cdot P_{\ell}(A)\)
        Compute \(\nu=\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)\).
        Compute \(B=S_{r, \ell}\left(Z-\mu\left(w_{t}\right) \mathbf{1}^{\top}\right)\)
        for \(i=1, \ldots, N\) do
            Let \(\widetilde{\tau}_{t, i}=\frac{1}{\nu}\left\|B_{: i}\right\|_{2}^{2}\)
        end for
        Let \(\tilde{q}_{t}=\sum_{i=1}^{N}\left(\tilde{\tau}_{t, i}-1\right)\).
        Compute \(\lambda \approx_{0.1}\left\|M\left(w_{t}\right)-I\right\|\) using power method with \(r\) many iterations.
        Output \(\tilde{\tau}_{t}, \tilde{q}_{t}, \tilde{\lambda}\).
    end procedure
```

Proof of Lemma E.5. By Lemma E.6, we can compute $A v$ for any $v$ in time $O(t T(N, d))$. By repeatedly multiplying $A$ on the left, we can compute $A^{k} v$ in time $O(k \cdot t T(N, d))$. Since we take $\ell=O(t)$, we can compute $P_{\ell}(A) v$ for any $v$ in time $\widetilde{O}\left(t^{2} \cdot T(N, d)\right)$. We can compute $J \cdot P_{\ell}(A)=\left(P_{\ell}(A) J^{\top}\right)^{\top}$ by multiply each column of $J^{\top}$ to $P_{\ell}(A)$. Thus, $S_{r, \ell}$ can be computed in time $\widetilde{O}\left(r \cdot t^{2} \cdot T(N, d)\right)=\widetilde{O}\left(t^{2} \cdot T(N, d)\right)$. We can compute $Z^{\top} v$ in $O(T(N, d))$ by multiplying each row of $S_{r, \ell}$ with $Z$. Therefore, we can compute matrix the $B$ in $\widetilde{O}(T(N, d))$.
Now, we consider how to compute $\nu$. Note that $\nu=\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)=\sum_{i=1}^{r}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)_{i, i}=$ $\sum_{i=1}^{r}\left\|S_{r, \ell} e_{i}\right\|_{2}^{2}$. Thus, $\nu$ can also be computed in time $\widetilde{O}(T(N, d))$. Once we have $B$ and $\nu$, then we can compute $\tilde{\tau}$ and $\tilde{q}$ in $\widetilde{O}(N)$ time.

Using power method (Lemma E.4), we can find a $1 \pm 0.1$ approximation of $\left\|M\left(w_{t}\right)-I\right\|$ using $O(\log d \log 1 / \delta)$ matrix-vector multiplications. By Lemma E.6, we can compute $(M(w)-I)$. $v$ for any $v$ in time $O(T(N, d))$. Then, the total runtime is $O(T(N, d)) \cdot O(\log d \log (1 / \delta))=$ $\widetilde{O}(T(N, d) \log (1 / \delta))$.
Thus, the algorithm runs in time $\widetilde{O}\left(t^{2} \cdot T(N, d) \log 1 / \delta\right)$.
Lemma E.7. The output of $\operatorname{ApproximateScore}\left(S, w_{1}, \ldots, w_{t}, \alpha\right)$ satisfies $\tilde{\tau} \approx_{0.1} \tau$ and $\tilde{q} \approx_{0.1} q$ with probability $1-\delta$.

The correctness proof directly follows by the original correctness proof in [DHL19]; for completeness, we prove it below.

Proof. We condition on two events occurring. Let $\widetilde{\lambda}$ be the output of Line 13 in Algorithm 4.

- $\left\|J \cdot P_{\ell}(A) v\right\|_{2} \approx_{0.01}\left\|P_{\ell}(A) v\right\|_{2}$ for all $v \in\left\{e_{1}, \ldots, e_{d^{2}}\right\} \cup\left\{X_{1}-\mu\left(w_{t}\right), \ldots, X_{N}-\mu\left(w_{t}\right)\right\}$.
- $\tilde{\lambda} \approx_{0.1} \lambda$.

Note that by our choice of parameters, both events occur individually with probability at $1-\frac{\delta}{3}$. Then, by a union bound over failure probability, these two events succeeds with probability at $1-\delta$.

The guarantee on $\tilde{\lambda}$ directly follows from correctness of power method. Now, we show $\tilde{\tau} \approx_{0.1} \tau$. Let $M=\frac{\alpha}{2} \sum_{i=1}^{t-1} M\left(w_{i}\right)$. Then, we have

$$
\tau_{i}=\left(X_{i}-\mu\left(w_{t}\right)\right)^{\top} \frac{\exp (2 M)}{\operatorname{tr} \exp (2 M)}\left(X_{i}-\mu\left(w_{t}\right)\right)=\frac{1}{\operatorname{tr} \exp (2 M)}\left\|\exp (M)\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2}
$$

where as $\tilde{\tau}_{i}=\frac{1}{\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)}\left\|S_{r, \ell}\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2}$.
Note that

$$
\begin{align*}
\left\|S_{r, \ell}\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2}=\left\|J \cdot P_{\ell}(A)\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2} & \approx_{0.01}\left\|P_{\ell}(A)\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2} \\
& \approx_{0.03}\left\|\exp (M)\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2} \tag{8}
\end{align*}
$$

where the first line follows by Lemma E. 2 and the last line follows by our choice of $\ell$ and Lemma E.1.
Similarly, we have $\exp (2 M)_{i, i}=\left\|\exp (M) e_{i}\right\|_{2}^{2}$ and $\left(S_{r, \ell} S_{r, \ell}^{\top}\right)_{i, i}=\left\|S_{r, \ell} e_{i}\right\|_{2}^{2}$.
By definition of $S_{r, \ell}$, we have

$$
\left\|S_{r, \ell} e_{i}\right\|_{2}^{2}=\left\|J \cdot P_{\ell}(A) \cdot e_{i}\right\|_{2}^{2} \approx_{0.01}\left\|P_{\ell}(A) \cdot e_{i}\right\|_{2}^{2} \approx_{0.03}\left\|\exp (M) e_{i}\right\|_{2}^{2}
$$

and this immediately implies $\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right) \approx_{0.03} \operatorname{tr}(\exp 2 M)$. Thus, we have $\tilde{\tau}_{i} \approx_{0.07} \tau_{i}$.
Now, we show $\tilde{q}$ is close to $q$. Rewriting $\tilde{q}$, we get

$$
\begin{aligned}
\tilde{q} & =\frac{1}{\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)} \sum_{i=1}^{N}\left(\left\|S_{r, \ell}\left(X_{i}-\mu\left(w_{t}\right)\right)\right\|_{2}^{2}-\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)\right) \\
& =\frac{1}{\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)}\left\langle M\left(w_{t}\right)-I, P_{\ell}(M) J J^{\top} P_{\ell}(M)\right\rangle \\
& =\frac{1}{\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right)}\left(\left\langle M\left(w_{t}\right)-I, \exp (2 M)\right\rangle+\xi\right),
\end{aligned}
$$

where $|\xi| \leq 0.02\left\|M\left(w_{t}\right)-I\right\| \cdot \operatorname{tr} \exp (2 M)$. We complete the proof by note that $\operatorname{tr}\left(S_{r, \ell} S_{r, \ell}^{\top}\right) \approx_{0.03}$ $\operatorname{tr}(\exp 2 M)$.

