Re-Analyze Gauss: Bounds for Private Matrix Approximation via Dyson Brownian Motion

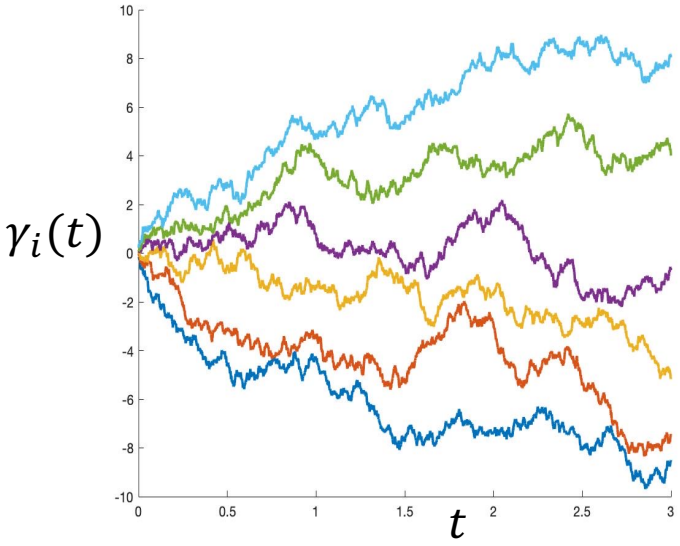

Oren Mangoubi
WPI

Nisheeth Vishnoi
Yale

NeurIPS 2022

Matrix Approximation

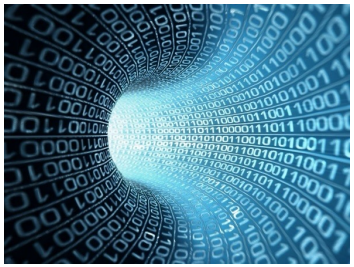

Matrix Approximation Problem:

Given $\lambda_1 \geq \cdots \geq \lambda_d$, $\Lambda := \mathrm{diag}(\lambda_1, \cdots, \lambda_d)$,

Symmetric $M \in R^{d \times d}$ with eigenvalues $\sigma_1 \geq \cdots \geq \sigma_d$ and diagonalization $M = V \Sigma V^\top$.

$$\min_{\hat{V} \in \mathrm{O}(d)} ||V \Lambda V^\top - \hat{V} \Lambda \hat{V}^\top||_F$$

- Without differential privacy constraints, optimal solution is $V = \hat{V}$ [Schur, 1923; Horn, 1954]
- With differential privacy constraints, need to output "noisy" $\hat{V}$ to hide private information

---

Special cases:

- Rank-k covariance approximation: $\lambda_i = \sigma_i, \quad i \leq k$
- Subspace Recovery: $\lambda_1 = \cdots = \lambda_k = 1, \quad \lambda_{k+1} = \cdots = \lambda_d = 0$

- Many applications to ML, statistics, medicine, engineering, etc.

Differential privacy

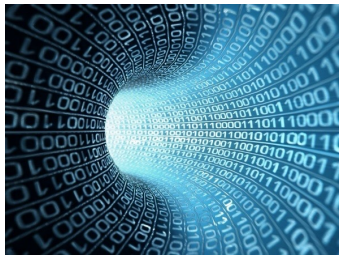

Applications:
- Medical data
- Census data
- etc.

- Covariance matrix may contain sensitive information about individuals
- Many examples of privacy breaches (e.g. Netflix problem)
- Need algorithms which output approximations H that hide private information, while still allowing researchers to learn from the data

---

(ε, δ)-differential privacy (DP) [Dwork, '06]:

Given $\varepsilon, \delta > 0$, a randomized mechanism $\mathcal{A}$ is (ε, δ)-DP if for any neighboring M, M', $\mathbb{P}[\mathcal{A}(M) \in S] \leq \exp(\varepsilon) \times \mathbb{P}[\mathcal{A}(M') \in S] + \delta$

(M, M' are "neighbors" if they differ by one datapoint: $M' - M = x\, x^*$, some $||x|| \leq 1$)

Previous work: Utility Upper bounds

Many works give algorithms for private matrix approximation under both (ε, δ)-DP, e.g., [Dwork, McSherry, Nissim, Smith, '06], [Hardt, Roth, '12, '13], and pure $(\varepsilon, 0)$-DP, e.g., [Kapralov, Talwar, '13], [Amin, Dick, Kulesza, Munoz, Vassilvitskii '19], [Leake, McSwiggen, Vishnoi, '21], [Mangoubi, Wu, Kale, Thakurta, Vishnoi, '22]

Best current utility upper bounds (w.r.t. Frobenius norm) under (ε, δ)-DP [Dwork, Talwar, Thakurta, Zhang '14], based on Gaussian mechanism of [Dwork, Kenthapadi, McSherry, Mironov, Naor, '06]:

- Rank-k covariance approximation: $||M_k - \widehat{M}_k||_F \leq \tilde{O}(k\sqrt{d})$ w.h.p.
- Subspace approximation: $\left|\left|V_k V_k^\top - \hat{V}_k \hat{V}_k^\top\right|\right|_F \leq \tilde{O}\left(\frac{\sqrt{kd}}{\sigma_k - \sigma_{k+1}}\right)$ w.h.p.

- $\tilde{O}(k\sqrt{d})$ utility bounds (for covariance approximation) not tight for all k, σ: e.g., for $k = d$, Gaussian mechanism gives $\left|\left|M - \widehat{M}\right|\right|_F \leq \tilde{O}(d) = \tilde{O}(\sqrt{k}\sqrt{d})$.
- Roughly, linear-in-k dependence is because perturbations to top-k eigenvectors are "added up" as a simple sum using trace inequalities

Can one obtain $\sqrt{k}$ improvement on utility for $k < d$, by adding up perturbations as (the square root of) a **sum-of-squares**?

Algorithm: Gaussian Mechanism for Matrix Approximation

Input: $M \in R^{d\times d}$, diagonal matrix $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_d)$

1. Add noise $\widehat{M} = G + G^{\mathsf{T}}$, where G has iid $\mathrm{N}(0,\ \varepsilon^{-1}\log d)$ entries.
2. Post-process: Diagonalize $\widehat{M} = \hat{V}\,\hat{\Sigma}\,\hat{V}^{\mathsf{T}}$
3. Output: $H = \hat{V}\Lambda\hat{V}^{\mathsf{T}}$

Main results: Utility bounds

Assumption(M): The top-k eigenvalues of M satisfy $\sigma_i - \sigma_{i+1} \geq \widetilde{\Omega}(\sqrt{d}) \quad \forall\, i \leq k$

Theorem: For $\varepsilon, \delta > 0, k \in [d]$, and given $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_d)$and a symmetric $M \in R^{d\times d}$ with eigenvalues $\sigma_1 \geq \cdots \geq \sigma_d$ and diagonalization $M = V\Sigma V^{\mathsf{T}}$. The above Gaussian Mechanism is (ε, δ)-DP and outputs $H = \hat{V}\Lambda\hat{V}^{\mathsf{T}}$ s.t., if Assumption(M) holds,

$$\mathbb{E}\left[\|\hat{V}\Lambda\hat{V}^{\mathsf{T}} - V\Lambda V^{\mathsf{T}}\|_F^2\right] \leq \tilde{O}\left(\sum_{i=1}^{k}\sum_{j=i+1}^{d}\frac{(\lambda_i - \lambda_j)^2}{(\sigma_i - \max(\sigma_j, \sigma_{k+1}))^2}\right)$$

Corollary (covariance approximation): $\mathbf{E}[||V\Sigma_k V^{\mathsf{T}} - \hat{V}\,\hat{\Sigma}_k\hat{V}^{\mathsf{T}}||_F] \leq \tilde{O}(\sqrt{k}\,\sqrt{d}\,\frac{\sigma_k}{\sigma_k - \sigma_{k+1}})$

- Improves by $\sqrt{k}$ on [Dwork et al'14], if $\sigma_i - \sigma_{i+1} \geq \widetilde{\Omega}(\sqrt{d})\ \forall\, i \leq k$ and $\sigma_k - \sigma_{k+1} \geq \Omega(\sigma_k)$

Corollary (subspace recovery): $\mathbf{E}\,[||V_k V_k^{\mathsf{T}} - \hat{V}_k\,\hat{V}_k{}^{\mathsf{T}}||_F] \leq \tilde{O}(\sqrt{d}\,\frac{\sigma_k}{\sigma_k - \sigma_{k+1}})$

- Improves by $\sqrt{k}$ on [Dwork et al '14], if we also have $\sigma_i - \sigma_{i+1} \geq \sigma_k - \sigma_{k+1} \forall\, i \leq k$

Dyson Brownian Motion

- View addition of Gaussian noise as a matrix-valued stochastic process

$$\widehat{M}(t) = M + B(t), \qquad t > 0$$

  - $B(t) = W(t) + W^{\top}(t)$
  - Each entry of $W(t)$ is a standard Brownian Motion

- Diagonalize $\widehat{M}(t) = U(t)\Gamma(t)U^{\top}(t)$
- Eigenvalues and eigenvectors evolve according to Stochastic Differential Equations (SDE) discovered by [Dyson, '62]:

$$\mathrm{d}\gamma_i(t) = \mathrm{d}B_{ii} + \sum_{i=1}^{d} \frac{1}{\gamma_i(t) - \gamma_j(t)}\,\mathrm{d}t$$

$$\mathrm{d}u_i(t) = \sum_{j \neq i} \frac{\mathrm{dB}_{ij}}{\gamma_i(t) - \gamma_j(t)}\,\mathrm{u}_j(t) - \frac{1}{2}\sum_{j \neq i} \frac{\mathrm{d}t}{(\gamma_i(t) - \gamma_j(t))^2}\,\mathrm{u}_i(t)$$

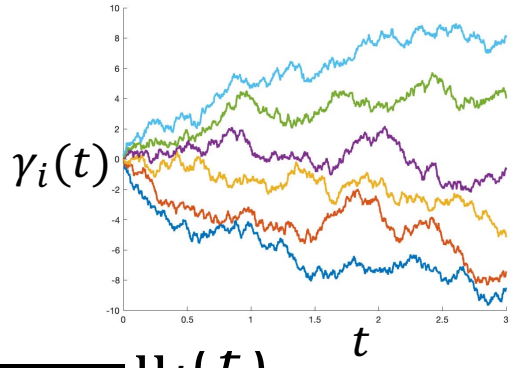

Using Dyson's SDEs to bound the utility

- Define projected process: $\Psi(t) = U(t)\Lambda U^{\top}(t)$
- We want to bound $||\Psi(T) - \Psi(0)||_F^2 = ||\int_0^T \mathrm{d}\Psi(t)||_F^2$

Step 1: use Dyson's equations to derive SDE for $\Psi(t)$:

$$\mathrm{d}\Psi(t) = \frac{1}{2}\sum_{i=1}^{d}\sum_{j\neq i}(\lambda_i - \lambda_j)\underbrace{\frac{\mathrm{d}B_{ij}(t)}{\gamma_i(t) - \gamma_j(t)}(u_i(t)u_j^{\top}(t) + u_j(t)u_i^{\top}(t))}_{\mathrm{d}\alpha_{ij}(t)} + \sum_{i=1}^{d}\sum_{j\neq i}(\lambda_i - \lambda_j)\frac{\mathrm{d}t}{(\gamma_i(t) - \gamma_j(t))^2}u_i(t)u_i^{\top}(t)$$

- $\mathrm{d}\Psi(t)$ is a sum of independent random terms $\mathrm{d}\alpha_{ij}(t)$:
  - $\mathrm{d}\alpha_{ij}(t)$ independent for all i, j, and independent of *past* times t
  - Their time-integrals $\int_0^T \mathrm{d}\alpha_{ij}(t)$ are <u>*not*</u> independent for all i, j!

Step 2: Use independence to "add up" Frobenius norms of $\mathrm{d}\alpha_{ij}(t)$ as sum-of-squares. Then use Ito's Lemma to integrate Frobenius norm over time:

$$\mathbb{E}[\|\Psi(T) - \Psi(0)\|_F^2] = 2\int_0^T \mathbb{E}\left[\sum_{i=1}^{d}\sum_{j\neq i}\frac{(\lambda_i - \lambda_j)^2}{(\gamma_i(t) - \gamma_j(t))^2}\mathrm{d}t\right] + T\int_0^T \mathbb{E}\left[\sum_{i=1}^{d}\left(\sum_{j\neq i}\frac{\lambda_i - \lambda_j}{(\gamma_i(t) - \gamma_j(t))^2}\right)^2\right]\mathrm{d}t.$$

Step 3: Use Weyl's inequality to bound eigenvalue gaps:

$$\gamma_i(t) - \gamma_j(t) \geq \sigma_i - \sigma_j - \big|\big|B(t)\big|\big|_2 \geq \frac{1}{2}(\sigma_i - \sigma_j) \quad \text{w.h.p.,}$$

as long as $\sigma_i - \sigma_{i+1} \geq \sqrt{d} \quad \forall\, i \leq k.$

Conclusion

Introduced new method of analyzing addition of noise by Gaussian mechanism as a matrix-valued Brownian motion

- In special case of rank-k covariance approximation ($\lambda_i = \sigma_i$, $i \leq k$) under (ε, δ)-DP, implies $\sqrt{k}$ **improvement** on utility bound $\mathrm{E}\left[\left|\left|M_k - \widehat{M}_k\right|\right|_F\right]$ if eigenvalues of M satisfy $\sigma_i - \sigma_{i+1} \geq \widetilde{\Omega}\left(\sqrt{d}\right)$, $i \leq k$ and $\sigma_{\boldsymbol{k}} - \sigma_{k+1} \geq \Omega(\sigma_k)$

- In special case of subspace recovery ($\lambda_i = 1$, $i \leq k$) under (ε, δ)-DP, implies $\sqrt{k}$ **improvement** on utility bound $\mathrm{E}\left[\left|\left|V_k V_k^\top - \hat{V}_k \hat{V}_k^\top\right|\right|_F\right]$ if eigenvalues of M also satisfy $\sigma_i - \sigma_{i+1} \geq \sigma_k - \sigma_{k+1}, \quad i \leq k$

Open problem: Can one obtain similar utility bounds *without* assumption that $\sigma_i - \sigma_{i+1} \geq \sqrt{d} \quad \forall\, i < k$?

Thanks!