Privacy-Preserving Visualization of Brain Functional Connectivity

Ye Tao[1], Anand D. Sarwate[1], Sandeep Panta[2], Sergey Plis[2], Vince D. Calhoun[2]

***Abstract*—Privacy protection is important in visualization due to the risk of leaking personal sensitive information. In this paper, we study the problem of privacy-preserving visualizations using differential privacy, employing biomedical data from neuroimaging as a use case. We investigate several approaches based on perturbing correlation values and characterize their privacy cost and the impact of pre- and post-processing. To obtain a better privacy/visual utility tradeoff, we propose workflows for connectogram and seed-based connectivity visualizations, respectively. These workflows successfully generate visualizations similar to their non-private counterparts. Experiments show that qualitative assessments can be preserved while guaranteeing privacy. These results show that differential privacy is a promising method for protecting sensitive information in data visualization.**

***Index Terms*—Differential privacy, visualization, neuroimaging data**

I. Introduction

DATA visualization [2], [3] is an important part of the research workflow that can uncover hidden patterns. Visualizations also assist in presenting research findings, making it easier for individuals to comprehend and analyze the results. In many applications, the data being used are private or sensitive. This is especially true in areas related to human health and behavior, where releasing "raw" data or even summary tables can breach the commitments made by data holders to individuals [4], [5]. Privacy harms can come from many different computations on the data [6], including visualization. For example, a scatter plot can reveal the actual data values. Visualizations can play several roles within a data analysis workflow. In many cases, visualizations are used to provide qualitative support for quantitative findings. This is particularly true in exploratory data analysis. Methods such as differential privacy offer protection at the cost of accuracy, or utility. This raises interesting challenges: how should we balance privacy and visual utility?

In this paper we explore this question via case studies of common visualizations used in neuroimaging research. Specifically, we explore the application of differential privacy [7] to the creation of functional network connectivity (FNC), connectogram, and seed-based connectivity (SBC) visualizations [8]–[10]. In neuroimaging research, privacy concerns represent a significant barrier to research collaborations, often resulting in small sample sizes and research "silos" that prevent researchers from leveraging larger sample sizes by performing joint analyses. Privacy-preserving data visualization can help demonstrate the benefits of collaboration, making researchers more willing (or able) to safely share access to sensitive data [11]–[13]. The visualization approach we propose here will be integrated into a federated learning system to allow for more effective and automated research collaboration in neuroimaging [14].

Differential privacy (DP) [7], [15], [16] has become one of the standard frameworks for quantifying the risk of sharing analyses performed on sensitive data. Most current applications focus on protecting privacy when publishing numerical outputs such as synthetic data, statistical estimates, or machine learning models. Under DP, algorithms are randomized and this randomization (often in the form of noise) affects the accuracy of the estimates, leading to a *privacy-utility tradeoff*. However, when thinking about visualization, the notion of "utility" can change: in some applications, a private visualization has high utility if an expert would draw the same conclusions as from the non-private (unprotected) visualization. In image processing, it is well-known that numerical measures such as mean squared error (MSE) may not be good measures of visual quality: the same phenomenon occurs in private visualization.

There are (at least) three potential issues to consider when using differential privacy for visualization: the mismatch between the specific computations needed in the visualization and existing (generic) DP methods, the issues of multiple data accesses coming from applications in exploratory data analysis, and the aforementioned challenge of defining the utility of data visualizations [3], [17], [18]. For our example of FNC, we adapt DP mechanisms to give a private estimate of the correlation matrix underlying the visualization. In the case of the connectogram, we introduce a DP workflow designed for conducting group comparisons. For SBC visualization, we propose a DP procedure to systematically identify and present the statistically significant connectivities. We use (adaptive) composition techniques [19]–[22] to quantify the risk of multi-

[1] Department of Electrical and Computer Engineering at Rutgers, The State University of New Jersey, Piscataway, NJ 08854. (yt371,ads221)@rutgers.edu.

[2] TReNDS Center, Georgia State, Georgia Tech, and Emory, Atlanta, GA 30303. (spanta,splis,vcalhoun)@gsu.edu.

Initial results from this paper were presented in part at the 21st IEEE International Symposium on Biomedical Imaging (ISBI 2024) [1].

The work of the authors was supported in part by the US National Institutes of Health under awards 2R01DA040487, R01DA049238, and R01MH121246, as well as by the National Science Foundation under award 2112455. The work of Sarwate was also supported by the National Science Foundation under award CNS-2148104 and is supported in part by funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program.

The authors wish to thank the FBIRN, BLSA, and ADNI teams. In addition, the authors acknowledge the Neuro Bureau, the ADHD 200 consortium, and Virginia Tech's ARC.

ple releases. Finally, to address the issues of utility, we provide both qualitative assessments and quantitative measurements to characterize the benefit of different approaches.

By delving into these issues, we outline a process for thinking though privacy-preserving visualization in other applications and highlight considerations for ensuring both high-quality visualization and data privacy. This process can be extended beyond neuroimaging to various domains. Typically, visualization involves extracting statistic or aggregating data from raw dataset but there are often several steps in a pipeline of data processing. The first step is to identify the statistics or aggregated data required for visualization. Next, we consider pre-processing methods to reduce the required noise. Different types of visualizations may require different pre-processing methods. Then, we design the differentially private method, which may includes selecting distributions, using composition methods, or constructing workflows, etc. Finally, we employ post-processing methods to mitigate the impact of perturbation. By following this process, we demonstrate how we create privacy-preserving visualizations in neuroimaging (see Section III and IV).

Prior Work. Differentially private visualization is a relatively new field, but prior works have explored several interesting scenarios. Avraam et al. [23] proposed the use of anonymization techniques, such as probabilistic anonymization, to create privacy-preserving visualizations. Probabilistic anonymization involves adding random noise to individual values. While similar to DP, this work does not quantitatively measure privacy protection. Nanayakkara et al. [24] and Thaker et al. [25] concentrated on interactive interfaces that visualize the relationship between privacy budget and accuracy to address *data privacy management*, as opposed to data visualization. Zhang et al. [26] apply the (DP) Laplace mechanism to four visualization types (e.g. bar charts) and explore methods to tune noise injection to improve visualization utility. Lee [27] applied the Laplace mechanism on timeseries data. They examine how five types of visualization products (bar graphs, pie charts, heatmaps, linear plots, and scatter plots) present information from statistical database queries. Panavas et al. [28] conducted a comparative study of visual utility among algorithms (e.g., AHP, DAWA) for generating differentially private scatter plots. Zhang et al. [17] proposed visualizing two-dimensional location data by constructing grids across the 2D domain and applying the Laplace mechanism to calculate noisy counts within each grid cell. Liu et al. [29] examined the privacy of aggregating eye tracking data as heatmaps and proposed the Laplace mechanism and the Gaussian mechanism as solutions to improve the privacy of aggregated gaze data. Ghazi et al. [30] worked on producing private heatmaps from users' aggregated data, with the primary algorithm being a differentially private procedure that produces an output close to the average of the input distributions in terms of earth mover's distance. Zhou et al. [18] introduced the PriVis model, a Bayesian network-based differential privacy model designed to produce private visualization charts (i.e., scatter, line, and bar charts) while preserving user-preferred patterns. They develop the DPVisCreator system to help data custodians implement their approach. The above works apply privacy protection for different visualization applications. By contrast to these works focusing on "generic" (application-independent) visualization, we focus on an important use case in neuroimaging: functional connectivity and perform a systematic evaluation of DP mechanisms, composition, and pre-/post-processing methods.

Our contributions. Our main contributions are: (1) We introduce several novel methods for DP visualization using neuroimaging data as a specific use case. (2) We compare their performance across different real-world datasets. (3) We show that different mechanisms have different optimal composition strategies depending on the privacy parameters used. (4) We propose different pre- and post-processing approaches to improve visualization quality. (5) We propose an end-to-end differentially private workflow for connectogram visualization, with a specific focus on group comparison, and a workflow for SBC visualization to identify the significant connectivities.

II. Background

A. Differential Privacy

Privacy research has become increasingly important in the last two decades, driven by advances in machine learning and the large-scale warehousing of data. Earlier syntactic approaches to privacy such as de-identification and k-Anonymity [31], [32] are susceptible to certain attacks such as linkage and membership inference [6], [33]–[35]. Differential privacy [16] has emerged as the state-of-the-art method by providing a quantitative definition of what privacy is and how it can be protected.

While differential privacy has been widely applied in many fields, its application in the field of visualization is still in the early stage. Visualization plays a crucial role in expressing discoveries and validating experimental results. However, if privacy considerations are not taken into account, it has the potential to inadvertently reveal an individual's sensitive information. For example, a scatter plot or a boxplot can reveal the actual data values. Another example involves Strava, a popular app used for tracking activity and exercise. Strava has previously released an anonymized heatmap of its global data. Interestingly, this heatmap revealed a scattering of little-known locations in war zones, which happen to be secret military bases of the United States [36]. Moreover, unprotected visualizations may be vulnerable to privacy attacks such as linkage attack. If an attacker knows that the trajectory depicted in the map from Fig. 1 belongs to Alice, they can deduce the latitude and longitude coordinates using landmarks. Additionally, if the attacker has access to the information from the table in Fig. 1, it becomes easy to infer that Alice's income is $85,000. As with more widely-studied queries, even if a data visualization only presents aggregated data there exists a risk of privacy leakage. Numerous studies suggest that, even with aggregated data, attackers can employ simple methods to extract the information they seek [37]–[39].

Differential privacy has been developed to enable statistical analyses on datasets while preserving the privacy of individuals whose records may be included. A differentially private algorithm/mechanism is a randomized function: the

Pseudo-identifier	Spatiotemporal points				Attributes		
3ED806C98F0E01	40.768477, -73.989071 2020/01/22 10:21	40.769886, -73.984356 2020/01/22 10:29	40.780404, -73.973086 2020/01/22 10:37	40.780190, -73.959346 2020/01/22 10:42	Female	Lawyer	$85,000
C0AD160C175446	40.762220, -73.972374 2020/01/22 07:11	40.761138, -73.983350 2020/01/22 07:29	40.756846, -73.982681 2020/01/22 07:38	40.751709, -73.990266 2020/01/22 07:53	Male	Accountant	$76,000
81FA0823E89531	40.805144, -73.956779 2020/01/22 11:54	40.809218, -73.948930 2020/01/22 11:59	40.812825, -73.945722 2020/01/22 12:02	40.815094, -73.943202 2020/01/22 12:06	Male	Engineer	$74,000
72BDAB014CAB1F	40.778801, -73.956132 2020/01/22 22:14	40.768538, -73.951561 2020/01/22 22:21	40.754988, -73.946112 2020/01/22 22:41	40.744272, -73.951005 2020/01/22 22:51	Female	Manager	$104,000

Fig. 1. Idea of a linkage attack using fake data in the context of privacy breaches.

randomization is designed to make it difficult for an adversary observing the output to discern private information. The kind of information which DP protects is captured by the notion of neighboring databases. For two datasets D and D' that are neighbors, denoted as $D \sim D'$, the randomization should make it difficult for the adversary to distinguish, using a hypothesis test, whether the output was generated from D or D' [19], [40]. Most commonly, we say $D \sim D'$ if they differ by the data of a single individual: this means that an adversary cannot reliably tell whether that individual was in the dataset or not.

For an algorithm to be differentially private, it must guarantee that the adversary's test is unreliable for all possible $D \sim D'$. The measure of unreliability is measured in DP by using two parameters, ϵ and δ.

Definition II.1 ((ϵ, δ)-DP [16])**.** *Let $\epsilon, \delta \geq 0$, a randomized algorithm $M : \mathcal{D} \to \mathcal{H}$ is (ϵ, δ)-differentially private if for all $\mathcal{Y} \subseteq \mathcal{H}$ and for all neighboring datasets $D, D' \in \mathcal{D}$,*

$$\mathbb{P}[M(D) \in \mathcal{Y}] \leq e^{\epsilon}\mathbb{P}[M(D') \in \mathcal{Y}] + \delta.$$

The parameter ϵ is the privacy budget that quantifies the maximum distance between the output distributions of algorithm M when it runs on any two neighboring datasets. A smaller ϵ will give stronger privacy but less accurate responses. The parameter δ represents the probability of information leakage. Typically, we are interested in values of δ that are far less than the inverse of the database size, $\delta \ll 1/N$. The definition of DP ensures that for all neighboring datasets D and D', the absolute value of the privacy loss [16] is bounded by ϵ with a probability of at least $1 - \delta$.

To design a (ϵ, δ)-DP algorithm that approximates a desired function $f : \mathcal{D} \to \mathbb{R}^d$, a common approach is to add random noise of appropriate magnitude to the output of $f(D)$. The amount of noise required depends not only on ϵ and δ, but also on the sensitivity of the function. The global ℓ_1 and ℓ_2 sensitivity of the function f are

$$\Delta_i = \max_{D \sim D'} \|f(D) - f(D')\|_i \tag{1}$$

for $i = 1$ or 2. "Pure" $(\epsilon, 0)$-DP can be achieved by applying the Laplace mechanism $M_L(D) = f(D) + Z$, where $Z \sim \text{Laplace}(0, \Delta_1/\epsilon \cdot I)$, although other noise distributions have also been proposed [41], [42]. For any $\epsilon, \delta \in (0, 1)$, the classical Gaussian mechanism $M_G(D)$ guarantees "approximate" (ϵ, δ)-DP when $M_G(D) = f(D) + Z$, where $Z \sim \mathcal{N}(0, \sigma^2 \cdot I)$ and $\sigma = \Delta_2\sqrt{2\log(1.25/\delta)}/\epsilon$. Balle et al. [43] proposed an "analytic" formula for the variance because this value of σ can be quite loose in the high privacy regime ($\epsilon \to 0$) and does not apply in the low privacy regime ($\epsilon \to \infty$).

Additive noise mechanisms work for numerical outputs, while the Exponential mechanism [44] is used for $(\epsilon, 0)$-DP randomized selection. Given a range $\mathcal{H}$, the Exponential mechanism is defined with respect to a utility function $u : \mathcal{D} \times \mathcal{H} \to \mathbb{R}$, which maps input and output pairs to utility scores. The Exponential mechanism $M_E(D)$ selects and outputs an element $H \in \mathcal{H}$ with a probability proportional to $\exp\left(\frac{\epsilon u(D,H)}{2\Delta_u}\right)$, where $\Delta_u = \max_{H \in \mathcal{H}} \max_{D \sim D'} |u(D, H) - u(D', H)|$.

Composition is one of the most crucial properties in DP. It provides a quantitative way to compute the accumulated privacy leakage over multiple analyses on the same dataset. There are several kinds of composition rules: basic, advanced, exact,

and Rényi differential privacy (RDP)-based compositions [16], [19], [22], [42], [45]–[56].

B. Functional Connectivity

In our work, we focus on neuroimaging data, specifically functional magnetic resonance imaging (fMRI), for several reasons: (1) Neuroimaging data is highly sensitive, as it can potentially reveal details about an individual's cognitive states. Protecting the privacy of participants in neuroimaging studies is crucial. (2) Neuroimaging data captures rich and complex brain activity, making it ideal for testing differential privacy techniques across diverse patterns. (3) Neuroimaging often involves sharing data across institutions. Implementing differential privacy enables collaborative efforts without concerns about unintentional data disclosures. (4) Neuroimaging data involves a wide variety of visualization types, facilitating our exploration and discussion of various visualization approaches. We believe that the exploration and discoveries we employ in the field of neuroimaging can also be applied or transferred to other domains.

Resting-state functional MRI (rs-fMRI) is a powerful neuroimaging technique that captures the blood oxygenation level dependent (BOLD) signal while subjects are in the scanner without performing specific tasks. It involves acquiring a series of MRI images over time, with each image containing information about the BOLD signal intensity at each voxel in the brain. We use $S(v,t)$ to denote the BOLD timeseries for voxel v at time t. MRI images typically consist of the order of $100,000$ voxels, and various parcellation schemes [57]–[59] are used to group these voxels into regions of interest (ROIs). Once ROIs are defined, ROI timeseries can be generated by averaging the timeseries across voxels within each ROI. The BOLD timeseries for an ROI is given by $S(r,t) = \frac{1}{M}\sum_{i=1}^{M} S(v_i,t)$, where r represents the ROI and M is the number of voxels in that ROI. We may abuse notation, but it will be clear from the context. Voxel-based or ROI-based timeseries analysis provides insight into brain functional connectivity at different levels, and various visualizations are necessary to effectively present different results.

Functional Network Connectivity (FNC) measures correlations between activity in different brain regions [9]. FNC is determined by calculating the Pearson correlation coefficient, which ranges from -1 to 1, between different ROI-based timeseries. Mathematically, the correlation between any two ROIs is given by

$$X_{ij} = \frac{\sum_{t=1}^{T} S(r_i,t)S(r_j,t)}{\sqrt{\sum_{t=1}^{T} S^2(r_i,t)}\sqrt{\sum_{t=1}^{T} S^2(r_j,t)}},$$

where T is the number of time points, and we assume that $S(r,t)$ is centered to 0. We represent the FNC between r brain regions for a subject n by a $r \times r$ adjacency matrix $X^{(n)}$ of correlation. We wish to visualize the average FNC $\bar{X} = \frac{1}{N}\sum_{n=1}^{N} X^{(n)}$ for a group of N subjects. The visualization normalizes the elements of $\bar{X}$ to the range $[0,1]$ and quantizes each entry $\bar{X}_{ij}$ to one of the color values in set $\mathcal{C}$ which is then mapped to a color triple (R_{ij}, G_{ij}, B_{ij}).

A differentially private visualization of the FNC matrix is a randomized function that takes $\bar{X}$ and produces an $r \times r$ RGB image. Since $\bar{X}$ is symmetric this involves computing $\frac{r(r-1)}{2}$ RGB triples. Since the entries of $\bar{X}$ are scalars, most visualization packages use a fixed color map from scalars to colors. We found that in this situation, randomizing the color map did yielded significantly worse results (see Remark 1). This is because the discrete output space $\mathcal{H}$ of all RGB images has $|\mathcal{H}| = |\mathcal{C}|^{r(r-1)/2}$ where $|\mathcal{C}|$ is the number of colors. The cardinality of $\mathcal{H}$ was too large for these algorithms to be effective.

In addition to visualizing the average FNC matrix, the connectogram is also commonly used to represent FNC analysis results. A connectogram is a graphical representation where nodes represent brain regions, and edges depict the strength or existence of connections between these regions.

While FNC concentrates on large-scale network interactions utilizing ROI-based timeseries, seed-based connectivity (SBC) offers detailed information about the connectivity of a particular seed across the rest of the brain. SBC involves selecting a specific seed voxel or ROI and analyzing its correlation with the timeseries of all other brain voxels. SBC spatial maps are commonly used to illustrate the resulting connectivity patterns [10].

III. FNC Visualization

A. FNC Matrix Visualization

In this section, we describe the process of visualizing the average FNC matrix as summarized in Fig. 2 along with pre- and post-processing approaches to help improve utility.

Matrix-level Private Algorithms. In FNC matrix visualization, the true outputs are the RGB triples for each pairwise correlation in the average matrix. To guarantee DP, we investigate the following strategy: computing the average FNC with DP and then using the standard quantization and color map. This strategy corresponds to *output perturbation* in which we compute the average FNC matrix and add noise:

$$\tilde{X} = \bar{X} + Z,$$

where the distribution of the noise Z depends on the privacy parameters. The numerical outputs $\tilde{X}$ need to be truncated to the valid range of $[-1,1]$ after perturbation. We propose two mechanisms:

- Pure: $(\epsilon,0)$-DP Laplace mechanism. $Z \sim \text{Laplace}(0, \frac{\Delta_1}{\epsilon} \cdot I)$, where $\Delta_1 = r(r-1)/N$ is the ℓ_1 sensitivity of $\bar{X}$.
- Approx: (ϵ,δ)-DP analytic Gaussian mechanism. Let $Z \sim \mathcal{N}(0, \sigma^2 \cdot I)$, where σ is the minimum standard deviation that satisfies the inequality in Theorem 8 from [43] using the ℓ_2 sensitivity $\Delta_2 = \sqrt{2r(r-1)}/N$ of $\bar{X}$.

Entry-level Private Algorithms. For the matrix-level methods, we calculate the overall sensitivity of the FNC matrix. Subsequently, we generate the corresponding perturbed matrix. We may find more effective approaches by applying composition methods. Given the total privacy parameters (ϵ,δ) and the size of sequence $k = r(r-1)/2$, we propose:

- Pure_k: We first compute the privacy parameters $(\epsilon',0)$ per entry using exact composition, and then apply the

Algorithm	Name	Strategy	Mechanism	Composition method
matrix-level	Pure	add perturbation	Laplace	—
	Approx	add perturbation	Gaussian	—
	Exp	random mapping	Exponential	—
entry-level	$\mathsf{Pure_k}$	add perturbation	Laplace	exact
	$\mathsf{Pure_r}$	add perturbation	Laplace	RDP-based
	$\mathsf{Approx_k}$	add perturbation	Gaussian	exact
	$\mathsf{Approx_r}$	add perturbation	Gaussian	RDP-based
	$\mathsf{Exp_k}$	random mapping	Exponential	exact

Fig. 2. A summary of the mechanisms explored in this paper. For certain tasks other noise distributions may have superior performance, but Laplace and Gaussian noise are most widely studied so we adopt them here for ease of exposition.

Laplace mechanism on each entry with $\Delta = 2/N$. Thus, $\tilde{X}_{ij} = \bar{X}_{ij} + Z_{ij}$ where $Z_{ij} \sim \text{Laplace}(0, \Delta/\epsilon')$.

- $\mathsf{Pure_r}$: Based on the overall privacy guarantee, we are able to compute scale parameter b of Laplace distribution using RDP-based composition. Then $\tilde{X}_{ij} = \bar{X}_{ij} + Z_{ij}$, where $Z_{ij} \sim \text{Laplace}(0, b)$.
- $\mathsf{Approx_k}$: We compute the privacy parameters (ϵ', δ') per entry using exact composition and apply the analytic Gaussian mechanism on the entries, each of which has sensitivity $\Delta = 2/N$. Thus, $\tilde{X}_{ij} = \bar{X}_{ij} + Z_{ij}$ where the variance of Z_{ij} satisfies the inequality in Theorem 8 from [43].
- $\mathsf{Approx_r}$: We compute the variance σ^2 of Gaussian distribution using RDP-based composition for the given privacy parameters. Then $\tilde{X}_{ij} = \bar{X}_{ij} + Z_{ij}$, where $Z_{ij} \sim \mathcal{N}(0, \sigma^2)$.

Visualization is often a preparatory step for research or exploratory data analysis. This means we may need to access raw data multiple times and create visualizations from various intermediate or resulting values obtained from raw data. Generally, (adaptive) composition plays a significant role in addressing this situation. The matrix-level Laplace mechanism provides a strong "pure" privacy guarantee ($\delta = 0$), but by accepting a slightly weaker guarantee ($\delta > 0$), we can enhance the privacy budget using composition methods, resulting in reduced noise perturbation. For the Gaussian mechanism, we are interested in comparing the matrix- and entry-level analytic Gaussian mechanisms.

Remark 1. We also investigated using the Exponential mechanism [44] to perform a randomized mapping from $\bar{X}$ to the set of all RGB images $\mathcal{H}$, either as a whole (which we denote by Exp) or entry-by-entry ($\mathsf{Exp_k}$). The Exp method guarantees $(\epsilon, 0)$-DP by sampling $H \in \mathcal{H}$ using the utility function $u(\bar{X}, H) = -\|\bar{X} - H\|_2$. This is impractical because $|\mathcal{H}|$ is too large. The entry-level approach uses an $(\epsilon', 0)$ mechanism with exact composition to map each $\bar{X}_{ij}$ to a color $C \in \mathcal{C}$ using the utility function $-|\bar{X}_{ij} - C|$ with sensitivity $2/N$. This produced significantly worse results than the other approaches.

Pre-processing Approaches. Given the privacy parameters ϵ and δ, some pre-processing approaches can be used to reduce the amount of required noise. One way to achieve this is by reducing the range. According to prior information about FNC, it is known that extremely high correlations (e.g., $|\bar{X}_{ij}| \geq 0.9, i \neq j$) are rare. Thus, it is possible to limit the range of $\bar{X}_{ij}$ to an interval $[-u, u]$, where $0 < u < 1$. As a result, the ℓ_1 and ℓ_2 sensitivity of $\bar{X}$ become $\Delta_1 = ur(r-1)/N$ and $\Delta_2 = u\sqrt{2r(r-1)}/N$, and the sensitivity of each entry $\bar{X}_{ij}$ is $\Delta = 2u/N$.

Algorithm 1 Method to choose clipping bound u

Input: A group of FNC matrices $X^{(1)}, X^{(2)}, ..., X^{(N)}$, privacy budget ϵ_u for each possible bound
Output: A list of $\{\tilde{s}_u\}$ for all possible u
for $u = 0.4, 0.45, 0.5, ..., 1$ **do**
1. $s_u = 0$
2. **for** $i = 1, ..., r; j > i; n = 1, ..., N$ **do**
   If $|X_{ij}^{(n)}| \leq u$: $s_u = s_u + X_{ij}^{(n)}$;
   If $X_{ij}^{(n)} > u$ or $X_{ij}^{(n)} < -u$: $s_u = s_u + u$ or $s_u = s_u - u$, respectively.
3. $\tilde{s}_u = s_u + \text{Laplace}(0, ur(r-1)/\epsilon_u)$

In this work, we employ a clipping method [60] to estimate the bounds of FNC range (see Algorithm 1). The idea is to establish upper and lower bounds for $X_{ij}^{(n)}$. The process involves gradually increasing the bound u from 0 to 1. If $X_{ij}^{(n)} > u$ or $X_{ij}^{(n)} < -u$, the value is clipped, setting $X_{ij}^{(n)}$ to u or $-u$ respectively. We then calculate the sum of all entries for each possible bound and use the Laplace mechanism to ensure the privacy of the total sum. Finally, we select a bound that results in a stable perturbed sum (see Fig. 3).

Alternative techniques can be employed to estimate the bounds. We can use prior domain knowledge, estimation from public data, or estimation using different DP methods. For example, we can use public datasets to choose the upper and lower bounds; we may consider estimating the differentially private quantiles and use them as the bounds [44], [61]; Liu et al. [29] propose an algorithm to estimate the bound by optimizing the expected mean squared error between perturbed aggregated data and the original counterpart. No matter which method is used to determine the bounds, there is a tradeoff between information loss and the amount of noise needed to guarantee differential privacy. When the upper and lower bounds are set closer to each other, less noise is needed to guarantee differential privacy due to a lower sensitivity.

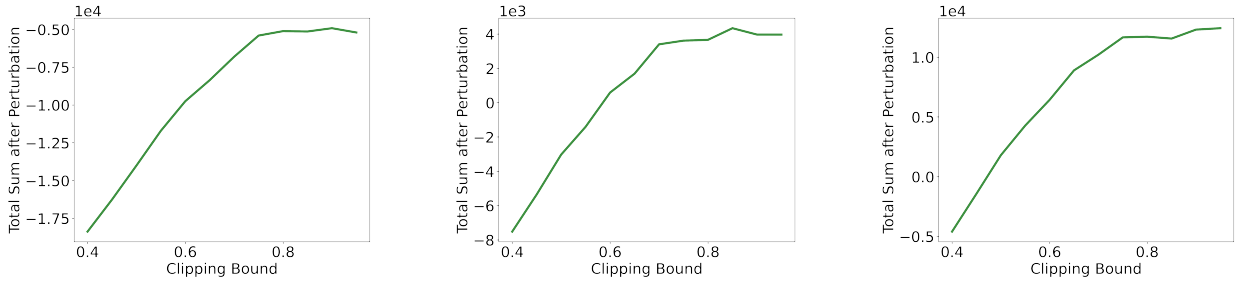

(a) FBIRN: $u = 0.75$ (b) BLSA: $u = 0.7$ (c) ADNI: $u = 0.75$

Fig. 3. Clipping method for choosing u on three datasets. The overall privacy budget used for bound estimation is 0.5.

However, using smaller clipping bounds results in the removal of a significant amount of information.

Post-processing Approaches. Differential privacy has a post-processing invariance property, meaning that any data-independent operations on differentially private outputs do not incur additional privacy loss. Hence, post-processing can be used to "clean up" images without compromising the privacy. This can help reduce the impact of perturbation on FNC or any other visualizations and potentially improve visual utility.

For FNC visualization, we explore two post-processing approaches: Singular Value Decomposition (SVD) and Haar wavelet transform [62]. The SVD of differentially private output is $\tilde{X} = U\Sigma V^\intercal$. We choose the first p components after observing the magnitudes of the singular values. If the added noise does not destroy the visualization pattern, we are able to observe that the first p singular values are significantly larger than the remaining ones. We use $\tilde{X}_s = U_p\Sigma_p V_p^\intercal$ for visualization after post-processing. In the wavelet transform, a transformation is performed to obtain the matrix $T = P^\intercal \tilde{X} P$, where P is the invertible Haar transformation matrix. In many cases, the wavelet transformed matrices are sparser than their original counterparts. We threshold the coefficients in matrix T to generate a new coefficient matrix, which we denote as T_w. Subsequently, we perform the inverse wavelet transform on T_w to obtain $\tilde{X}_w = (P^\intercal)^{-1} T_w P^{-1}$. $\tilde{X}_w$ represents the processed data used for visualization.

B. Connectogram Visualization

In this section, we introduce a novel workflow for DP connectogram visualization in group comparison, which uses a nonparametric test to identify significant connectivity.

Non-private Workflow for Group Comparison. Connectogram visualizations is an effective way for representing group differences in brain regions and connectivity (e.g. comparing healthy controls (HC) with schizophrenia (SZ) patients). The typical workflow to generate it involves transforming each element in FNC matrix into Fisher's Z-scores. The significance of the connectivity strength between HC and SZ groups is determined by conducting a two-tailed two-sample t-test ($p < 0.05$ with Bonferroni correction) [63].

Private Workflow for Group Comparison. We propose a novel differentially private connectogram workflow. The most intuitive approach is to add perturbation to either the p-value or the t-statistic. However, this can be challenging because computing the sensitivity of the t-statistic is challenging and noise can significantly change the p-value. Our experiments showed that under differential privacy it is more effective

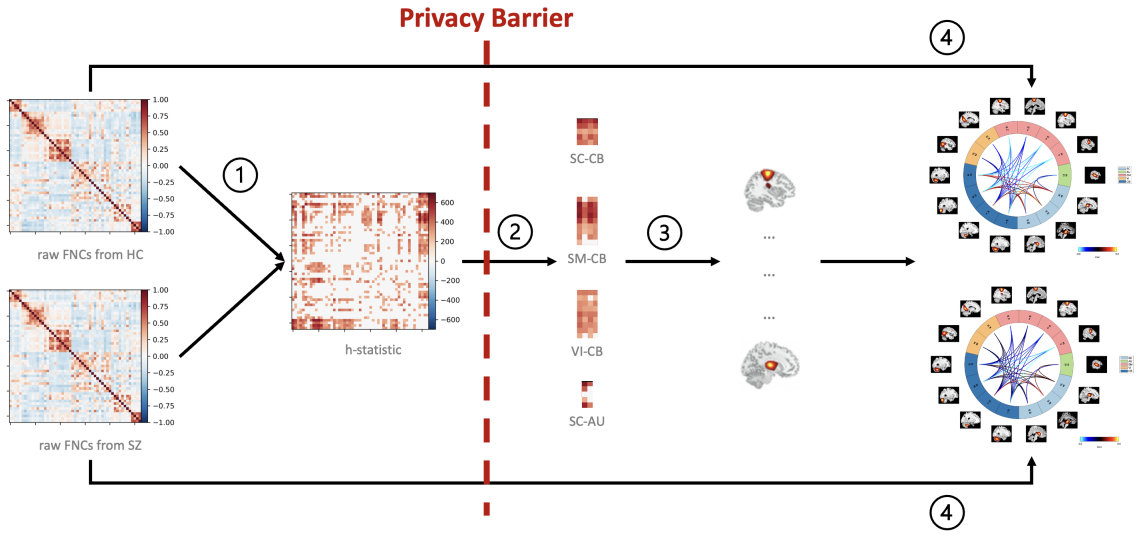

Fig. 4. Privacy-preserving workflow for connectogram visualization.

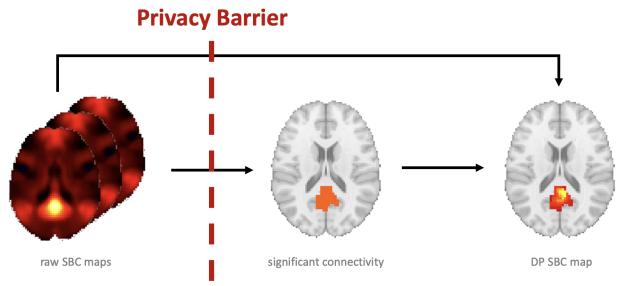

Fig. 5. Privacy-preserving workflow for SBC map visualization.

to use nonparametric tests [64] for constructing the connectogram. We therefore propose a private workflow based on a nonparametric absolute Kruskal-Wallis (AKW) test. Consider a database X with a size of N and g groups. Let N_i be the size of each group and r_{ij} be the rank of the j-th element of group i. Define $\bar{r}_i = \frac{1}{N_i}\sum_{j=1}^{N_i} r_{ij}$ as the mean rank of group i and $\bar{r} = \frac{N+1}{2}$ as the average of all the ranks. The absolute Kruskal-Wallis h-statistic is

$$h = (N-1)\frac{\sum_{i=1}^{g} N_i|\bar{r}_i - \bar{r}|}{\sum_{i=1}^{g}\sum_{j=1}^{N_i}|r_{ij} - \bar{r}|} \tag{2}$$

and has a sensitivity of 8 [64, Algorithm $\widetilde{\text{KW}}_{\text{P}}$].

Our workflow (see Fig. 4) includes the following steps: (1) We perform the AKW test using FNC matrices from two groups to obtain an h-statistic matrix. (2) We calculate the average h-statistic for each h-statistic submatrix. The submatrix is obtained by partitioning the matrix based on different brain domains [63, Table 2]. Then we apply the Report Noisy Max (RNM) algorithm [7] to privately select the first p_d largest submatrices. (3) We average the h-statistic for each brain region within the selected brain domains and again use the RNM algorithm to privately select the first p_r largest brain regions. (4) Finally, we obtain perturbed connectivities from the average FNC of each group. With privacy-preserving nodes and edges in place, we are able to visualize the connectogram for group comparisons.

IV. SBC Visualization

A. Non-private Workflow for Significant Connectivity

SBC characterizes the connectivity pattern between a predefined seed and the entire brain. This involves computing Fisher-transformed bivariate correlation coefficients between the seed BOLD timeseries and the timeseries of each individual voxel [10]. The correlation between the seed and any voxel v is defined as

$$\frac{\sum_{t=1}^{T} S(\cdot,t)S(v,t)}{\sqrt{\sum_{t=1}^{T} S^2(\cdot,t)}\sqrt{\sum_{t=1}^{T} S^2(v,t)}},$$

where $S(\cdot,t)$ is the BOLD timeseries from the seed (either a single voxel or an ROI) at time t, and T is the number of time points in the experiment. The correlation coefficient is then converted to a z-score using the Fisher Z-transform. Subsequently, a one-sample t-test is applied to assess whether the average connectivity differs significantly from zero, indicating a significant positive or negative connectivity.

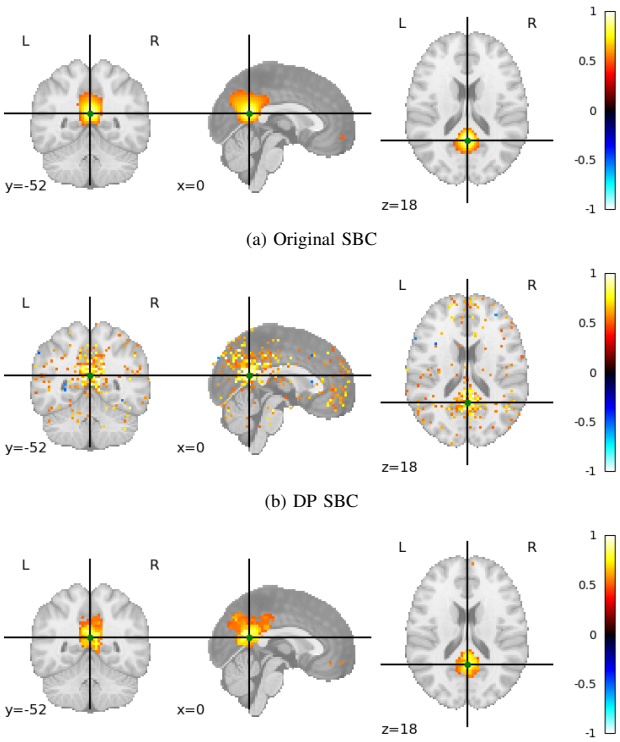

(a) Original SBC

(b) DP SBC

(c) DP SBC after applying a median filter

Fig. 6. SBC maps on the NITRC dataset using the Posterior Cingulate Cortex (PCC) as the seed (indicated by the green dot), with parameters $N = 1082, \epsilon = 2.5$, and $\delta = N^{-3/2}$. Correlations with absolute values greater than 0.5 are being displayed.

B. Private Workflow for Significant Connectivity

The privacy-preserving SBC visualization process is quite similar to connectogram visualization and also involves two main components (see Fig. 5): privately selecting significant connectivities and privately displaying z-scores. To identify statistically significant connectivities, we employ the Wilcoxon signed-rank (WSR) test, a nonparametric alternative to the one-sample t-test, as nonparametric tests are often more effective than their parametric counterparts. For a given database X, let s_i be the sign of the data x_i, and r_i be the rank of $|x_i|$, where the rank is determined based on the magnitude of the data. The WSR test statistic is calculated as $w = \sum s_i r_i$, and we select the connectivities with a significance level of $p < 0.05$ after Bonferroni correction [64, Algorithm $\widetilde{\text{WP}}_{\text{P}}$].

To protect z-scores, we utilize the analytic Gaussian mechanism for noise generation, adding it to correlations, and then apply the Fisher Z-transform. As mentioned earlier, while we could reduce the range of correlations to minimize the required noise, such an approach is not suitable here due to the expectation of presenting high correlations (see Fig. 6a). When displaying correlations, we observe that the privacy-protected correlation map exhibits salt-and-pepper noise (see Fig. 6b), a phenomenon not identified in the FNC visualization. One possible reason could be that the connections within the same brain network are typically strong, leading to relatively stable and similar correlation values. Therefore, we use a median filter [65] as a post-processing method to remove the noise (see Fig. 6c). With the identification of significant connectivities

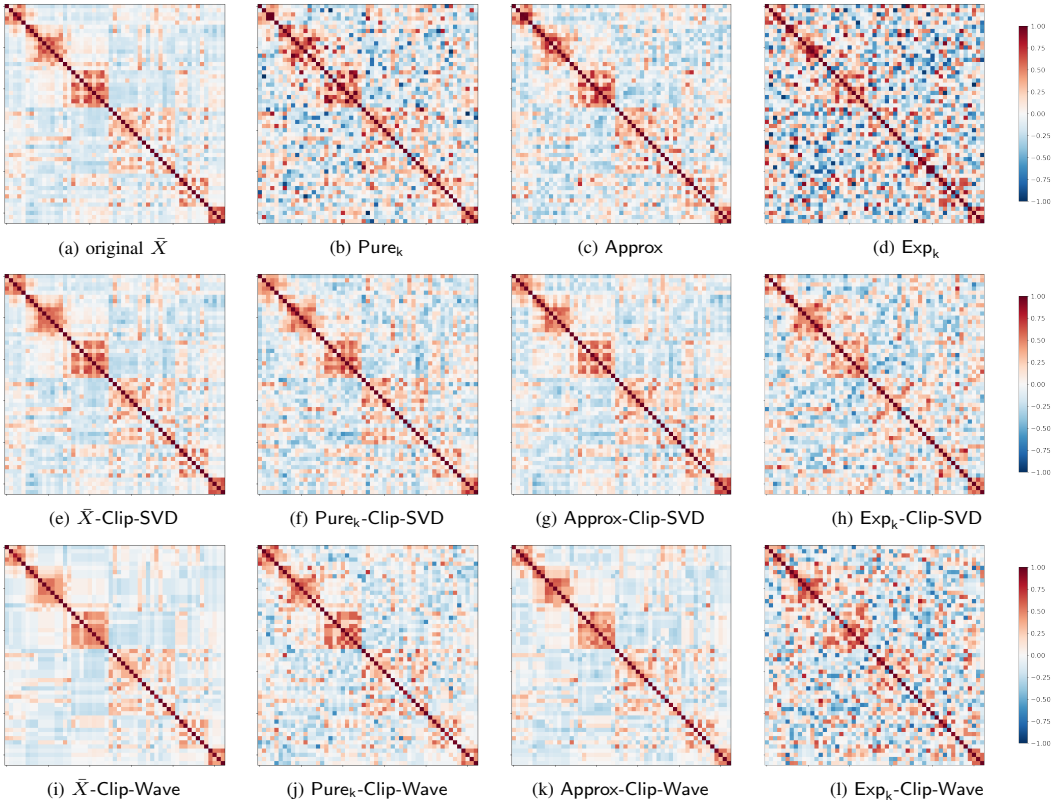

Fig. 7. Visualization of the FBIRN dataset with $N = 1056$, clipping bound $u = 0.75$, number of components used in SVD is 10, threshold used in wavelet transform is 0.2, and the number of colors is 256. The first row represents the methods without pre- or post-processing, the second row represents the methods with clipping and SVD post-processing, and the third row represents the methods with clipping and wavelet post-processing. The overall privacy guarantee for all perturbed FNC matrices is $(\epsilon = 2, \delta = N^{-3/2})$-DP. Unsurprisingly, the Laplace mechanism performs worse than the analytic Gaussian mechanism since the ℓ_1 sensitivity is larger than the ℓ_2 sensitivity in high dimensions. Even though we apply composition methods to reduce the noise, Pure$_k$ still does not outperform Approx. The Exponential mechanism is barely usable. Clipping and SVD post-processing can significantly improve the visual fidelity of the private FNC matrices. However, wavelet post-processing can introduce additional visual distortion. This indicates that using classical image processing techniques developed for natural images (e.g. photos) is not suitable for this type of visualization.

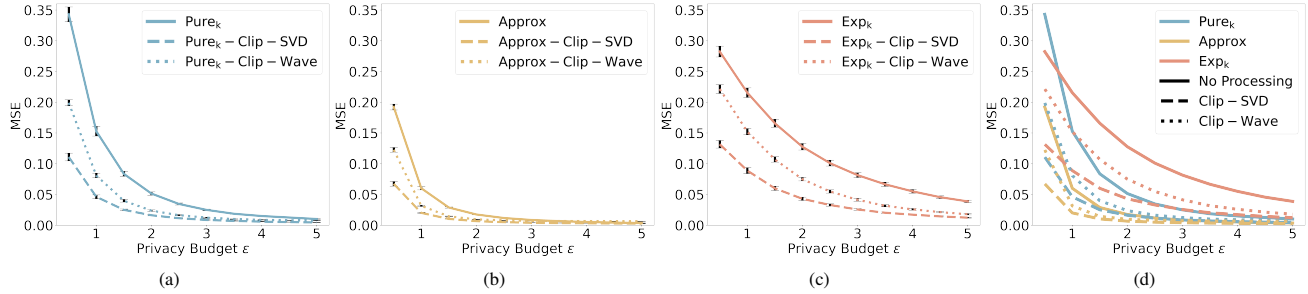

Fig. 8. MSE between the original average FNC and its private counterparts on FBIRN dataset. Comparison among (a) Pure$_k$-; (b) Approx-; (c) Exp$_k$-related methods; (d) all methods. Each method run 50 times, and the error bars are provided.

and their corresponding values, we are able to construct the privacy-preserving SBC map.

V. Experiments and Results

We aim to investigate the following questions: (1) Which algorithm is most effective for visualizing functional connectivity? (2) Are there potential benefits from pre- and post-processing methods? (3) Can the differentially private workflow for connectogram visualization identify significant brain regions and connections as effectively as its non-private counterpart? (4) Is the differentially private workflow for SBC visualization as effective as its non-private counterpart in identifying significant connectivities? (5) How do numerical utility measures like MSE correspond to the quality of visualizations? We conducted privacy-preserving FNC visualization experiments using datasets from the Function Biomedical Informatics Research Network (FBIRN) study, the Baltimore Longitudinal Study of Aging (BLSA) study, and the Alzheimer's Disease Neuroimaging Initiative (ADNI) study [63], [66], [67]. Our findings and results are consistent across all the datasets. We tested our connectogram visualizations using the FBIRN dataset, which compare HC and SZ. Additionally, we tested

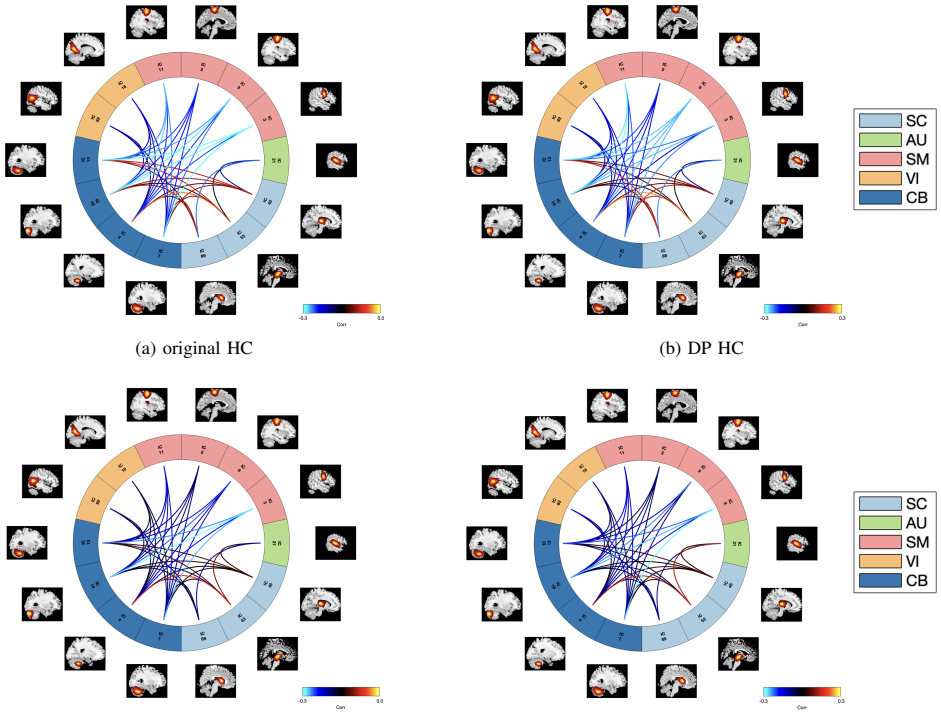

(a) original HC (b) DP HC

(c) original SZ (d) DP SZ

Fig. 9. Privacy-preserving connectogram visualization with $\epsilon = 2, p_d = 5, p_r = 14$. We allocate 45% of the privacy budget to identify brain domains, another 45% to pinpoint brain regions, and reserve the remaining 10% to safeguard the connectivities.

the SBC visualizations using the NITRC dataset.

A. Comparing Methods

In general, we found that Gaussian mechanisms worked better than Laplace mechanisms, which were in turn better than Exponential mechanisms (see Fig. 7, Fig. A.2, and Fig. A.3). Overall, Approx-Clip-SVD has the best performance. Using composition methods such as $\mathsf{Approx_r}$-Clip-SVD could also achieve competitive visualization.

Because the ℓ_2 sensitivity is lower than the ℓ_1 sensitivity in high dimensions, the analytic Gaussian mechanism adds less noise compared to the Laplace mechanism, resulting in better performance. The Exponential mechanism exhibits the worst visualization performance due to its high level of randomness. Specifically, the PDF of the output from $\mathsf{Pure_k}$ for each entry is proportional to $\exp(-N\epsilon'|z_{ij} - \bar{X}_{ij}|/2)$, while the PDF of the output of $\mathsf{Exp_k}$ is proportional to $\exp(-N\epsilon'|c - \bar{X}_{ij}|/4)$. In general, the Exponential mechanism is more suitable for scenarios such as determining the selling price of products where a small perturbation in the determined price may lead to significant changes in revenue (utility). However, for FNC matrices, small changes in correlation do not result in significant changes in color assignment.

In terms of composition, exact composition and RDP-based composition perform better than advanced composition, regardless of which mechanism is used. For the Laplace mechanism, $\mathsf{Pure_k}$ requires less noise than $\mathsf{Pure_r}$, while the opposite is true for the analytic Gaussian mechanism, $\mathsf{Approx_k}$ requires more noise.

B. Processing Effect

We find that pre- and post-processing significantly improve privacy-preserving data visualization. For FNC visualization, Fig. 7g and 7k use different processing methods, and both of them exhibit the best visualization results, i.e., they are visually closest to the original average FNC shown in Fig. 7a. To quantitatively analyze this finding, we also calculate the MSE between the original average FNC and its private counterpart. As shown in Fig. 8 and Fig. A.1, pre- and post-processing have a more significant effect when the total privacy budget ϵ is small (high privacy), but the advantages are less significant if ϵ is relatively large. The post-processing method using SVD is better than wavelet transform. For SBC visualization, pre-processing methods are not applicable in this case. We apply the median filter technique as the post-processing method, resulting in a noticeable improvement in visualization quality (see Fig. 6). It is evident that different visualization types may require distinct pre- and/or post-processing methods.

C. Workflow Effectiveness

For connectogram visualization, we observe that our proposed workflow is able to produce similar visualizations as its non-private counterpart (see Fig. 9). Our work also finds that brain functional abnormalities in SZ were mainly located in the sub-cortical (SC), auditory (AU), sensorimotor (SM), visual (VI), and cerebellar (CB) domains [63]. Given the importance of identifying these brain domains and regions rather than determining the exact connectivity values in this context, we allocate more privacy budget to determine the brain domains

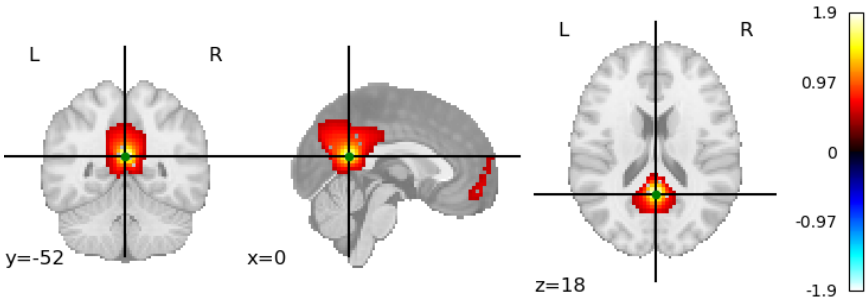

(a) Original SBC

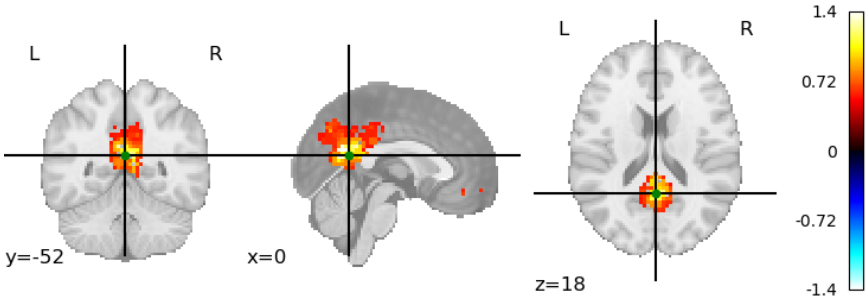

(b) DP SBC after applying a median filter

Fig. 10. SBC maps on the NITRC dataset using the Posterior Cingulate Cortex (PCC) as the seed (indicated by the green dot), with parameters $N = 1082, \epsilon = 5$, and $\delta = N^{-3/2}$. We allocate 50% of the privacy budget to identify significant correlations and 50% to protect the z-scores.

and regions. Similar outcomes are observed in SBC visualization. The privacy-preserving SBC map closely aligns with the SBC map that does not consider privacy protection. For instance, Fig. 10 illustrates strong connectivities between the selected seed PCC and its neighboring voxels.

D. Visual Utility

It should be noted that there is no standard way for quantitatively assessing the visualization quality for privacy-preserving data. As discussed by Zhang et al. [17], making a precise notion of "visual utility" is challenging. Errors are often measured using metrics such as the ℓ_1 or ℓ_2 norms, but these may not align with the perceived quality. Moreover, in many cases, norm metrics may not be applicable for assessing visualization quality. For example in connectogram visualization, comparing hypothesis testing errors (seeing irrelevant domains or regions or failing to notice significant domains or regions) could provide a better evaluation method because our primary concern is identifying domains or regions that exhibit significant differences in group comparisons. When it comes to SBC visualization, finding quantitative methods to assess visual utility becomes even more challenging. So in many cases, we may still rely on professionals to conduct assessments.

VI. Future Directions

Data visualizations are meant to allow qualitative assessments of quantitative data. When visualizing sensitive data, this gap between qualitative and quantitative offers many interesting options for guaranteeing privacy. In this study, we examined some of the basic approaches. When considering federated learning, the differences among these methods would be much smaller because all these methods would require less noise and provide better utility due to the increased sample size. There are many other directions to explore for future work in this and other visualizations. For example, exploring alternative processing methods such as using Fourier transform and bit plane slicing method. We can also transform the data to obtain sparse coefficients, apply privacy-preserving techniques to protect those coefficients, and then use the inverse transform to get the privacy-preserving data. In addition to refining these approaches, manipulating the choice of color map, the number of colors $|\mathcal{C}|$, and the quantization scheme can all potentially improve the visual effect.

Visualizing uncertainty has gained significant attention in the past, and numerous studies have discussed the necessity of presenting uncertain information in visualizations [68]–[70]. Visualizing uncertainty for privacy-preserving visualizations is also a potential research direction. Traditional techniques such as error bars and confidence intervals can be adopted, making them applicable to privacy-preserving settings, to convey the level of uncertainty associated with patterns in a visualization.

We briefly discussed the risks of sensitive information leakage in visualization. Because of this, the importance of privacy protection in visualization is growing. Current research has demonstrated various privacy attacks in different fields, such as machine learning [71]. Adversaries may aim to obtain the training dataset, model parameters, hyperparameters, architecture, etc. Inspired by these considerations, we believe that privacy attacks in visualization are also a field worth exploring.

Acknowledgments

We acknowledge the assistance of ChatGPT for English language editing.

REFERENCES

[1] Y. Tao, A. D. Sarwate, S. Panta, S. Plis, and V. D. Calhoun, "Privacy-preserving visualization of brain functional network connectivity," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2024, pp. 1–5.

[2] C. Kostelnick, "The visual rhetoric of data displays: The conundrum of clarity," *IEEE Transactions on Professional Communication*, vol. 51, no. 1, pp. 116–130, 2008.

[3] K. Bhattacharjee, M. Chen, and A. Dasgupta, "Privacy-preserving data visualization: reflections on the state of the art and research opportunities," in *Computer Graphics Forum*, vol. 39, no. 3. Wiley Online Library, 2020, pp. 675–692.

[4] L. Sweeney, "Weaving technology and policy together to maintain confidentiality," *The Journal of Law, Medicine & Ethics*, vol. 25, no. 2-3, pp. 98–110, 1997.

[5] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 2008, pp. 111–125.

[6] C. Dwork, A. Smith, T. Steinke, and J. Ullman, "Exposed! a survey of attacks on private data," *Annual Review of Statistics and Its Application*, vol. 4, pp. 61–84, 2017.

[7] C. Dwork, "Differential privacy," in *International colloquium on automata, languages, and programming*. Springer, 2006, pp. 1–12.

[8] M. P. Van Den Heuvel and H. E. H. Pol, "Exploring the brain network: a review on resting-state fMRI functional connectivity," *European neuropsychopharmacology*, vol. 20, no. 8, pp. 519–534, 2010.

[9] R. Mohanty, W. A. Sethares *et al.*, "Rethinking measures of functional connectivity via feature extraction," *Scientific reports*, vol. 10, no. 1, p. 1298, 2020.

[10] S. E. Joel, B. S. Caffo, P. C. Van Zijl, and J. J. Pekar, "On the relationship between seed-based and ICA-based measures of functional connectivity," *Magnetic resonance in medicine*, vol. 66, no. 3, pp. 644–657, 2011.

[11] S. M. Plis, A. D. Sarwate, D. Wood, C. Dieringer, D. Landis, C. Reed, S. R. Panta, J. A. Turner, J. M. Shoemaker, K. W. Carter *et al.*, "COINSTAC: a privacy enabled model and prototype for leveraging and processing decentralized brain imaging data," *Frontiers in neuroscience*, vol. 10, p. 204805, 2016.

[12] A. D. Sarwate, S. M. Plis, J. A. Turner *et al.*, "Sharing privacy-sensitive access to neuroimaging and genetics data: a review and preliminary validation," *Frontiers in neuroinformatics*, vol. 8, p. 35, 2014.

[13] J. Ming, E. Verner, A. Sarwate *et al.*, "COINSTAC: Decentralizing the future of brain imaging analysis," *F1000Research*, vol. 6, 2017.

[14] K. Rootes-Murdy, H. Gazula *et al.*, "Federated analysis of neuroimaging data: a review of the field," *Neuroinformatics*, pp. 1–14, 2022.

[15] C. Dwork, "Differential privacy: A survey of results," in *International conference on theory and applications of models of computation*. Springer, 2008, pp. 1–19.

[16] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[17] D. Zhang, M. Hay, G. Miklau, and B. O'Connor, "Challenges of visualizing differentially private data," *Theory and Practice of Differential Privacy*, vol. 2016, pp. 1–3, 2016.

[18] J. Zhou, X. Wang *et al.*, "Dpviscreator: Incorporating pattern constraints to privacy-preserving visualizations via differential privacy," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 1, pp. 809–819, 2022.

[19] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," in *International conference on machine learning*. PMLR, 2015, pp. 1376–1385.

[20] J. Dong, D. Durfee, and R. Rogers, "Optimal differential privacy composition for exponential mechanisms," in *International Conference on Machine Learning*. PMLR, 2020, pp. 2597–2606.

[21] J. Dong, A. Roth, and W. J. Su, "Gaussian differential privacy," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 84, no. 1, pp. 3–37, 2022.

[22] I. Mironov, "Rényi differential privacy," in *2017 IEEE 30th computer security foundations symposium (CSF)*. IEEE, 2017, pp. 263–275.

[23] D. Avraam, R. Wilson *et al.*, "Privacy preserving data visualizations," *EPJ data science*, vol. 10, no. 1, p. 2, 2021.

[24] P. Nanayakkara, J. Bater *et al.*, "Visualizing privacy-utility trade-offs in differentially private data releases," *arXiv preprint arXiv:2201.05964*, 2022.

[25] M. Budiu, P. Thaker *et al.*, "Overlook: Differentially private exploratory visualization for big data," *Journal of Privacy and Confidentiality*, vol. 12, no. 1, 2022.

[26] D. Zhang, A. Sarvghad, and G. Miklau, "Investigating visual analysis of differentially private data," *IEEE transactions on visualization and computer graphics*, vol. 27, no. 2, pp. 1786–1796, 2020.

[27] H. B. Lee, "Visualization and differential privacy," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2017.

[28] L. Panavas, T. Crnovrsanin *et al.*, "Investigating the Visual Utility of Differentially Private Scatterplots," *IEEE Transactions on Visualization and Computer Graphics*, 2023.

[29] A. Liu, L. Xia *et al.*, "Differential privacy for eye-tracking data," in *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, 2019, pp. 1–10.

[30] B. Ghazi, J. He *et al.*, "Differentially private heatmaps," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 6, 2023, pp. 7696–7704.

[31] S. Garfinkel *et al.*, *De-identification of Personal Information:*. US Department of Commerce, National Institute of Standards and Technology, 2015.

[32] L. Sweeney, "k-anonymity: A model for protecting privacy," *International journal of uncertainty, fuzziness and knowledge-based systems*, vol. 10, no. 05, pp. 557–570, 2002.

[33] M. M. Merener, "Theoretical results on de-anonymization via linkage attacks." *Trans. Data Priv.*, vol. 5, no. 2, pp. 377–402, 2012.

[34] J. Henriksen-Bulmer and S. Jeary, "Re-identification attacks—A systematic literature review," *International Journal of Information Management*, vol. 36, no. 6, pp. 1184–1192, 2016.

[35] M. Fiore, P. Katsikouli *et al.*, "Privacy in trajectory micro-data publishing: a survey," *Transactions on Data Privacy*, vol. 13, no. 2, pp. 91–149, 2020.

[36] C. T. Allen, "Finding the Enemy on the Data-Swept Battlefield of 2035," *Military Review*, 2020.

[37] N. Homer, S. Szelinger *et al.*, "Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays," *PLoS genetics*, vol. 4, no. 8, p. e1000167, 2008.

[38] C. Dwork, A. Smith *et al.*, "Robust traceability from trace amounts," in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE, 2015, pp. 650–669.

[39] R. Shokri, M. Stronati *et al.*, "Membership inference attacks against machine learning models," in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 3–18.

[40] L. Wasserman and S. Zhou, "A statistical framework for differential privacy," *Journal of the American Statistical Association*, vol. 105, no. 489, pp. 375–389, 2010. [Online]. Available: http://dx.doi.org/10.1198/jasa.2009.tm08651

[41] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1176–1184, 2015.

[42] W. Alghamdi, S. Asoodeh, F. P. Calmon *et al.*, "Cactus mechanisms: Optimal differential privacy mechanisms in the large-composition regime," in *2022 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2022, pp. 1838–1843.

[43] B. Balle and Y.-X. Wang, "Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *International Conference on Machine Learning*. PMLR, 2018, pp. 394–403.

[44] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, 2007, pp. 94–103.

[45] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 2010, pp. 51–60.

[46] J. Murtagh and S. Vadhan, "The complexity of computing the optimal composition of differential privacy," in *Theory of Cryptography Conference*. Springer, 2015, pp. 157–175.

[47] M. Abadi, A. Chu, I. Goodfellow *et al.*, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.

[48] B. Balle, G. Barthe *et al.*, "Hypothesis testing interpretations and renyi differential privacy," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2496–2506.

[49] J. Dong, A. Roth, and W. J. Su, "Gaussian differential privacy," *arXiv preprint arXiv:1905.02383*, 2019.

[50] D. Sommer, S. Meiser, and E. Mohammadi, "Privacy loss classes: The central limit theorem in differential privacy," *Cryptology ePrint Archive*, 2018.

[51] W. Alghamdi, S. Asoodeh, F. P. Calmon *et al.*, "The saddle-point accountant for differential privacy," *arXiv preprint arXiv:2208.09595*, 2022.

[52] A. Koskela and A. Honkela, "Learning rate adaptation for differentially private learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2465–2475.

[53] S. Gopi, Y. T. Lee, and L. Wutschitz, "Numerical composition of differential privacy," *Advances in Neural Information Processing Systems*, vol. 34, pp. 11 631–11 642, 2021.

[54] A. Koskela, J. Jälkö, and A. Honkela, "Computing tight differential privacy guarantees using fft," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2560–2569.

[55] B. Ghazi, P. Kamath *et al.*, "Faster privacy accounting via evolving discretization," in *International Conference on Machine Learning*. PMLR, 2022, pp. 7470–7483.

[56] V. Doroshenko, B. Ghazi *et al.*, "Connect the dots: Tighter discrete approximations of privacy loss distributions," *arXiv preprint arXiv:2207.04380*, 2022.

[57] N. Tzourio-Mazoyer, B. Landeau *et al.*, "Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain," *Neuroimage*, vol. 15, no. 1, pp. 273–289, 2002.

[58] R. C. Craddock, G. A. James *et al.*, "A whole brain fMRI atlas generated via spatially constrained spectral clustering," *Human brain mapping*, vol. 33, no. 8, pp. 1914–1928, 2012.

[59] E. M. Gordon, T. O. Laumann *et al.*, "Generation and evaluation of a cortical area parcellation from resting-state correlations," *Cerebral cortex*, vol. 26, no. 1, pp. 288–303, 2016.

[60] J. P. Near and C. Abuah, "Programming differential privacy," *URL: https://uvm*, 2021.

[61] A. Smith, "Privacy-preserving statistical estimation with optimal convergence rates," in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 2011, pp. 813–822.

[62] C. Mulcahy, "Image compression using the Haar wavelet transform," *Spelman Science and Mathematics Journal*, vol. 1, no. 1, pp. 22–31, 1997.

[63] Y. Du, Z. Fu, J. Sui *et al.*, "NeuroMark: An automated and adaptive ICA based pipeline to identify reproducible fMRI markers of brain disorders," *NeuroImage: Clinical*, vol. 28, p. 102375, 2020.

[64] S. Couch, Z. Kazan, K. Shi *et al.*, "Differentially private nonparametric hypothesis testing," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 737–751.

[65] U. Erkan, L. Gökrem, and S. Enginoğlu, "Different applied median filter in salt and pepper noise," *Computers & Electrical Engineering*, vol. 70, pp. 789–798, 2018.

[66] M. S. Sendi, E. Zendehrouh, C. A. Ellis *et al.*, "Aberrant dynamic functional connectivity of default mode network in schizophrenia and links to symptom severity," *Frontiers in Neural Circuits*, vol. 15, p. 649417, 2021.

[67] L. Ferrucci, "The Baltimore Longitudinal Study of Aging (BLSA): a 50-year-long journey and plans for the future," pp. 1416–1419, 2008.

[68] A. M. MacEachren, "Visualizing uncertain information," *Cartographic perspectives*, no. 13, pp. 10–19, 1992.

[69] A. M. MacEachren, A. Robinson, S. Hopper *et al.*, "Visualizing geospatial information uncertainty: What we know and what we need to know," *Cartography and Geographic Information Science*, vol. 32, no. 3, pp. 139–160, 2005.

[70] K. Potter, J. Kniss, R. Riesenfeld, and C. R. Johnson, "Visualizing summary statistics and uncertainty," in *Computer Graphics Forum*, vol. 29, no. 3. Wiley Online Library, 2010, pp. 823–832.

[71] M. Rigaki and S. Garcia, "A survey of privacy attacks in machine learning," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–34, 2023.

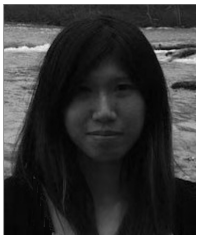

Ye Tao received the BS degree in Electrical Information Engineering from South China University of Technology, China, in 2014, and the MS degree in Electrical and Computer Engineering from the University of Rochester, US, in 2019. From 2019 to 2020, she worked at CogT Lab in University of Rochester as a research assistant, information analyst and data manager. She is currently working toward her PhD degree in the Department of Electrical and Computer Engineering at Rutgers, The State University of New Jersey, USA. Her research interests include differential privacy, signal processing, machine learning and visualization.

Anand D. Sarwate (S'99–M'09–SM'14) received the B.S. degrees in electrical engineering and computer science and mathematics from the Massachusetts Institute of Technology (Cambridge, MA, USA), in 2002, and the M.S. and Ph.D. degrees in electrical engineering from the Department of Electrical Engineering and Computer Sciences (EECS), University of California, Berkeley (Berkeley, CA, USA). He is a currently an Associate Professor with the Department of Electrical and Computer Engineering, The State University of New Jersey (New Brunswick, NJ, USA), where he has been since January 2014. He was previously a Research Assistant Professor from 2011 to 2013 with the Toyota Technological Institute at Chicago (Chicago, IL, USA); prior to this, he was a Postdoctoral Researcher from 2008 to 2011 with the University of California, San Diego (La Jolla, CA, USA). His research interests include information theory, machine learning, signal processing, optimization, and privacy and security. Dr. Sarwate is a Distinguished Lecturer of the IEEE Information Theory Society for 2024–2025, received the Rutgers Board of Governors Research Fellowship for Scholarly Excellence in 2020, and the NSF CAREER award in 2015. He is a member of Phi Beta Kappa and Eta Kappa Nu.

Sandeep Panta received his M.S. in Computer Engineering from University of New Mexico (Albuquerque, NM, USA) in 2008. He has worked on automation of various pre-processing pipelines used in the neuroscience community for analyzing human brain MRI scans. He is currently the Software Development Manager for COINSTAC (Collaborative Informatics and Neuroimaging Suite Toolkit for Anonymous Computation) software. COINSTAC is software to foster collaborative research, removing large barriers to traditional data-centric collaboration approaches. It enables groups of users to run common analyses on their own machines over their own datasets with ease.

Sergey M. Plis received the Ph.D. degree in Computer Science in 2007 from The University of New Mexico, Albuquerque, NM, USA. Currently, he is an associate professor of CS at the Georgia State University, and the Director of Machine Learning core at the TReNDS institute. His research interests within the fields of machine learning, AI, and data science lie in developing novel and applying existing techniques and approaches to analyzing large scale datasets. One of his key goals is to take advantage of the strengths of multiple data modalities and infer structure and patterns that are hard to obtain non-invasively and/or that are unavailable for direct observation. His ongoing work is focused on inferring multimodal probabilistic and causal descriptions based on fusion of fast and slow imaging modalities. This includes feature estimation via deep learning-based pattern recognition and learning causal graphical models. Besides academia, Dr. Plis has lead AI research teams developing deep tech AI products including NLP and speech processing solutions that are currently shipped to customers. Dr. Plis is supported by grants from NSF and NIH, he has authored more than a 100 publications, and serves on program committees of 6 major ML conferences.

Vince D. Calhoun is founding director of the tri-institutional Center for Translational Research in Neuroimaging and Data Science (TReNDS) where he holds appointments at Georgia State, Georgia Tech and Emory. He is the author of more than 1000 full journal articles. His work includes the development of flexible methods to analyze neuroimaging data including blind source separation, deep learning, multimodal fusion and genomics, neuroinformatics tools. Dr. Calhoun is a fellow of the Institute of Electrical and Electronic Engineers, The American Association for the Advancement of Science, The American Institute of Biomedical and Medical Engineers, The American College of Neuropsychopharmacology, The Organization for Human Brain Mapping (OHBM) and the International Society of Magnetic Resonance in Medicine. He currently serves on the IEEE BISP Technical Committee and is also a member of IEEE Data Science Initiative Steering Committee as well as the IEEE Brain Technical Committee.

APPENDIX

The appendix contain additional details and figures omitted from the main manuscript.

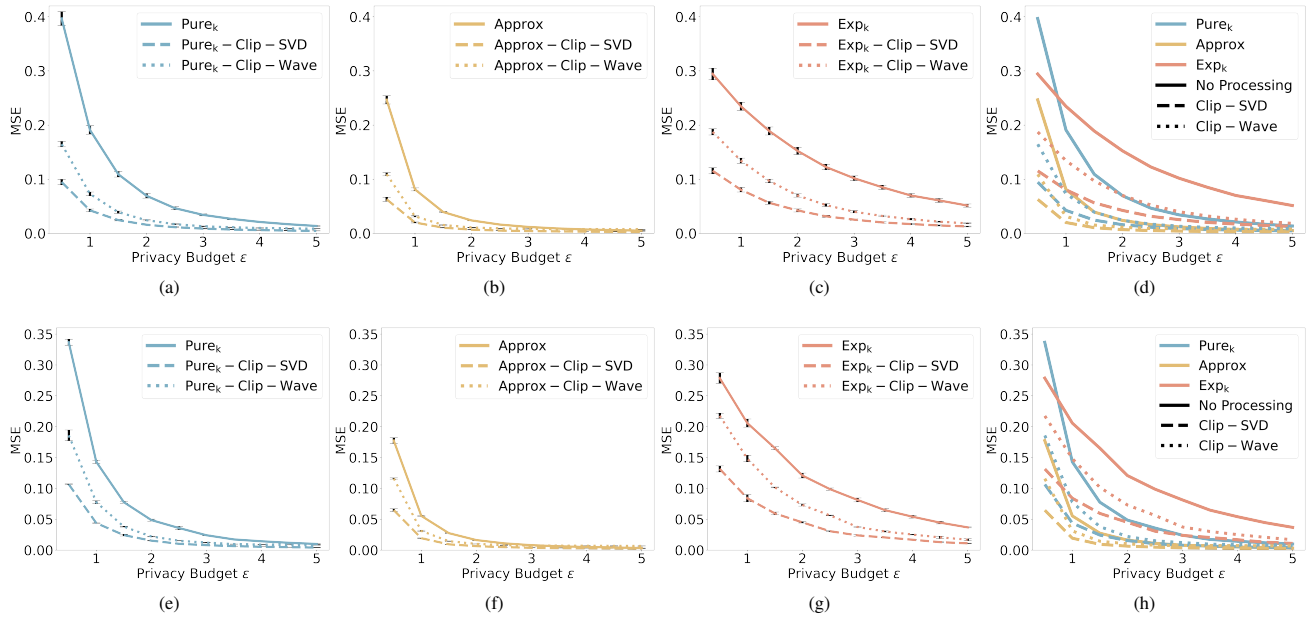

Fig. A.1. MSE between the original average FNC and its private counterparts on BLSA (see (a) - (d)) and ADNI (see (e) - (h)) datasets. Comparison among (a)(e) Pure$_k$-; (b)(f) Approx-; (c)(g) Exp$_k$-related methods; (d)(h) all methods. Each method run 50 times, and the error bars are provided.

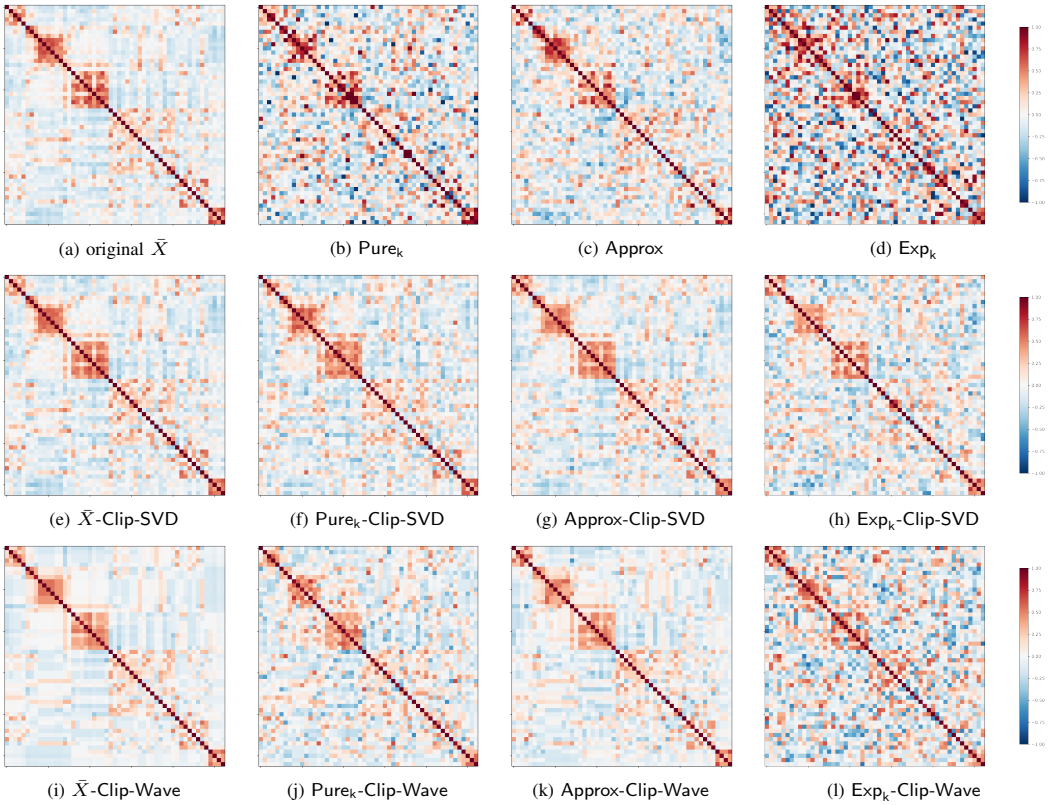

Fig. A.2. Visualization of the BLSA dataset with $N = 891, \epsilon = 2, \delta = N^{-3/2}$, clipping bound $u = 0.7$, number of components used in SVD is 10, threshold used in wavelet transform is 0.2, and the number of colors is 256.

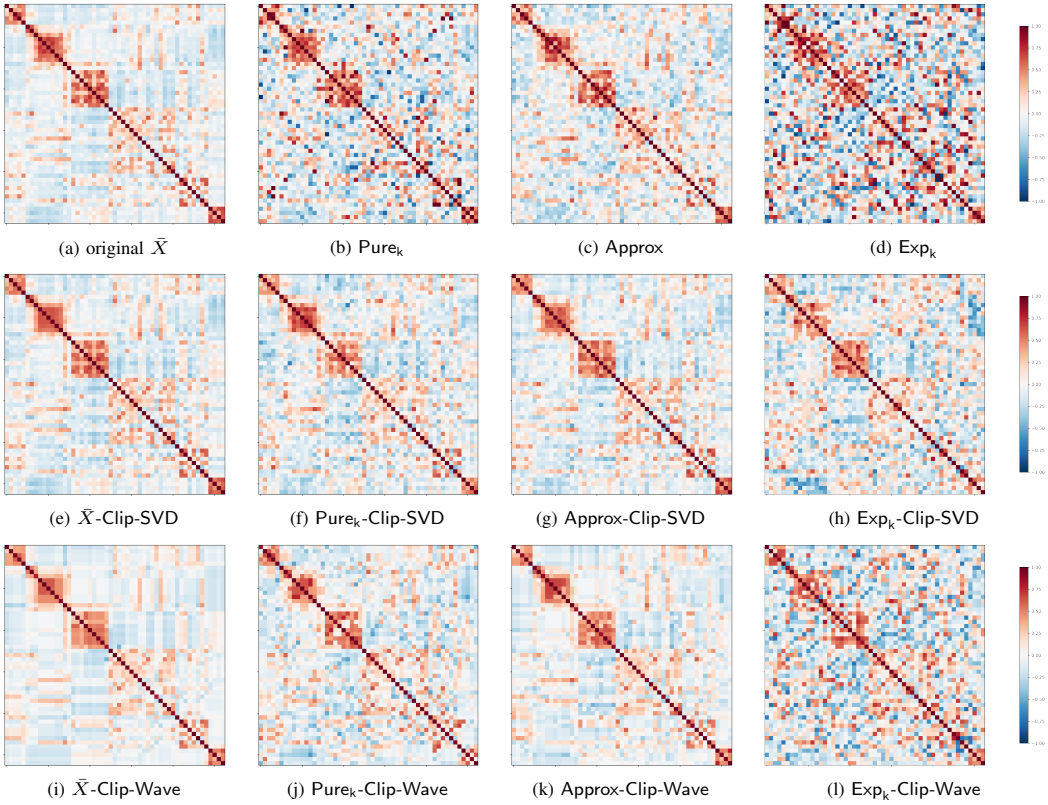

(a) original $\bar{X}$ (b) Pure_k (c) Approx (d) Exp_k

(e) $\bar{X}$-$\mathsf{Clip\text{-}SVD}$ (f) Pure_k-$\mathsf{Clip\text{-}SVD}$ (g) $\mathsf{Approx\text{-}Clip\text{-}SVD}$ (h) Exp_k-$\mathsf{Clip\text{-}SVD}$

(i) $\bar{X}$-$\mathsf{Clip\text{-}Wave}$ (j) Pure_k-$\mathsf{Clip\text{-}Wave}$ (k) $\mathsf{Approx\text{-}Clip\text{-}Wave}$ (l) Exp_k-$\mathsf{Clip\text{-}Wave}$

Fig. A.3. Visualization of the ADNI dataset with $N = 1091, \epsilon = 2, \delta = N^{-3/2}$, clipping bound $u = 0.75$, number of components used in SVD is 10, threshold used in wavelet transform is 0.2, and the number of colors is 256.