

SoK: Privacy-Preserving Data Synthesis

Yuzheng Hu*¹ Fan Wu*¹ Qinbin Li² Yunhui Long¹ Gonzalo Munilla Garrido³
Chang Ge⁴ Bolin Ding⁵ David Forsyth¹ Bo Li¹ Dawn Song²
¹University of Illinois Urbana-Champaign ²UC Berkeley
³Technische Universität München ⁴University of Minnesota ⁵Alibaba Group

Abstract—As the prevalence of data analysis grows, safeguarding data privacy has become a paramount concern. Consequently, there has been an upsurge in the development of mechanisms aimed at privacy-preserving data analyses. However, these approaches are task-specific; designing algorithms for new tasks is a cumbersome process. As an alternative, one can create synthetic data that is (ideally) devoid of private information. This paper focuses on privacy-preserving data synthesis (PPDS) by providing a comprehensive overview, analysis, and discussion of the field. Specifically, we put forth a *master recipe* that unifies two prominent strands of research in PPDS: *statistical methods* and *deep learning (DL)-based methods*. Under the master recipe, we further dissect the statistical methods into choices of modeling and representation, and investigate the DL-based methods by different generative modeling principles. To consolidate our findings, we provide comprehensive reference tables, distill key takeaways, and identify open problems in the existing literature. In doing so, we aim to answer the following questions: What are the *design principles* behind different PPDS methods? How can we categorize these methods, and what are the *advantages* and *disadvantages* associated with each category? Can we provide *guidelines* for method selection in different real-world scenarios? We proceed to benchmark several prominent DL-based methods on the task of private image synthesis and conclude that DP-MERF is an all-purpose approach. Finally, upon systematizing the work over the past decade, we identify future directions and call for actions from researchers.

1. Introduction

Recent years have witnessed an array of groundbreaking scientific achievements across various sectors, including but not limited to medical analyses [1, 2], home automation and robots [3, 4], face recognition [5, 6], and malware detection systems [7, 8]. These significant strides in practical applications can be attributed to the evolution of hardware technology, refinement of algorithms, and most importantly—the availability of large-scale datasets for learning tasks or analytic workflows.

While data serves as the critical fuel, the analysis of such large-scale datasets can often lead to privacy breaches. Examples range from exposing a driver’s working hours and location through the data collected for training autonomous

vehicles [9, 10], to the reconstruction of data used in training machine learning classifiers [11, 12], and even the potential exposure of personal identities from medical images or records [13, 14]. The consequences are severe: the average cost of a healthcare data breach was of \$9.2M [15].

In response to the escalating pace of data collection and analysis as well as the underlying privacy risks, regulatory bodies worldwide have enacted numerous legislations. Notably, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 introduced the privacy rules for disclosing medical records [16], and the Federal Rules of Civil Procedure guide the disclosure of court records [17]. These legislations have exerted considerable impact on various industries. For instance, non-compliance with HIPAA has resulted in substantial financial penalties, with fines reaching as high as \$1.25M [18] and \$6.9M [19]. Amidst the convergence of regulations, penalties, cybersecurity threats, and heightened user awareness [20], the academic community has embarked on various research endeavors aimed at facilitating data analysis in a *privacy-preserving* manner. These investigations encompass a range of paradigms, such as interactive and non-interactive query processing [21, 22], data mining [23–25], machine learning [26–30], among others.

However, creating privacy-preserving solutions tailored to each specific data analysis task is a formidable challenge that necessitates extensive customization. An alternative involves synthesizing a dataset devoid of private information that can be used for *any* downstream task, enabling “*exploratory, open-ended, and iterative*” [31] analyses. In contrast to the aforementioned paradigms, synthetic data permits the execution of a variety of downstream tasks without requiring modifications to the original algorithms to meet privacy constraints. This offers flexibility in terms of future task options, whilst mitigating concerns over privacy leakage.

In this paper, we focus on such privacy-preserving data synthesis (PPDS) approaches. It’s important to note that the uninformed application of privacy mechanisms to the original data could result in significant utility loss [32]. To this end, PPDS mechanisms require delicate design choices to maintain the privacy-utility trade-off.

Motivation. Research of PPDS in the past decade has demonstrated a shift of focus compared to earlier works due to the following *emerging* challenges:

- **High-Dimensional and Unstructured Data.** Classical PPDS approaches designed for tabular data [33, 34] encounter obstacles in coping with the rising dimensionality

1. * indicates equal contribution. Correspondence to: Yuzheng Hu <yh46@illinois.edu> and Fan Wu <fanw6@illinois.edu>.

of data and the inherent complexity involved in handling unstructured data such as images and text.

- **Advanced Attacks.** Traditional anonymization-based PPDS approaches which adopt either the syntactic privacy [35] or less rigorous privacy principles are known to suffer from various privacy attacks [36–40].

To alleviate concerns associated with the data-related challenges, statistical models (*e.g.*, hidden Markov model (HMM) [41] and Bayesian network [42]) and tools (*e.g.*, copula [43]) have been used to facilitate the modeling and deconstruction of high-dimensional structured data, while deep generative models [44–46] arise as powerful tools to handle unstructured data. To remediate concerns on the privacy front, research is focused on developing rigorous privacy notions (*e.g.*, differential privacy (DP) [47]) with corresponding privacy mechanisms, as well as their adaptations and applications in different scenarios. With abundant work exploring the intersection of these two lines of research at hand, we believe it is time for us to take a step back and systematically review the established toolkits and the lessons learned.

Standing Point and Contributions. Our work joins a small family of surveys (see “Related Work” below) that study PPDS. While we are not the first to review this topic, we position our work as **the first to systematize the research effort over the past decade which tackles the emerging challenges using the most up-to-date toolkits**. Specifically, we make the following contributions:

- We present a comprehensive overview of PPDS (Section 2), considering different privacy notions and mechanisms, input data types as well as utility evaluation protocols.
- We distinguish two primary research directions in PPDS: statistical methods and deep learning (DL)-based methods. Furthermore, we present a master recipe (Fig. 1) that unifies these approaches, encompassing the algorithms falling within each category.
- Under the master recipe, we dissect the statistical methods (Section 3) into choices of modeling and representation, and investigate the DL-based methods (Section 4) by different generative modeling principles. To consolidate our findings, we provide comprehensive reference tables, distill key takeaways, and identify open problems in the existing literature.
- We benchmark the DL-based methods for private image synthesis (Section 5), and identify DP-MERF [48] as the best practice. We further note that the pessimistic conclusion on synthetic tabular data by Stadler *et al.* [49] does not fully extend to the image domain.
- We delineate future directions for researchers (Section 6).

Related Work. Prior surveys related to PPDS can be classified into two categories: (i) privacy-preserving data publishing (PPDP) [40, 50–53] which focuses on the anonymization of structured data, and (ii) privacy-preserving deep generative models [54, 55]. Apart from these works, we also identify a few benchmarks [56, 57] for some representative algorithms on private tabular data synthesis, as well as a general introduction to synthetic data [58].

Scope. In terms of the *timeline*, we collect papers after 2010 which allows us to concentrate on DP, and also distance ourselves from a prior survey [53]. Our *ontology* is as follows: Section 3 primarily concerns the intersection of database and privacy, whereas Section 4 concerns the intersection of DL and privacy. In terms of the *target*, we exclude works that *only* facilitate learning tasks or analytic workflows [59–62], or those that publish aggregated or summary statistics [63, 64] rather than generating synthetic data.

2. Preliminaries

Given an input dataset D , we aim to design a PPDS algorithm P , such that $P(D) \rightarrow D'$ with D' being the synthesized dataset which is both *private* and *utilitarian*. Formally, P satisfies a specific **privacy** definition (Section 2.1); and D' attains high **utility** (Section 2.2).

2.1. Privacy Notions and Mechanisms

There are numerous definitions of privacy spanning different applications and time-periods [65–68]. In the context of PPDS, we adopt the perspective of F. T. Wu [67], which posits that privacy remains intact if no violation, in the form of *re-identification*, occurs. We deal with a broad sense of re-identification beyond human identities [69–72] and consider the exposure of any sensitive information within data as re-identification. This could range from unveiling connections among users in social networks [73], to revealing sensitive facial attributes, such as the eyes, nose, or mouth, in facial image data [74].

We focus on *differential privacy* (DP), which provides mathematical guarantees against re-identification. Privacy definitions such as syntactic anonymity [35] (*e.g.*, k -anonymity [75] and its variants) and less rigorous privacy principles such as obfuscation [76] were widely adopted in PPDS [40, 53]. However, recently, there is a clear tendency to prefer DP [77–79] because of (i) its ability to quantify privacy leakage, and (ii) the stronger protection it offers [80]. **Differential Privacy (DP).** Along with a variety of benign properties including composability [81], immune to post-processing [82] and future-proof [83], DP has become the gold standard of privacy. DP quantifies and limits the amount of knowledge one could learn by observing the algorithm’s output. Formally:

Definition 1 ((ϵ, δ) -DP [82]). Let \mathcal{X} be the space of the dataset, a randomized algorithm \mathcal{M} with domain $\mathbb{N}^{|\mathcal{X}|}$ is (ϵ, δ) -differentially private iff for all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$ and for any neighboring datasets D and D' differing by one record: $\Pr[\mathcal{M}(D) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{M}(D') \in \mathcal{S}] + \delta$.

The standard definition assumes a centralized database with trusted server and independence between participants [35]. Rényi DP (RDP) [84] is a variant of DP that uses Rényi divergence [85] to measure the distance between two probability distributions. Other variants include edge-DP and node-DP [86] for graphs (see Section 3.3).

Privacy Mechanisms. For algorithms with numeric output, DP and its variants are typically enforced by adding noise (Laplace and Gaussian mechanism [82]). For non-numeric

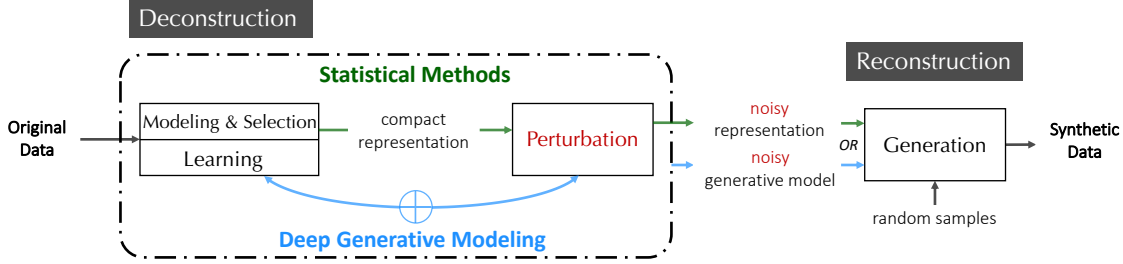


Figure 1: **The master recipe.** In the deconstruction stage, statistical methods obtain a compact representation through *modeling & selection*, which is then subject to perturbations. On the other hand, DL-based methods typically integrate *learning* with perturbation. The resulting noisy representation or generative model is then used to generate synthetic data in the reconstruction stage.

purpose, conventional wisdom is to perform probabilistic selection based on a score function (exponential mechanism [87]). For algorithms that involve iterative optimization of a parametrized model, DP-SGD [88] adds Gaussian noise to the clipped gradient. Moments Accountant [88] and RDP accountant [84] track the privacy cost throughout iterative optimization.

2.2. Data Types and Utility

PPDS has engaged extensively with both *structured* (S) and *unstructured* (U) data, each necessitating distinct approaches. Statistical methods have been devised primarily for structured data, whereas DL-based methods are more suitable for unstructured data. We focus on the following data types: (S1) *tabular data*, which arrange elements in columns and rows with each column corresponding to an attribute and each row representing a record; (S2) *trajectory data*, which comprise sequences of locations such as (latitude, longitude) pairs and reflect the movement of individual over time; (S3) *graph data*, which consist of nodes connected by edges where the nodes denote certain entities and the edges represent their interconnections; (U1) *image data*, which are comprised of pixels.

Data utility is characterized by a metric that depends on two primary factors: (i) the type of data, and (ii) the nature of the data analysis conducted. We take tabular data as a representative example. When the analysis task involves the construction of a histogram with the data, utility can be assessed by the similarity (e.g., ℓ_∞ -norm) w.r.t. the outcome of queries, such as random linear queries [89], range queries [90–92], parity queries [93], and aggregation queries [94]. Conversely, when the analysis task entails training a machine learning model with the data, utility can be evaluated by the model’s performance on downstream tasks, for instance, classification accuracy [48, 95, 96]. Due to space constraints, we defer a detailed summary of utility metrics to our GitHub repository².

Another concept that bears a close relationship to data utility is *fidelity*. It refers to the extent to which inherent data properties (e.g., statistical or structural) are preserved, and is frequently applied in the literature to evaluate synthetic

data. In assessing fidelity, similarity functions are often used for structured data, while pre-trained neural networks are employed for unstructured data. We will delve further into the correlation between utility and fidelity in Section 5.

3. PPDS via Statistical Methods

Overview. In this section, we introduce the statistical methods for PPDS. Unlike DL-based methods, which are capable of performing end-to-end synthesis, statistical methods are typically tailored to specific data types. Consequently, it is generally impractical to expect a single algorithm to function effectively across various scenarios without adjustments. To that end, we structure this section based on different data types. Our analysis encompasses three categories of structured data (S1, S2, S3) as outlined in Sections 3.1 to 3.3, respectively, with Table 1 providing a comprehensive overview of the literature from multiple perspectives. We exclude unstructured data from our discussion due to the lack of sufficient references.

Master Recipe. Despite the variations in tools and implementation details, the majority of works discussed in this section can be encapsulated within a unified framework (refer to Figure 1), which consists of two primary components: deconstruction and reconstruction. We refer to this as the *master recipe* in the subsequent discussion. In the context of statistical methods, the deconstruction component involves modeling the original data and identifying a compact representation, potentially safeguarded by privacy measures. The subsequent noisy representation, derived through the application of Laplace or Gaussian mechanisms, is then passed to the reconstruction component for synthetic data generation. It is worth noting that the deconstruction step can effectively tackle the challenge posed by high-dimensionality by reducing the amount of noise required to ensure DP.

Most naturally, we unfold our discussion by examining how various approaches can be interpreted as instantiations of the master recipe, with a focus on the modeling of the original data as well as the choice of compact representation.

3.1. PPDS for Tabular Data

We begin with tabular data, a prevalent data type in practical applications. Due to its highly-structured nature,

2. https://sok-ppds.github.io/data_utility_and_fidelity.html

TABLE 1: **Statistical methods for PPDS.** Generality: ●=general, ○=specific. Scalability: ●=good, ◐=moderate, ○=poor. ✓=with, ✗=without. / =not applied.

Data		Methodology		Privacy			Properties		Evaluation		References
Data Type	Data Characteristics	Approach/Model	Tool	Privacy Guarantee	Privacy Mechanism	Privacy Test	Utility Metrics	Utility Guarantee	Generality	Scalability	
Tabular Data	categorical attribute	marginal-based	Bayesian network	DP	LM, (EM)	✓	S1-F1, F5, A1	✗	●	◐	[97, 98]
			Markov random field	DP	LM / GM	✗	S1-F1, F2, A1	✗	●	◐	[99–102]
	numerical attribute	query-based	/	DP	EM, (LM)	✗	S1-F1, F2	✓	○	○	[90, 93] [103, 104]
			Gaussian copula	DP	LM	✗	S1-F1, F2	✗	●	○	[91, 105]
		projection-based	vine copula	DP	LM	✓	S1-F1, F3, F6, A1, A2	✗	●	○	[106]
			random projection	attribute-DP	GM	✗	S1-F2, F4, A1	✓	○	●	[89, 107]
Trajectory Data	sequential data	tree-based	/	DP	LM	✓	S2-F1~F7	✗	●	○	[111–115]
			distribution-based	/	DP	LM	✓	S2-F1, F2, F4~F7	✗	●	◐
	/	SKG	/	edge-DP	LM	✗	S3-F1, F4, F8, F9	✗	●	○	[121]
	Graph Data	/	dK-Graph	/	edge-DP	LM	✗	S3-F1~F4, F6, F8, F9, A1, A2	✗	●	◐
/		ERGM	/	edge-DP	LM	✗	S3-F7, F8	✗	●	○	[124]
attributed graph		AGM	/	edge-DP	EM, LM	✗	S3-F1, F3, F9	✗	●	◐	[125]
				edge-DP	LM	✗	S3-F1, F4, F5, F8	✗	●	◐	[126, 127]

numerous approaches target tabular data from distinct perspectives while adhering to the same master recipe.

3.1.1. Paper Summaries. We identify four approaches of DP tabular data synthesis based on marginal, copula, projection and query. At a high level, marginal-based and copula-based approach model the rows of tabular data as *realizations of a joint distribution*; projection-based approach models tabular data holistically as a *matrix*; query-based approach belongs to a line of theoretical research and does not perform any modeling.

Marginal-Based Approach. The low-order marginals are the go-to choice of statistics for tabular data. They are adept at capturing the low-dimensional structure of real-world distributions and have low sensitivity, making them an ideal candidate for the compact representation in Fig. 1. In line with the master recipe, algorithms fall under the marginal-based approach roughly follow the same routine: select marginals, measure with noise, and generate synthetic data from the joint distribution estimated by the noisy marginals. To facilitate marginal selection and/or generation, the procedure is implemented within graphical models such as Bayesian network [42] or Markov random field [128].

Bayesian Network. Zhang *et al.* [97] propose PrivBayes, which simultaneously selects marginals and determines the topology of the Bayesian network by iteratively maximizing information gain through the exponential mechanism. Synthetic data can subsequently be sampled from the joint distribution, which is computed based on the Bayesian network topology. Bindschaedler *et al.* [98] adopt a similar procedure as PrivBayes, albeit applied to only a random subset of attributes for *partial* data synthesis. Additionally, the authors propose a privacy test for the release of synthesized records, which leads to the guarantee of plausible deniability [129].

Markov Random Field. The Markov random field (MRF) is capable of capturing more complex correlations within

the data compared to a Bayesian network. Chen *et al.* [99] construct a graph that explore pairwise dependence between attributes and apply the junction tree algorithm to obtain the MRF, from which the noisy marginals are generated and the synthetic data are sampled. The probabilistic graphical-model (PGM) based estimation [130] does not fulfill PPDS per se, but is rather a reconstruction algorithm that can effectively learn a MRF that best matches the noisy marginals via mirror descent [131]. As a matter of fact, PGM serves as a crucial building-block in the *state-of-the-art* mechanisms of tabular data synthesis: it gives rise to the top-ranked teams in the NIST 2018 [78] and 2020 [79] contests [100, 102] when coupled with sophisticated techniques of marginal selection. More recently, Graduate Update Method (GUM) [101] arises as an alternative reconstruction algorithm to handle graphical models with dense interactions.

The marginal-based approach explicitly models the correlations between different attributes within graphical models, which makes the synthetic data capable of handling a wide range of tasks with satisfactory utility.

Copula-Based Approach. A copula, as defined within the realms of probability theory and statistics, refers to the joint cumulative distribution function which describes the *dependence* among random variables. Once integrated with univariate marginal distributions, it is capable of determining the joint distribution (see Sklar’s Theorem [132]), which forms the foundation for sampling synthetic data. Therefore, a combination of the univariate marginal distributions and a copula would serve as the compact representation in Fig. 1, and the key question is to determine the copula (the marginal distributions can be estimated from the histograms of the original data). We will discuss two classes of copula respectively, namely, Gaussian copula [133] and vine copula [134].

Gaussian Copula. Gaussian copula enjoys a closed-form and can be uniquely determined by the correlation matrix. DPCopula [91] considers two methods: maximum likelihood

and Kendall’s τ coefficient [135] to estimate the correlation matrix, perturb with noise and derive the Gaussian copula thereof. To overcome the limited applicability of copula on discrete attributes, Asghar *et al.* [105] propose an additional step of pre-processing which can be applied to both categorical and numerical attributes before the estimation.

Vine Copula. Despite the simplicity and convenience of the Gaussian copula, it has been noted that the tail dependencies might not be fully captured by this model, and can lead to a significant loss of information [136]. Therefore, COPULA-SHIRLEY [106] adopts *vine copula*, which decomposes the multivariate function into multiple bivariate copulas, and applies a top-down greedy algorithm from [137] for copula selection. While being less computational efficient, the customizable framework as well as the ability of vine copula in modeling complex distributions lead to improved applicability and utility.

Similar to the marginal-based approach, copulas can adapt to the inherent structure of the data, enabling algorithms to attain relatively high utility. But copula-based approach comes with the downside of high complexity—quadratic in the number of attributes in the case of Gaussian copula selection and even higher for vine copula. This renders them less favorable compared to the marginal-based approach.

Projection-Based Approach. Under the projection-based approach, tabular data is plainly modeled as a matrix, and the projections onto low-dimensional subspace naturally serve as the compact representation in Fig. 1. Furthermore, the selection of these projections typically leverages certain geometric properties inherent to the high-dimensional space.

Random Projection. There are a few works [89, 107] that *partially* align with the master recipe due to the absence of a reconstruction process. Specifically, the synthetic data are obtained via projecting the original data by row using a Gaussian random matrix followed by some perturbation. The pairwise L_2 -distances between individuals are preserved by the Johnson-Lindenstrauss lemma [138], making the synthetic data suitable for tasks such as clustering and nearest neighbors which only rely on such information.

Others. Apart from random projection, other algorithms that fall under this category are consistent with the master recipe. Li *et al.* [108] obtain a sparse representation of the original data using a Bernoulli random matrix, and reconstruct the synthetic data from the perturbed representation via compressive sensing [139]. Jiang *et al.* [109] adopt the noisy top principal components to construct the synthetic data, which are computed based on the perturbed first and second moments. Chanyaswad *et al.* [110] observe that the *orthonormal* projection of high-dimensional data are nearly Gaussian [140], and apply a Gaussian generative model on top of the noisy projection to produce the synthetic data.

On the positive side, algorithms based on projection are easy to implement and usually come with some utility *guarantee*: they aim at preserving specific properties of the original data, and the corresponding utility bounds arise from the similarity in statistical measures. However, they usually fail to capture the underlying structure of the original data, mainly due to the coarse selection of the compact

representation which leads to a significant loss of information. As a consequence, the synthetic data generated from these algorithms are incapable of accommodating generic usage.

Query-Based Approach. The query-based approach [90, 93, 103, 104] is built upon a series of prior work that lay down the theoretical foundation of differentially private query processing [34, 141], with the goal of privately releasing a synthetic dataset that can accurately answer an *assigned* workload of linear queries instead of making it generically useful. While queries are a generalization of marginals and can be similarly interpreted as the compact representation in Fig. 1, there is no explicit modeling of the original data, and the reconstruction step deviates from the master recipe. As such, we will not discuss these algorithms in detail.

Takeaways. *Trade-off between specificity and generality—Algorithms tailored to specific data usage, such as random projection, typically offer provable utility guarantees but lack versatility across different tasks. Conversely, marginal-based and copula-based approaches, while being heuristic in utility, excel at capturing attribute correlations. Moreover, their joint-probability modeling enables the generation of an infinite amount of data.*

Open Problems.

- **Fine-grained structures.** *State-of-the-art algorithms for differentially private tabular data synthesis primarily rely on low-order marginals, which may not adequately capture the high-order correlations between attributes (e.g., denial constraints [142] and (conditional) functional dependencies [143, 144]). Integrating these fine-grained structures into the design of algorithms holds the potential to enhance data utility, a prominent example being Kamino [145].*
- **Mixed attributes.** *In real-world scenarios, datasets commonly consist of a combination of categorical and numerical attributes. However, most existing algorithms are specialized for handling only one type of attribute (marginal-based approaches for categorical attributes and copula-based approaches for numerical attributes). Consequently, there is a pressing need for a thoughtful pre-processing design to generate tabular data with mixed attributes.*

3.2. PPDS for Trajectory Data

While trajectory data typically take a tabular form, their ubiquitous collection in real-world scenarios [146, 147] has fostered the development of privacy-preserving trajectory data synthesis as a distinct subject of interest, which attracts substantial research efforts from the academic community. Compared to tabular data, trajectory data 1) embody *temporal dependencies*, akin to general sequential data; 2) incorporate *key features* such as start and end points, trajectory length, and so forth. As such, in addition to discussing different approaches under the master recipe, we will also underscore how these features are integrated into the algorithmic design.

3.2.1. Paper Summaries. Following [40], we discuss two approaches of DP trajectory data synthesis: the tree-based approach and the distribution-based approach. At a high level, these two approaches diverge in their modeling of

trajectory data. The former primarily addresses the temporal dependency, viewing trajectory data broadly as a sequence. In contrast, the latter interprets trajectories as the combination of certain key features.

Tree-Based Approach. The tree-based approach attempts to create a noisy *prefix tree* based on the original data, from which the synthetic trajectory data will be derived. The prefix tree—a hierarchical structure that groups trajectories with the same prefix into the same branch, can therefore be interpreted as the compact representation in Fig. 1 and encodes the temporal dependency of the trajectory data.

General Sequential Data Synthesis. A few works on differentially private trajectory data synthesis are designed for general sequential data. Chen *et al.* [111] apply multiple levels of location generalization to the original prefix tree and add Laplacian noise to each generalized node, whose expansion terminates when the noisy count fails a certain threshold or a predefined tree height is reached. A subsequent work [112] constructs the noisy prefix tree based on the n -gram model, which makes use of the Markov independence assumption to estimate the probability of a current location based on the previous $(n-1)$ ones. Compared to [111] whose majority of nodes have small counts due to the uniqueness of patterns, the n -gram model captures more temporal dependency and is therefore more robust to noise. Building upon the Markov model, Zhang *et al.* [113] employ a *prediction suffix tree* [148] which bears resemblance to the prefix tree explored in [111]. Nonetheless, their algorithm brings several improvements to the table: it permits adaptability in the tree’s height and incorporates a more sophisticated strategy for node expansion, which considers the diversity of potential future child nodes. Both of these adaptations contribute to enhanced utility.

Fine-grained Trajectory Synthesis. The aforementioned algorithms only work with coarse trajectories on small location domain. To produce fine-grained trajectories over real-world geographical span, follow-up works seek to leverage additional properties or features of trajectory data beyond temporal dependency. He *et al.* [114] propose DPT, which takes into account the *non-uniformity* of real trajectories: the algorithm discretizes the trajectories using multiple location sampling resolutions to capture individual movements at different speeds, and maintains a prefix tree for each resolution. Wang and Sinnott [115] establish the noisy prefix tree based on a private reference system. This construction calibrates original trajectories against a selection of anchor points, adeptly reflecting both overarching mobility patterns and detailed attributes such as turning and congestion points.

The tree-based approach compactly summarizes the temporal dependency between consecutive locations and works more generally for sequential data. However, algorithms that primarily rely on the tree structure usually fail to preserve useful properties such as the trajectory length. Furthermore, they typically exhibit limited scalability. For instance, the discretization of a vast location domain elevates both the size of the prefix tree and the necessary noise [111]. Correspondingly, algorithms that depend on the Markov model [112, 113] exhibit increasing inefficiencies as the

trajectory length expands.

Distribution-Based Approach. Distribution-based approach first derives the probability distributions associated with key attributes of the original data, samples from the perturbed distributions, and generates synthetic trajectories thereof. The probability distributions can be viewed as the compact representation in Fig. 1.

Mir *et al.* [116] put forth DP-WHERE, an approach that estimates various distributions of an attribute set, capable of comprehensively determining human movement trajectories. Roy *et al.* [117], on the other hand, consider the interactions between different attributes by grouping strongly correlated attributes into non-disjoint sets and constructing a corresponding distribution for each set. More recently, Gursoy *et al.* (2018) introduce DP-Star, an approach that deconstructs trajectory data into four principal features: space, start and end points, transition probabilities, and route length. Each feature is associated with a probability distribution that is perturbed with noise. Synthetic trajectories are then reconstructed from the features sampled from the noisy distributions. Two subsequent studies, AdaTrace [119] and OptaTrace [120], adopt similar pipelines as [118] while offering respective enhancements in terms of privacy and utility. Specifically, AdaTrace enforces deterministic constraints on the generated trajectories while maintaining DP, thereby bolstering the algorithm’s resilience against common threats in trajectory data (*e.g.*, Bayesian inference). OptaTrace employs Bayesian optimization to identify the optimal allocation of the privacy budget across the four features, aiming to minimize the output error relative to a specific metric (*e.g.*, query error).

Distribution-based approach leverages domain-specific features, equipping algorithms under this category to generate trajectories with high fidelity. One minor disadvantage of the distribution-based approach is that it can *occasionally* be outperformed by the tree-based approach [112, 114] under stringent privacy settings (*e.g.*, $\epsilon = 0.1$). This occurs because the total privacy budget is divided into multiple shares, and the required noise escalates super-linearly with the decreasing budget. As such, each component incurs substantial error when ϵ is small.

Takeaways. *The power of domain information*—Through a comparative analysis of the tree-based and distribution-based approaches, as well as the two subcategories within the tree-based approach, we observe that the inclusion of fine-grained properties (*e.g.*, non-uniformity) or features (*e.g.*, anchor points) consistently enhances utility and applicability. This serves as compelling evidence for the influential role of domain information in algorithmic design.

Open Problems.

- **Expanding formats.** *Real-world trajectory data often encompasses supplementary columns that extend beyond spatiotemporal information, such as user attributes (*e.g.*, occupation). The increasing prevalence of drones and airborne sensors will push data synthesis from 2D planes to 3D spaces. These diverse data formats present researchers with novel challenges.*
- **Task-centric utility metrics.** *Currently, the evaluation of*

synthetic trajectories is statistical in nature: the utility metrics primarily focus on fidelity but fail to capture the overall usefulness of the generated data. We envision the establishment of task-centric utility metrics as a vital step towards achieving equitable quality assessment across different algorithms.

3.3. PPDS for Graph Data

Graphs surface in numerous real-life scenarios such as social networks and communication data [77]. Two variants of DP are proposed for this data structure: *edge-DP* and *node-DP* [86], and we focus on the former since node-DP is scarce in the literature.

3.3.1. Paper Summaries. In line with the master recipe, most algorithms for DP graph data synthesis share the same pipeline: extract graph statistics, which serve as the compact representation in Fig. 1, and generate the synthetic graph based on the perturbed statistics. Additionally, the graph statistics are usually computed within a specific graph model, which reflects the *belief* of the underlying generation process. As such, we organize the discussion based on different graph models and omit works that do not adopt any [149–152].

Stochastic Kronecker Graph Model. The Stochastic Kronecker Graph Model (SKG) [153] is generated through the Kronecker power of an adjacent matrix and reduces graph synthesis to parameter estimation. Mir and Wright [121] incorporate the Laplace mechanism into the moment-based estimation of SKG [154] to ensure edge-DP. While being simple to implement, SKG cannot accurately capture the structural properties of real graph data since the generation process is determined by a *single* parameter.

DK-Graph Model. The dK-graph model [155] condenses the graph into the degree distribution of connected components of size K (dK-series). Notably, this can also serve to reconstruct the graph when $K = 2$, making the dK-2-series a primary research focus. Sala *et al.* [122] observe that the global sensitivity of the dK-2-series is large and propose a partitioning-based algorithm, which divides the dK-2-series into different sub-series and adds *non-uniform* noise based on their local sensitivity [156]. However, it is highlighted in [157] that adding noise based on local sensitivity might inadvertently expose private information. As such, a subsequent work by Wang and Wu [123] calibrates the noise based on the smooth sensitivity [157], thereby providing a rigorous edge-DP guarantee. However, despite these improvements, the privacy budget is still required to be unreasonably large (e.g., $\epsilon \geq 100$).

Exponential Random Graph Model. The Exponential Random Graph Model (ERGM) [158] is characterized by a vector of sufficient statistics, capable of encapsulating a broad spectrum of structural and attribute information. Analogous to KGM, graph synthesis can be reformulated as a problem of parameter estimation in ERGM. Lu and Miklau [124] first privately compute the alternating graph statistics [158], and then estimate the model parameters using Bayesian inference [159]. However, ERGM can only work with a few

thousand nodes and fail to scale to moderately-large graphs due to the heavy computational cost of Bayesian inference.

Hierarchical Random Graph. The hierarchical random graph (HRG) [160] model represents a graph with a dendrogram which reflects its hierarchical structure, and some connection probabilities associated with the internal nodes of the dendrogram. Xiao *et al.* [125] privately sample a dendrogram via Markov-Chain Monte Carlo (MCMC) [161], and propose a thresholding strategy to calculate the noisy connection probabilities. Compared to other works, the HRG model achieves the smallest global sensitivity and therefore requires the least amount of noise to ensure edge-DP. This showcases the power of encoding structural information in the form of edge probabilities.

The models discussed so far are designed for unlabeled graphs and focus on the *graph structure* in isolation. In contrast, real-world graph data such as social networks attach attributes to each vertex (e.g. affiliations on Facebook and Twitter), and their correlations are reflected by the graph structure in the form of *homophily*—the tendency for nodes with similar attributes to form connections. This observation guides us to the subsequent graph model.

Attributed Graph Model. The attributed graph model (AGM) [162] attaches binary attributes (e.g. smoker v.s non-smoker) to nodes and performs conditional sampling of the graph structure based on the correlations between attributes. AGM requires a generative model to capture the structural properties of the input graph, and contains three sets of parameters that characterize the distribution of attributes and edges as well as the correlations between attributes and edges. Jorgensen *et al.* [126] propose a structural model called TriCycLe, which more effectively captures the clustering property of the input graph compared to its predecessors [163, 164]. The authors integrate TriCycLe into AGM and generate the synthetic graph based on privately-estimated parameters. A follow-up work by Chen *et al.* [127] proposes the Community-Preserving Attributed Graph Model (C-AGM), which preserves the community structure while retaining the advantages of the AGM model.

Takeaways. *Absence of a universal solution*—A universally applicable differentially private graph synthesis algorithm has yet to be found and might not even exist. Each method demonstrates distinctive trade-offs related to privacy, utility, and the preservation of specific graph properties. As a result, the selection of an appropriate algorithm is contingent on the demands and constraints inherent to each application.

Open Problems.

- **Leveraging prior knowledge.** The selection of graph models often lacks principled justification (except for AGM, which explicitly encodes attribute information). To enhance model selection in diverse scenarios, a combination of comprehensive understanding of each model and prior knowledge regarding the underlying generation process can prove invaluable.
- **Node-DP.** Existing research has primarily centered around Edge-DP, which safeguards the relationship between users. However, node-DP offers a more robust notion of adjacency,

where neighboring graphs differ in both nodes and associated edges. This introduces additional challenges in terms of selecting and fitting graph models, as it significantly impacts the graph statistics. Addressing this challenge remains an intriguing open problem for future investigation.

4. PPDS via DL-Based Methods

Deep generative models are renowned for their capability and flexibility in handling high-dimensional data and generating diverse, high-fidelity samples [45, 46, 188, 189]. We refer to these strategies as Deep Learning (DL)-based methods, given the prevalent use of modern neural networks (NNs) as their primary computational framework. Although these models afford a degree of empirical privacy protection [190, 191], the substantial capacity of NNs can give rise to the “memorization” phenomenon [192], thus rendering them vulnerable to membership inference attacks [193–195]. Consequently, it is crucial to develop deep generative models with formal privacy guarantees.

Master Recipe. We discuss how privacy-preserving deep generative modeling aligns with the master recipe. In the deconstruction stage, a generative model is learned from the training data. In the reconstruction stage, synthetic data is generated by feeding random samples to the generative model. DP is achieved by incorporating perturbations into the learning procedure.

The instantiations of the master recipe for DL-based methods markedly deviate from those for statistical methods across various dimensions, owing to the inherent differences in their working mechanisms. Both the *modeling* and the *representation selection* in statistical methods are contingent upon the data type. In contrast, the *generative modeling principle* plays a central role in DL-based methods. Once the principle is determined and the model architecture is chosen (e.g., deep convolutional GANs for image data), the representation selection is automatically accomplished within the learning process. Therefore, privacy enforcement in DL-based methods is marginally associated with the data type, but closely related to the generative modeling principle. Additionally, the iterative nature of the optimization algorithms employed in DL makes DP-SGD [88] the default privacy mechanism for DL-based methods.

Overview. We organize the section based on different principles of generative modeling, including the auto-encoding mechanism, minimax training, optimal transport, feature alignment, and stochastic simulation (Sections 4.1 to 4.5). Notably, some algorithms leverage multiple principles to perform PPDS. In this case, we categorize them based on the primary principle driving their design. Table 2 provides a comprehensive overview of the literature.

4.1. PPDS Based on Auto-Encoding Mechanism

4.1.1. Preliminaries. Auto-encoders (AEs) [196] are designed to distill the data into a low-dimensional representation known as the *latent space*, which captures the core features of the data. The generation of new samples can be accomplished

by sampling points from this latent space and subsequently decoding them back into the original data space. The latent space plays a central role in the design of AE-based approaches; lacking structure can lead to the creation of unrealistic samples.

4.1.2. Paper Summaries. We review how different works impose structure on the latent space.

Gaussian Latent Space. Variational Autoencoders (VAEs) [46] impose a Gaussian prior on the latent space and optimize the parameters to minimize a combined loss, comprising the reconstruction error which measures the dissimilarity between the input and reconstructed data, as well as the KL divergence that quantifies the disparity between the prior and posterior distributions of the latent variables. This ensures a continuous and smooth structure in the latent space. One can apply DP-SGD to the end-to-end training of the VAE to achieve DP [165, 166].

Gaussian Mixture Latent Space. Imposing a Gaussian prior on the latent space of VAEs can restrict their expressiveness. To address this issue, several studies model the latent space with more flexible distributions. DP-SYN [167] trains an AE using DP-SGD, and then models its latent space as a Gaussian Mixture Model (GMM) via DP Expectation Maximization (DP-EM) [197].

GAN-Learned Latent Space. Rather than imposing a prior on the latent distribution, an automated approach uses GANs (see Section 4.2) to learn the latent distribution of the trained AE [168–171]. One can sample from the GAN to obtain latent representations of the AE, and then decode them to generate the synthetic data. DP-SGD is applied to the training of both the AE and the GAN. This line of research aims to harness the strengths of both GANs and AEs to effectively generate *discrete* data, such as tabular data with categorical attribute. GANs facilitate the sampling from the AE’s latent representation, and AEs permit the learning of discrete data, a task that cannot be accomplished solely by GANs.

Takeaways. Different types of latent spaces entail unique trade-offs. Gaussian latent spaces excel in structural simplicity and smoothness, but struggle in modeling complex data. Enhancing the expressiveness and flexibility of the latent space comes at the cost of increased parametrization, leading to additional privacy cost. Therefore, the selection of latent space hinges on the specific use case.

Open Problems. AEs, designed for efficient data representation, implicitly capture important features of the data distribution in the latent space. This characteristic is particularly suitable for PPDS, in which the utility of the synthesized data often outweighs their fidelity (i.e., similarity in data distribution). To amplify this advantage, one promising approach could be to explicitly enhance the utility of the latent space. For instance, beyond the conventional reconstruction task, one could introduce an additional classification loss within the training objective. This strategy encourages the latent representation to acquire useful features for classification tasks, which will subsequently be embodied in the synthesized data.

TABLE 2: DL-based methods for PPDS. Generality: ●=general, ○=specific. Scalability: ●=good, ◐=moderate, ○=poor. Privacy Test: ✓=with, ✗=without.

Generative Modeling Principles	Privacy		Instantiations		Properties		Evaluation		References
	Privacy Mechanism	Privacy Accountant	Data Type	Model Type	Generality	Scalability	Utility Evaluation	Privacy Test	
Auto-Encoding Mechanism	DP-SGD (+DP- <i>k</i> means)	MA	tabular, images	ensemble+VAE	●	◐	S2-F1, U1-A1	✗	[165, 166]
	DP-SGD+DP-EM	MA	tabular	ensemble+AE	●	○	S1-F1, S1-A1	✗	[167]
	DP-SGD	MA	time series	GAN+AE	●	○	S1-A1, S1-A3	✗	[168]
		RDP	tabular, network traffic	GAN+AE	●	○	S1-F1, S1-A1	✗	[169–171]
Minimax Training	DP-SGD	MA	tabular, network traffic	GAN, WGAN, improved WGAN	●	○	S1-F1, F3, F4, A1, A2, A3	✓	[94, 171–175]
				AC-GAN	●	○	S1-A1	✗	[176]
				CTGAN	○	◐	S1-A1	✗	[177, 178]
	RDP	RDP	images	WGAN, improved WGAN	●	○	U1-F1, F4, A1	✓	[174, 179]
				CGAN	●	◐	U1-A1, A3	✗	[95]
	Partial Gradient Sanitization	RDP	images	improved WGAN	●	◐	U1-F1, F2, A1	✗	[96]
	Private Knowledge Aggregation and Transfer	MA	tabular	teacher-student+WGAN	●	○	S1-A1, A2	✗	[180]
RDP		tabular, images	teacher-student+DCGAN	●	◐	S1-A1, A2, U1-F1, F2, A1	✗	[181]	
		images	teacher-student+DCGAN	●	●	U1-A1	✗	[182]	
Optimal Transport	Partial Gradient Sanitization	RDP	images	generator	○	●	U1-F2, U1-A1	✗	[183]
Feature Alignment	GM	RDP	tabular, images	generator	●	◐	S1-A1, U1-A1, U1-F2, U1-F3	✗	[48, 184, 185]
Stochastic Simulation	RR	N/A	images	SGM	○	●	U1-F1, F2, A1	✗	[186]
	DP-SGD	RDP	images	Score SDE	○	●	U1-F2, U1-A1	✗	[187]

4.2. PPDS Based on Minimax Training

4.2.1. Preliminaries. Minimax training involves training two components in opposition to each other in a zero-sum game. A prominent example in the context of generative modeling is the Generative Adversarial Network (GAN) [45]. GANs consist of a generator network that produces samples resembling real data, and a discriminator network that distinguishes between the generated and real data.

In contrast to other categories of generative modeling, a distinctive characteristic of GAN is that it consists of two interdependent models that are alternatively updated during the training procedure. The escalation in system complexity warrants a reflection on the threat model, specifically, what information the adversary could access in order to launch privacy attacks. This is critical in designing DP algorithms, as any information accessible to the adversary, and only that information, should be protected. The strongest adversary can access the entire information flow (*e.g.*, outputs, gradients) throughout training. A slightly weaker adversary can access only the generator; this is a reasonable assumption given that only the generator is typically released for public use. The weakest adversary sees only the synthetic data.

4.2.2. Paper Summaries. We review two strategies of achieving DP, corresponding to the first two types of adversaries. The discussion of the model-agnostic weakest adversary is deferred to Section 6.

Perturbing the Discriminator. Existing literature primarily focuses on the strongest adversary. We start with a preliminary method that is implemented in several parallel studies [172–174, 176, 179], which employs DP-SGD exclusively

to the discriminator. This method effectively ensures DP for the entire system at all times, as it makes all information flow private. Importantly, given that only the discriminator interacts with the private data, the generator also satisfies DP once the discriminator does so, due to the post-processing property. Additionally, confining DP-SGD to the discriminator circumvents introducing excessive noise to the generator, which could potentially impede the training process.

Extending GAN from the image domain to tabular data, Frigerio *et al.* [175] adopt the one-hot representation for the categorical inputs and adjust the generator’s output accordingly. CTGAN [198] proposes a new normalization approach for the numerical inputs by modeling it as a variational Gaussian mixture model (VGM) [199]. These adjustments in data representation and model architecture enable effective training of GAN and integration of DP-SGD [175, 177, 178]. Another line of work leverages the auto-encoding mechanism to automatically learn the representation for mixed-type tabular data, thereby circumventing the need for manual design of such representations. A detailed discussion of these works is offered in Section 4.1.

Perturbing the Discriminator’s Signal. In the case where only the generator is accessible to the adversary, it suffices to build a *privacy barrier* between the discriminator and the generator. Concretely, the training of the generator relies on the signal from the discriminator; perturbing the signal provides a privacy guarantee for the generator. Different works focus on different types of signal and different approaches of perturbation.

Partial gradient sanitization. GS-WGAN [96] points out that the gradient of the generator can be decomposed into two components via the chain rule: an upstream gradient

(discriminator loss w.r.t. the generated image) and a local gradient (generated image w.r.t. the generator’s parameters). The authors perturb the upstream gradient via the Gaussian mechanism, as it is the sole bearer of private information. This method effectively sanitizes the generator’s gradient without inducing substantial distortion, leading to more informative gradient updates and improved training outcomes.

Private knowledge aggregation and transfer. In the teacher-student framework [200], the knowledge of the teacher ensemble is aggregated and transferred to the student. In GANs, the discriminator acts as the teacher, providing signal (knowledge) to the generator (student). By privatizing this knowledge transfer, a privacy barrier is established between the discriminator and the generator.

Existing works consider two types of knowledge and employ different architectures for aggregation. PATE-GAN [180] perceives the discriminator as a classifier and utilizes its output (true/false) as knowledge. It extends PATE [201, 202], a DP classification scheme, to privatize the discriminator’s output. Specifically, they train multiple teacher discriminators along with a single student discriminator. The student learns from noisy labels that are obtained through privately aggregating the teachers’ votes. The update of the generator is then guided by the signal from the student discriminator. In comparison, G-PATE [181] and DataLens [182] focus on the discriminator’s gradient as knowledge. They employ an architecture that comprises an ensemble of teacher discriminators and a student generator. The student generator receives private gradients that are obtained through the private aggregation of the teachers’ gradients. To mitigate the impact of noise, G-PATE [181] performs dimensionality reduction on the gradients using random projection, whereas DataLens [182] conducts stochastic gradient compression in selective dimensions to reduce sensitivity.

Takeaways. *Researchers should tailor the development of PPDS algorithms according to the specific privacy requirements and threat models of each scenario. Applying existing privacy mechanisms indiscriminately to complex systems could lead to unwarranted reductions in data utility.*

Open Problems. *Existing works commonly attain DP in minimax training by applying DP-SGD to conventional generative models. An open problem, however, lies in the design or identification of architectures that are better suited to DP-SGD. Leveraging Neural Architecture Search (NAS) [203] could serve as a compelling approach, optimizing not only task performance but also compatibility with DP-SGD. This might involve a search strategy related to factors such as stable convergence, smooth gradient, and model sensitivity.*

4.3. PPDS Based on Optimal Transport

4.3.1. Preliminaries. Generative modeling based on optimal transport [204–206] adopt the Wasserstein distance (W-distance) [207] to measure the discrepancy between probability distributions. Minimizing this metric facilitates the generation of samples that resemble a target distribution. The

Sinkhorn divergence [204] provides a smooth approximation of the W-distance by incorporating a regularization term, eliminating the need for using a discriminator to estimate the W-distance as in WGANs [208]. This results in more efficient and stable training of a generative model, making it more robust to noise addition and suitable for DP.

4.3.2. Paper Summary. DP-Sinkhorn [183] minimizes the Sinkhorn divergence between the real and the generated distributions in a differentially private manner via partial gradient sanitization [96] (see Section 4.2). To balance the bias-variance trade-off, the authors introduce a semi-debiased Sinkhorn loss. This effectively tackles the limitations of previous empirical Sinkhorn losses, which are either biased [209] or have large variance [205]. The semi-debiased loss is computed by partitioning the generated data and evaluating both the internal divergence (self-term loss) as well as their divergence with the real data (cross-term loss). Privacy is ensured by sanitizing the gradient of the cross-term loss with respect to one partition of the data. The power of optimal transport and the Sinkhorn divergence makes DP-Sinkhorn excel at capturing multiple modes and converge stably.

Takeaways. *Utilizing Sinkhorn divergence enables a tractable and efficient approximation of the W-distance. In turn, this allows for the direct alignment of distributions, thereby eliminating the need for a discriminator and bypassing the instability issues inherent in minimax training. Moreover, the formulation of a distribution-aligning objective has the added benefit of mitigating mode collapse, a common pitfall in GANs.*

Open Problems. *The cost function in optimal transport plays a crucial role in data generation. Conventional choices involve using L_1 and L_2 distances, which promote image-level similarity but may not align with the desired properties for PPDS. An open problem is thus to design task-specific cost functions, potentially using kernel-based methods to capture complex features and enhance data utility. This direction, which goes beyond simple pixel-level resemblance, could lead to a more effective balance between data utility, privacy preservation, and generation quality.*

4.4. PPDS Based on Feature Alignment

4.4.1. Preliminaries. Generative models underpinned by feature alignment aim to learn a mapping from a noise distribution to the data distribution. This is achieved by aligning the features of the synthesized data with those of the real data, typically employing feature-level distances. This approach holds particular merit for PPDS as it ensures that the learned representation is useful for downstream tasks, without enforcing similarity in the data space.

In feature alignment, DP is achieved by adding noise to the features of the real data distribution, a technique known as *feature sanitization*. The generator built on these sanitized features is DP due to the post-processing property.

4.4.2. Paper Summaries. Kernel Mean Embeddings (KMEs) [210] transform a probability distribution into a reproducing kernel Hilbert space (RKHS) [211]

by computing the mean of kernel evaluations. Maximum Mean Discrepancy (MMD) [212] quantifies the distance between two probability distributions within the RKHS, thus offering a mechanism to impose distributional constraints in generative modeling. To achieve DP, one can add noise to the KME of the real data, then minimize the MMD between this noisy KME and the KME of the generated data [48, 184, 185].

The computation of KMEs is resource-intensive. Fortunately, an efficient approximation of KMEs can be achieved by employing the inner product of feature vectors [213]. DP-MERF [48] adopts the random Fourier features (RFFs) [214] to approximate KMEs. However, to achieve good approximation quality, a large number of random features are required. This translates into a high-dimensional feature vector, resulting in a significant amount of noise. DP-HP [184] circumvents this challenge by using Hermite polynomial features (HPFs) [215], which encapsulate a higher degree of information within a smaller order of features compared to RFFs. In the image domain, perceptual features (PFs), *i.e.*, the concatenation of activations from all layers of a pre-trained deep convolutional neural network (CNN) [216], have been effectively employed for generative modeling [217]. Building on PFs obtained from publicly pre-trained networks, DP-MEPF [185] generates samples of high visual quality on CIFAR [218] and CelebA [219] datasets, while maintaining small ϵ values such as 1 and 0.2, respectively.

Takeaways. *Feature alignment offers several advantages:*

- *Feature sanitization is a one-time process and eliminates the concern of accumulating privacy costs during subsequent training.*
- *The objective functions are simple and can be computed and optimized efficiently. For instance, DP-MERF [48] completes training within seconds.*

Open Problems. *In the realm of feature alignment based data synthesis, a captivating research opportunity lies in the exploration of dynamic, data-dependent feature selection. Could we build a meta-learning framework that identifies optimal features for different datasets or applications, like selecting disease-specific features for medical image generation? By evolving feature selection from a static decision into a context-aware, learning-based process, we could push the boundaries of privacy-preserving data synthesis.*

4.5. PPDS Based on Stochastic Simulation

4.5.1. Preliminaries. Diffusion Models (DMs) leverage stochastic simulations to represent complex, high-dimensional data distributions, thereby displaying remarkable capabilities in image synthesis [188, 189, 220–222]. A defining property of DMs is their iterative generation mechanism, starting with random noise and progressively refining it into a high-fidelity image that resembles the target distribution.

4.5.2. Paper Summaries. Different types of DMs vary in their approach to the *forward process*, in which data is gradually transformed into noise, and the *reverse process*,

where data is generated from noise through a series of denoising steps. We review two variants of DMs [223] and discuss how existing works achieve DP within these models.

Score-Based Generative Models (SGMs) [220, 221] utilize a neural network to approximate the gradient of the log-density, known as the *score* [224], of a target distribution. During the training phase, SGMs simulate the forward process by perturbing the data with Gaussian noise. A noise-conditioned score network (NCSN) is then trained to estimate the score function for the noisy data distributions. During the generation phase, SGMs employ score-based sampling approaches that leverage the score network, which guides the iterative refinement in the reverse direction. Common score-based sampling approaches include Langevin Monte Carlo [220, 225].

DPGEN [186] introduces a novel approach for generating DP synthesized data by learning DP scores. It begins with the creation of a sanitized dataset, consisting of Gaussian-perturbed samples and their respective private scores. Then, an NCSN is trained on this sanitized data. To privatize the scores, the authors leverage Randomized Response (RR) [226], which probabilistically associates each perturbed sample with a nearby neighbor rather than its original non-perturbed source. An advantage of RR is that it incurs minimal error for $\epsilon > 1$. Additionally, once the sanitized dataset is created, extending the training duration solely enhances the generation quality without introducing additional privacy costs. Putting together, DPGEN can synthesize perceptually realistic images of high dimensionality (CelebA [219] and LSUN [227] of size $128 \times 128 \times 3$) at $\epsilon = 5$.

Stochastic Differential Equations (Score SDEs) [189, 228] establish a framework where the forward diffusion process is governed by an SDE consisting of a drift term and a diffusion term. Any such diffusion process can be reversed by solving another reverse-time SDE [229], which relies on the score function. Once the score function at each step is learned, we can solve the reverse-time SDE with numerical methods, effectively transforming the noise into real data.

DPDM [187] achieves DP by applying DP-SGD to the denoiser (similar to a score network). The authors note that the denoiser often has low capacity, and is thus amenable to DP-SGD since the magnitude of noise grows linearly with the number of model parameters. They demonstrate state-of-the-art (SOTA) performance on various benchmarks.

Takeaways. *The sequential nature of DMs grants them a distinct advantage in PPDS. Unlike other generative modeling approaches that opt for a single-pass synthesis, DMs utilize a sequence of gradual refinements, fostering a smoother and more controlled generation process. The modest progression allows DMs to seamlessly incorporate the distortions necessitated by privacy requirements, minimizing the impact on the utility of the synthesized data.*

Open Problems. *A critical open problem is to develop a privacy accounting scheme tailored to the sequential nature of diffusion models. Specifically, exploring step-wise privacy loss accounting during model training stands as a promising*

direction. Such an approach has the potential to deliver a more accurate distribution of the privacy budget. It recognizes that different steps might contribute unevenly to overall privacy leakage, allowing for more precise tracking and mitigation.

5. Evaluation

We benchmark DL-based methods on the task of private image data synthesis, using MNIST [230] and FashionMNIST [231], two datasets commonly adopted in the literature. The results are presented in Table 3. We make two additional comments here. First, PPDS significantly lags behind its non-private counterpart in the image domain. As pointed out by Chen *et al.* [186], very few works can generate meaningful 32×32 images under a reasonable privacy budget. Second, we narrow our focus to image data to distinguish our study from a prior work [57], which set a benchmark for DP tabular data synthesis. We refer readers to the Appendix for general method selection criteria in real-world scenarios.

5.1. Experimental Setups

Inclusion & Exclusion Criteria. We aim to build a comprehensive benchmark covering all PPDS milestone approaches from the papers accepted at top ML & security conferences. We first narrowed down the approaches by available code repositories on GitHub. We then made an initial attempt to run each approach, following the instructions in their README and applying the hyper-parameters reported in their evaluation settings. We exclude approaches that either 1) fail to work; or 2) demonstrate significant discrepancy in our re-implementation, in order to prevent any misrepresentation of their method. In the end, we obtain a list of six approaches, with three (DPGAN [174], DP-CGAN [95], and GS-WGAN [96]) under the category of minimax training (Section 4.2), one (DP-MERF [48]) under feature alignment (Section 4.4), one (DPGEN [186]) under stochastic simulation (Section 4.5), and one (DP-Sinkhorn [183]) under optimal transport (Section 4.3).

Evaluation Scenarios. We evaluate the performance of these approaches on two standard scenarios ($\epsilon = 10$, $\epsilon = 1$) and two challenging scenarios (stringent privacy— $\epsilon = 0.2$, and smaller population—half training set size under $\epsilon = 10$). We also evaluate the non-private scenario using the non-private version of each approach. This is used as a baseline.

Evaluation Metrics. We synthesize a dataset of the same size as the training set. We follow the literature [48, 96, 183] to evaluate the fidelity and utility of the synthesized data. For fidelity, we measure the Fréchet Inception Distance (FID) [232] and the Inception Score (IS) [233]. For utility, we train classifiers on the synthesized dataset and test their performance on the real test data. We adopt 13 classifiers in all, including convolutional neural network (CNN), multi-layer perceptron (MLP), and 11 other scikit-learn [234] classifiers. We measure the performance of these classifiers using classification *accuracy*. Due to limited space, we report the averaged accuracy on the 11 scikit-learn classifiers, denoted as “Avg” in Table 3. For each combination of

DDPS approach and evaluation scenario, we report the results averaged over five runs of dataset generation and classifier training following [48, 96].

We refer readers to our website for detailed information on the experimental setups (*e.g.*, hyper-parameter tuning) and additional experimental results (*e.g.*, runtime comparisons).

5.2. Highlighted Conclusions

DP-MERF is the best practice. Among the evaluated approaches, DP-MERF [48] stands out as the only one that consistently delivers high classification accuracy across all scenarios. It demonstrates minimal performance loss compared to its non-private counterpart, indicating its compatibility with DP. In contrast, DPGEN degrades significantly in the small epsilon regime. Additionally, DP-MERF exhibits remarkable computational efficiency. Thus, we highly recommend DP-MERF as an all-purpose approach.

We note that the visual quality (measured by FID and IS) of DP-MERF is not the best; yet it effectively learns the key features for classification. This trait is beneficial for classification, but caution is needed when the synthesized dataset is applied to other types of downstream tasks.

The pessimistic conclusion in Stadler *et al.* [49] does not fully extend to image data. Stadler *et al.* conclude that synthetic tabular data either has poor privacy or utility. However, our experimental results suggest a different narrative. Although the general trend across all approaches is that one has to sacrifice a significant amount of utility for strong privacy, two methods (DP-Sinkhorn and DP-MERF) yield notable performance at $\epsilon = 1$, which is typically viewed as stringent privacy that meets practical requirements [235]. In particular, DP-MERF demonstrates the possibility of simultaneously achieving high utility and strong privacy in synthetic data, implying that such solutions exist—it’s simply a matter of discovering them.

Fidelity and utility are positively correlated. We compute the Kendall rank correlation coefficient [236] between utility and fidelity measures for the above approaches, using an average derived from six pairs of measures (three for utility and two for fidelity). This is further averaged across five scenarios (non-private, two standard, and two challenging). The obtained value is 0.47 for MNIST and 0.65 for FashionMNIST, indicating a strong positive correlation between fidelity and utility in image data.

We refer readers to our website for more conclusions.

Recommended Practices. *Across the course of our replication study, we have pinpointed a number of concerns that permeate various works in the field. These issues are elaborated upon on our website, leading us to propose several best practice recommendations:*

- Report the runtime and discuss the computational efficiency.
- Acknowledge and report any limitations, such as the incapability of generating labels.
- Furnish rigorous privacy analysis both in the paper and the associated code; elucidate the factors or parameters influencing the privacy analysis.
- Be cautious about inadvertent privacy leakage in algorithmic

TABLE 3: **Main results.** We evaluate six DL-based approaches on two datasets under five scenarios. We measure the utility (by classifier accuracy) and fidelity (by FID and IS) of the synthesized data. ‘‘Avg’’ refers to the average accuracy over 11 scikit-learn classifiers. Each number is averaged over five runs of synthetic data generation and classifier training. The best performing approach is presented in **bold**, while the second best performing one is shown in *italics*.

MNIST	Non-private					Standard										Challenging									
						$\epsilon = 10$					$\epsilon = 1$					$\epsilon = 0.2$					Half ($\epsilon = 10$)				
	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑
DPGAN	85.5	88.9	52.0	63.1	8.6	74.8	74.7	54.7	268.8	2.0	40.0	58.6	27.0	387.1	1.0	15.1	8.9	12.3	420.2	1.3	69.7	71.2	48.2	295.1	1.7
DP-CGAN	84.6	91.6	71.3	39.7	9.4	60.5	64.8	51.7	174.2	4.2	9.3	25.3	15.5	281.8	1.5	9.2	12.9	10.2	292.0	1.0	62.5	65.0	49.8	248.1	2.3
GS-WGAN	85.7	84.9	60.3	64.3	9.8	79.5	78.3	55.7	60.6	8.2	27.0	13.2	16.3	392.5	1.0	8.3	10.5	9.7	489.6	1.0	73.2	77.7	46.2	59.2	7.7
DP-MERF	83.4	87.0	67.0	103.7	5.5	<i>80.4</i>	<i>83.2</i>	<i>69.4</i>	103.5	6.3	<i>81.8</i>	<i>83.9</i>	66.6	112.0	6.4	<i>75.6</i>	<i>79.1</i>	54.5	130.9	5.7	<i>80.0</i>	<i>82.6</i>	<i>65.2</i>	106.8	6.4
DP-Sinkhorn	<i>89.4</i>	92.3	66.1	52.3	9.8	77.0	79.6	52.5	78.0	5.9	61.5	61.9	33.6	200.0	3.3	56.6	48.8	34.8	208.9	1.1	75.2	78.2	44.8	87.7	5.2
DPGEN	95.7	98.2	83.5	6.2	7.9	95.0	98.2	78.5	9.3	6.7	92.3	97.6	48.5	125.5	1.3	79.8	96.1	36.5	<i>183.1</i>	1.0	96.0	98.4	84.1	5.5	8.0

Fashion-MNIST	Non-private					Standard										Challenging									
						$\epsilon = 10$					$\epsilon = 1$					$\epsilon = 0.2$					Half ($\epsilon = 10$)				
	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑	MLP	CNN	Avg	FID↓	IS↑
DPGAN	79.3	78.7	60.2	<i>64.0</i>	7.7	63.9	59.5	50.1	290.8	3.7	44.8	54.1	31.2	433.7	1.4	15.7	20.6	12.8	438.2	1.4	61.9	58.7	47.0	316.0	2.3
DP-CGAN	64.3	64.5	56.6	147.6	5.2	53.5	49.0	37.9	256.5	3.3	24.5	18.0	21.0	339.1	1.8	12.0	12.8	11.2	331.2	1.7	36.4	37.8	35.4	351.3	1.3
GS-WGAN	76.4	71.1	58.7	83.7	8.0	65.9	64.7	50.7	125.9	5.7	21.6	12.5	19.8	334.1	1.0	9.7	16.0	9.4	490.5	1.0	67.8	65.2	53.4	118.3	5.6
DP-MERF	74.8	73.0	<i>61.0</i>	101.0	5.4	<i>73.5</i>	<i>71.7</i>	58.2	<i>102.3</i>	5.5	72.6	72.4	58.9	97.1	5.3	70.1	69.7	51.0	139.8	4.2	73.6	73.7	56.9	105.8	5.1
DP-Sinkhorn	77.6	77.5	59.1	90.8	7.5	69.4	65.4	50.2	164.3	4.5	58.0	56.3	38.5	216.2	2.9	<i>45.9</i>	<i>44.3</i>	<i>30.7</i>	260.5	2.6	70.2	67.8	48.6	164.0	4.9
DPGEN	80.2	84.5	71.2	14.4	7.3	79.5	84.0	70.2	11.0	7.3	36.3	41.7	21.3	<i>113.2</i>	2.0	27.1	33.2	17.3	<i>190.0</i>	1.2	80.2	84.6	70.6	12.2	7.3

design (e.g., utilizing a classifier trained on private data for label assignment) and implementation (see [49]).

- Provide explicit descriptions of all algorithmic steps.

6. Future Directions

Emerging Applications. Current PPDS algorithms remain inadequate for addressing the challenges posed by the high-dimensional unstructured data, which are prevalent in contemporary everyday applications. Such data types are exemplified by high-resolution images [237, 238], audio files [239], videos [240], and text [241]. To mitigate this issue, it would be beneficial for researchers to leverage the most recent advancements in diffusion models [222] and large language models (LLMs) [242], integrating them with privacy mechanisms to produce a more robust solution.

Privacy Attacks. Conventional membership inference attacks primarily operate on model predictions or low-dimensional feature vectors such as logits [243]. However, PPDS systems yield high-dimensional synthetic records, inherently richer in information. This poses the question of whether such increased information enables novel attacks on PPDS systems, thereby facilitating record linkage. Stadler *et al.* [49] initiated attempts in this direction with synthetic tabular data. We advocate for further development of innovative attacks that can be applied to image data.

Privacy Auditing. Given a synthetic dataset released by a PPDS system, we would like to quantitatively assess its adherence to the claimed privacy guarantees. Previous work has demonstrated the feasibility of achieving such goal for a given private model through data poisoning [244]. It warrants exploration as to whether comparable techniques can be effectively applied to PPDS systems.

Privacy Notions. When applying PPDS to pragmatic domains such as healthcare and IoT, DP may fall short due to its lack of intuitive interpretation. In contrast, the

level of privacy with respect to downstream tasks is often more interpretable and consequently frequently adopted. For context-aware tasks of this nature, the enforcement of DP may lead to unnecessary utility degradation. Therefore, it is crucial to develop privacy notions that are tailored for specific applications or tasks.

Threat Models. Existing literature has demonstrated that making assumptions about an adversary’s knowledge could permit the use of a larger ϵ , thereby enhancing utility [245]. In the context of PPDS, it is reasonable to formulate several plausible assumptions about the adversary, such as lack of access to the training process or even the trained model, with only the released data being available. These assumptions could pave the way for a refined privacy analysis.

Latent Issues of PPDS. PPDS has been demonstrated to degrade fairness, leading to disproportionate impacts on underrepresented groups [246, 247]. Furthermore, the ranking of feature importance within the synthetic data is not sufficiently preserved [248]. There are likely additional latent issues with PPDS that require thorough investigation.

Finally, we wish to underscore a dilemma pervading the research community: namely, *the lack of a consensus on the selection of appropriate utility metrics and privacy notions for the evaluation of synthetic data*. Privacy notions, even when restricted to DP, are frequently misapplied; a scant number of studies incorporate privacy tests. Each paper appears to employ its unique suite of utility metrics, thereby risking an evaluation bias in favor of their specific algorithm. Moreover, there is a conspicuous dearth of task-centric utility metrics for trajectory and graph data. Consequently, we urgently advocate for the establishment of a comprehensive benchmark, replete with a consistent and exhaustive array of utility metrics and privacy tests.

Acknowledgments. We are extremely grateful to Varun Chandrasekaran for his help in improving this work.

References

- [1] J. A. Sidey-Gibbons and C. J. Sidey-Gibbons, "Machine learning in medicine: a practical introduction," *BMC medical research methodology*, vol. 19, no. 1, pp. 1–18, 2019.
- [2] I. Kononenko, "Machine learning for medical diagnosis: history, state of the art and perspective," *Artificial Intelligence in medicine*, vol. 23, no. 1, pp. 89–109, 2001.
- [3] S. Chik, C. Yeong, E. Su, T. Lim, Y. Subramaniam, and P. Chin, "A review of social-aware navigation frameworks for service robot in dynamic human environments," *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, vol. 8, no. 11, pp. 41–50, 2016.
- [4] K. Siau and W. Wang, "Building trust in artificial intelligence, machine learning, and robotics," *Cutter business technology journal*, vol. 31, no. 2, pp. 47–53, 2018.
- [5] Y. Kortli, M. Jridi, A. Al Falou, and M. Atri, "Face recognition systems: A survey," *Sensors*, vol. 20, no. 2, p. 342, 2020.
- [6] M. Wang and W. Deng, "Deep face recognition: A survey," *Neurocomputing*, vol. 429, pp. 215–244, 2021.
- [7] D. Gavriluț, M. Cimpoșu, D. Anton, and L. Ciortuz, "Malware detection using machine learning," in *2009 International Multiconference on Computer Science and Information Technology*, pp. 735–741, IEEE, 2009.
- [8] D. Gibert, C. Mateu, and J. Planes, "The rise of machine learning for detection and classification of malware: Research developments, trends and challenges," *Journal of Network and Computer Applications*, vol. 153, p. 102526, 2020.
- [9] D. J. Glancy, "Privacy in autonomous vehicles," *Santa Clara L. Rev.*, vol. 52, p. 1171, 2012.
- [10] J. Boeglin, "The costs of self-driving cars: Reconciling freedom and privacy with tort liability in autonomous vehicle regulation," *Yale JL & Tech.*, vol. 17, p. 171, 2015.
- [11] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1322–1333, 2015.
- [12] Y. Zhang, R. Jia, H. Pei, W. Wang, B. Li, and D. Song, "The secret revealer: Generative model-inversion attacks against deep neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 253–261, 2020.
- [13] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart, "Privacy in pharmacogenetics: An {End-to-End} case study of personalized warfarin dosing," in *23rd USENIX Security Symposium (USENIX Security 14)*, pp. 17–32, 2014.
- [14] C. G. Schwarz, W. K. Kremers, T. M. Therneau, R. R. Sharp, J. L. Gunter, P. Vemuri, A. Arani, A. J. Spychalla, K. Kantarci, D. S. Knopman, et al., "Identification of anonymous mri research participants with face-recognition software," *New England Journal of Medicine*, vol. 381, no. 17, pp. 1684–1686, 2019.
- [15] I. report, "Cost of a data breach report." <https://www.ibm.com/downloads/cas/OJDVQGRY>. Online; accessed 11 April 2022.
- [16] R. Nosowsky and T. J. Giordano, "The health insurance portability and accountability act of 1996 (hipaa) privacy rule: implications for clinical research," *Annu. Rev. Med.*, vol. 57, pp. 575–590, 2006.
- [17] C. Tobias, "Public law litigation and the federal rules of civil procedure," *Cornell L. Rev.*, vol. 74, p. 270, 1988.
- [18] J. Davis, "Banner health pays \$1.25m penalty over hipaa failures from 2016 breach — sc media." <https://www.scmagazine.com/analysis/compliance/banner-health-pays-1-25m-penalty-over-hipaa-failures-from-2016-breach>.
- [19] H. Landi, "Premera blue cross pays 2nd-largest hipaa fine for 2014 breach — fierce healthcare." <https://www.fiercehealthcare.com/tech/premera-blue-cross-to-pay-6-9m-to-hhs-for-2014-data-breach>.
- [20] D. Huth, F. Burmeister, F. Matthes, and I. Schirmer, "Empirical Results on the Collaboration Between Enterprise Architecture and Data Protection Management during the Implementation of the GDPR," in *HICSS*, pp. 1–10, 2020.
- [21] R. Brand, "Microdata protection through noise addition," in *Inference control in statistical databases*, pp. 97–116, Springer, 2002.
- [22] A. Blum, C. Dwork, F. McSherry, and K. Nissim, "Practical privacy: The sulp framework," in *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 128–138, 2005.
- [23] J. Vaidya and C. Clifton, "Privacy-preserving data mining: Why, how, and when," *IEEE Security & Privacy*, vol. 2, no. 6, pp. 19–27, 2004.
- [24] C. C. Aggarwal and P. S. Yu, "A general survey of privacy-preserving data mining models and algorithms," in *Privacy-preserving data mining*, pp. 11–52, Springer, 2008.
- [25] R. Mendes and J. P. Vilela, "Privacy-preserving data mining: methods, metrics, and applications," *IEEE Access*, vol. 5, pp. 10562–10582, 2017.
- [26] Z. Ji, Z. C. Lipton, and C. Elkan, "Differential privacy and machine learning: a survey and review," *arXiv preprint arXiv:1412.7584*, 2014.
- [27] P. Mohassel and Y. Zhang, "Secureml: A system for scalable privacy-preserving machine learning," in *2017 IEEE symposium on security and privacy (SP)*, pp. 19–38, IEEE, 2017.
- [28] M. Al-Rubaie and J. M. Chang, "Privacy-preserving machine learning: Threats and solutions," *IEEE Security & Privacy*, vol. 17, no. 2, pp. 49–58, 2019.
- [29] A. Boulemtafes, A. Derhab, and Y. Challal, "A review of privacy-preserving techniques for deep learning," *Neurocomputing*, vol. 384, pp. 21–45, 2020.
- [30] B. Liu, M. Ding, S. Shaham, W. Rahayu, F. Farokhi, and Z. Lin, "When machine learning meets privacy: A survey and outlook," *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–36, 2021.
- [31] M. G. Hay, *Enabling accurate analysis of private network data*. University of Massachusetts Amherst, 2010.
- [32] F. K. Dankar and K. El Emam, "The application of differential privacy to health data," in *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, pp. 158–166, 2012.
- [33] N. Mohammed, R. Chen, B. C. Fung, and P. S. Yu, "Differentially private data release for data mining," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 493–501, 2011.
- [34] A. Blum, K. Ligett, and A. Roth, "A learning theory approach to noninteractive database privacy," *Journal of the ACM (JACM)*, vol. 60, no. 2, pp. 1–25, 2013.
- [35] C. Clifton and T. Tassa, "On syntactic anonymity and differential privacy," in *2013 IEEE 29th International Conference on Data Engineering Workshops (ICDEW)*, pp. 88–93, IEEE, 2013.
- [36] A. Narayanan and V. Shmatikov, "How to break anonymity of the netflix prize dataset," *arXiv preprint cs/0610105*, 2006.
- [37] R. C.-W. Wong, A. W.-C. Fu, K. Wang, and J. Pei, "Minimality attack in privacy preserving data publishing," in *Proceedings of the 33rd international conference on Very large data bases*, pp. 543–554, 2007.
- [38] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pp. 111–125, IEEE, 2008.
- [39] P. Ohm, "Broken promises of privacy: Responding to the surprising failure of anonymization," *UCLA L. Rev.*, vol. 57, p. 1701, 2009.
- [40] M. Fiore, P. Katsikouli, E. Zavou, M. Cunche, F. Fessant, D. Le Hello, U. Aivodji, B. Olivier, T. Quertier, and R. Stanica, "Privacy in trajectory micro-data publishing: a survey," *Transactions on Data Privacy*, vol. 13, pp. 91–149, 2020.
- [41] S. R. Eddy, "What is a hidden markov model?," *Nature biotechnology*, vol. 22, no. 10, pp. 1315–1316, 2004.
- [42] N. Friedman, D. Geiger, and M. Goldszmidt, "Bayesian network classifiers," *Machine learning*, vol. 29, no. 2, pp. 131–163, 1997.
- [43] R. B. Nelsen, *An introduction to copulas*. Springer science & business media, 2007.
- [44] S. Bond-Taylor, A. Leach, Y. Long, and C. G. Willcocks, "Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.
- [45] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- [46] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations*, 2014.
- [47] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography (S. Halevi and T. Rabin, eds.)*, (Berlin, Heidelberg), pp. 265–284, Springer Berlin Heidelberg, 2006. Online; accessed 30 December 2021.
- [48] F. Harder, K. Adamczewski, and M. Park, "Dp-merf: Differentially private mean embeddings with random features for practical privacy-preserving data generation," in *International Conference on Artificial Intelligence and Statistics*, pp. 1819–1827, PMLR, 2021.
- [49] T. Stadler, B. Oprisanu, and C. Troncoso, "Synthetic data-anonymisation groundhog day," *arXiv preprint arXiv:2011.07018*, 2021.
- [50] B. Zhou, J. Pei, and W. Luk, "A brief survey on anonymization techniques for privacy preserving publishing of social network data," *ACM Sigkdd Explorations Newsletter*, vol. 10, no. 2, pp. 12–22, 2008.
- [51] X. Wu, X. Ying, K. Liu, and L. Chen, "A survey of privacy-preservation of graphs and social networks," in *Managing and mining graph data*, pp. 421–453, Springer, 2010.
- [52] B.-C. Chen, D. Kifer, K. LeFevre, A. Machanavajjhala, et al., "Privacy-preserving data publishing," *Foundations and Trends® in Databases*, vol. 2, no. 1–2, pp. 1–167, 2009.
- [53] B. C. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Computing Surveys (Csur)*, vol. 42, no. 4, pp. 1–53, 2010.
- [54] L. Fan, "A survey of differentially private generative adversarial networks," in *The AAAI Workshop on Privacy-Preserving Artificial Intelligence*, p. 8, 2020.
- [55] Z. Cai, Z. Xiong, H. Xu, P. Wang, W. Li, and Y. Pan, "Generative adversarial networks: A survey toward private and secure applications," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–38, 2021.
- [56] C. M. Bowen and J. Snoke, "Comparative study of differentially private synthetic data algorithms from the nist psr differential privacy synthetic data challenge," *Journal of Privacy and Confidentiality*, vol. 11, no. 1, 2021.
- [57] Y. Tao, R. McKenna, M. Hay, A. Machanavajjhala, and G. Miklau, "Benchmarking differentially private synthetic data generation algorithms," in *Proceedings of the Third AAAI Workshop on Privacy-Preserving Artificial Intelligence (PPAI-22)*, 2022.
- [58] J. Jordon, L. Szpruch, F. Houssiau, M. Bottarelli, G. Cherubin, C. Maple, S. N. Cohen, and A. Weller, "Synthetic data—what, why and how?," *arXiv preprint arXiv:2205.03257*, 2022.
- [59] K. Liu, K. Das, T. Grandison, and H. Kargupta, "Privacy-preserving data analysis on graphs and social networks," in *Next generation of data mining*, pp. 443–462, Chapman and Hall/CRC, 2008.
- [60] Y. Wang, X. Wu, and L. Wu, "Differential privacy preserving spectral graph analysis," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 329–340, Springer, 2013.
- [61] C. Li, G. Miklau, M. Hay, A. McGregor, and V. Rastogi, "The matrix mechanism: optimizing linear counting queries under differential privacy," *The VLDB journal*, vol. 24, no. 6, pp. 757–781, 2015.
- [62] R. McKenna, G. Miklau, M. Hay, and A. Machanavajjhala, "Optimizing error of high-dimensional statistical queries under differential privacy," *Proceedings of the VLDB Endowment*, vol. 11, no. 10, pp. 1206–1219, 2018.
- [63] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar, "Privacy, accuracy, and consistency too: a holistic solution to contingency table release," in *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 273–282, 2007.
- [64] W. Qardaji, W. Yang, and N. Li, "Privity: practical differentially private release of marginal contingency tables," in *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pp. 1435–1446, 2014.
- [65] A. F. Westin, "Privacy and freedom," *Washington and Lee Law Review*, vol. 25, no. 1, p. 166, 1968.
- [66] S. Warren and L. Brandeis, "The right to privacy," in *Killing the Messenger*, pp. 1–21, Columbia University Press, 1989.

- [67] F. T. Wu, "Defining privacy and utility in data sets," *U. Colo. L. Rev.*, vol. 84, p. 1117, 2013.
- [68] A. Cohen and K. Nissim, "Towards formalizing the gdpr's notion of singling out," *Proceedings of the National Academy of Sciences*, vol. 117, no. 15, pp. 8344–8352, 2020.
- [69] C. C. Porter, "De-identified data and third party data mining: The risk of re-identification of personal information," *Shidler JL Com. & Tech.*, vol. 5, p. 1, 2008.
- [70] K. El Emam, E. Jonker, L. Arbuckle, and B. Malin, "A systematic review of re-identification attacks on health data," *PLoS one*, vol. 6, no. 12, p. e28071, 2011.
- [71] B. Malin and L. Sweeney, "How (not) to protect genomic data privacy in a distributed network: using trail re-identification to evaluate and design anonymity protection systems," *Journal of biomedical informatics*, vol. 37, no. 3, pp. 179–192, 2004.
- [72] L. Rocher, J. M. Hendrickx, and Y.-A. De Montjoye, "Estimating the success of re-identifications in incomplete datasets using generative models," *Nature communications*, vol. 10, no. 1, pp. 1–9, 2019.
- [73] E. Zheleva and L. Getoor, "Preserving the privacy of sensitive relationships in graph data," in *International workshop on privacy, security, and trust in KDD*, pp. 153–171, Springer, 2007.
- [74] E. M. Newton, L. Sweeney, and B. Malin, "Preserving privacy by de-identifying face images," *IEEE transactions on Knowledge and Data Engineering*, vol. 17, no. 2, pp. 232–243, 2005.
- [75] 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.
- [76] F. Brunton and H. Nissenbaum, *Obfuscation: A user's guide for privacy and protest*. MIT Press, 2015.
- [77] S. Ji, P. Mittal, and R. Beyah, "Graph data anonymization, de-anonymization attacks, and de-anonymizability quantification: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 1305–1326, 2016.
- [78] NIST, "Nist 2018 differential privacy synthetic data competition." <https://www.nist.gov/ct/pscr/open-innovation-prize-challenges/past-prize-challenges/2018-differential-privacy-synthetic>. Online; accessed 28 July 2022.
- [79] NIST, "Nist 2020 differential privacy temporal map challenge." <https://www.nist.gov/ct/pscr/open-innovation-prize-challenges/past-prize-challenges/2020-differential-privacy-temporal>. Online; accessed 28 July 2022.
- [80] A. Cohen, "Attacks on deidentification's defenses," in *31st USENIX Security Symposium (USENIX Security 22)*, (Boston, MA), pp. 1469–1486, USENIX Association, Aug. 2022.
- [81] C. Dwork and J. Lei, "Differential privacy and robust statistics," in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 371–380, 2009.
- [82] 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.
- [83] S. Garfinkel, "Implementing differential privacy for the 2020 census," *USENIX Association*, 2021.
- [84] I. Mironov, "Rényi differential privacy," in *Computer Security Foundations Symposium (CSF)*, 2017 *IEEE 30th*, pp. 263–275, IEEE, 2017.
- [85] A. Rényi et al., "On measures of entropy and information," in *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, vol. 1, Berkeley, California, USA, 1961.
- [86] M. Hay, C. Li, G. Miklau, and D. Jensen, "Accurate estimation of the degree distribution of private networks," in *2009 Ninth IEEE International Conference on Data Mining*, pp. 169–178, IEEE, 2009.
- [87] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pp. 94–103, IEEE, 2007.
- [88] M. Abadi, A. Chu, I. J. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [89] C. Xu, J. Ren, Y. Zhang, Z. Qin, and K. Ren, "Dppro: Differentially private high-dimensional data release via random projection," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 12, pp. 3081–3093, 2017.
- [90] M. Hardt, K. Ligett, and F. McSherry, "A simple and practical algorithm for differentially private data release," in *Advances in Neural Information Processing Systems*, pp. 2339–2347, 2012.
- [91] H. Li, L. Xiong, and X. Jiang, "Differentially private synthesis of multi-dimensional data using copula functions," in *Advances in database technology: proceedings. International conference on extending database technology*, vol. 2014, p. 475, NIH Public Access, 2014.
- [92] S. Zhang, W. Ni, and N. Fu, "Differentially private graph publishing with degree distribution preservation," *Computers & Security*, vol. 106, p. 102285, 2021.
- [93] M. Gaboardi, E. J. G. Arias, J. Hsu, A. Roth, and Z. S. Wu, "Dual query: Practical private query release for high dimensional data," in *International Conference on Machine Learning*, pp. 1170–1178, PMLR, 2014.
- [94] J. Fan, J. Chen, T. Liu, Y. Shen, G. Li, and X. Du, "Relational data synthesis using generative adversarial networks: A design space exploration," *Proceedings of the VLDB Endow.*, vol. 13, no. 12, pp. 1962–1975, 2020.
- [95] R. Torkzadehmahani, P. Kairouz, and B. Paten, "Dp-cgan: Differentially private synthetic data and label generation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [96] D. Chen, T. Orekondy, and M. Fritz, "Gs-wgan: A gradient-sanitized approach for learning differentially private generators," *Neural Information Processing Systems (NeurIPS)*, 2020.
- [97] J. Zhang, G. Cormode, C. M. Procopiuc, D. Srivastava, and X. Xiao, "Privbayes: Private data release via bayesian networks," *ACM Transactions on Database Systems (TODS)*, vol. 42, no. 4, pp. 1–41, 2017.
- [98] V. Bindshaedler, R. Shokri, and C. A. Gunter, "Plausible deniability for privacy-preserving data synthesis," *Proceedings of the VLDB Endow.*, vol. 10, no. 5, pp. 481–492, 2017.
- [99] R. Chen, Q. Xiao, Y. Zhang, and J. Xu, "Differentially private high-dimensional data publication via sampling-based inference," in *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 129–138, 2015.
- [100] R. McKenna, G. Miklau, and D. Sheldon, "Winning the nist contest: A scalable and general approach to differentially private synthetic data," *Journal of Privacy and Confidentiality*, vol. 11, no. 3, 2021.
- [101] Z. Zhang, T. Wang, N. Li, J. Honorio, M. Backes, S. He, J. Chen, and Y. Zhang, "Privsyn: Differentially private data synthesis," in *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.
- [102] K. Cai, X. Lei, J. Wei, and X. Xiao, "Data synthesis via differentially private markov random fields," *Proceedings of the VLDB Endowment*, vol. 14, no. 11, pp. 2190–2202, 2021.
- [103] G. Vietri, G. Tian, M. Bun, T. Steinke, and S. Wu, "New oracle-efficient algorithms for private synthetic data release," in *International Conference on Machine Learning*, pp. 9765–9774, PMLR, 2020.
- [104] S. Aydoore, W. Brown, M. Kearns, K. Kenthapadi, L. Melis, A. Roth, and A. A. Siva, "Differentially private query release through adaptive projection," in *International Conference on Machine Learning*, pp. 457–467, PMLR, 2021.
- [105] H. Asghar, M. Ding, T. Rakotoarivelo, S. Mrabet, and D. Kaafar, "Differentially private release of high-dimensional datasets using the gaussian copula," *Journal of Privacy and Confidentiality*, vol. 10, no. 2, pp. 1–38, 2020.
- [106] S. Gamsb, F. Ladouceur, A. Laurent, and A. Roy-Gaumont, "Growing synthetic data through differentially-private vine copulas," *Proc. Priv. Enhancing Technol.*, vol. 2021, no. 3, pp. 122–141, 2021.
- [107] K. Kenthapadi, A. Korolova, I. Mironov, and N. Mishra, "Privacy via the johnson-lindenstrauss transform," *Journal of Privacy and Confidentiality*, vol. 5, no. 1, 2013.
- [108] Y. D. Li, Z. Zhang, M. Winslett, and Y. Yang, "Compressive mechanism: Utilizing sparse representation in differential privacy," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*, pp. 177–182, 2011.
- [109] X. Jiang, Z. Ji, S. Wang, N. Mohammed, S. Cheng, and L. Ohno-Machado, "Differentially-private data publishing through component analysis," *Transactions on Data Privacy*, vol. 6, no. 1, p. 19, 2013.
- [110] T. Chanyaswad, C. Liu, and P. Mittal, "Ron-gauss: Enhancing utility in non-interactive private data release," *Proc. Priv. Enhancing Technol.*, vol. 2019, no. 1, pp. 26–46, 2019.
- [111] R. Chen, B. C. Fung, B. C. Desai, and N. M. Sossou, "Differentially private transit data publication: a case study on the montreal transportation system," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 213–221, 2012.
- [112] R. Chen, G. Acs, and C. Castelluccia, "Differentially private sequential data publication via variable-length n-grams," in *Proceedings of the 2012 ACM conference on Computer and communications security*, pp. 638–649, 2012.
- [113] J. Zhang, X. Xiao, and X. Xie, "Privtree: A differentially private algorithm for hierarchical decompositions," in *Proceedings of the 2016 International Conference on Management of Data*, pp. 155–170, 2016.
- [114] X. He, G. Cormode, A. Machanavajhala, C. M. Procopiuc, and D. Srivastava, "Dpt: differentially private trajectory synthesis using hierarchical reference systems," *Proceedings of the VLDB Endowment*, vol. 8, no. 11, pp. 1154–1165, 2015.
- [115] S. Wang and R. O. Sinnott, "Protecting personal trajectories of social media users through differential privacy," *Computers & Security*, vol. 67, pp. 142–163, 2017.
- [116] D. J. Mir, S. Isaacman, R. Cáceres, M. Martonosi, and R. N. Wright, "Dp-where: Differentially private modeling of human mobility," in *2013 IEEE international conference on big data*, pp. 580–588, IEEE, 2013.
- [117] H. Roy, M. Kantarcioglu, and L. Sweeney, "Practical differentially private modeling of human movement data," in *IFIP Annual Conference on Data and Applications Security and Privacy*, pp. 170–178, Springer, 2016.
- [118] M. E. Gursoy, L. Liu, S. Truex, and L. Yu, "Differentially private and utility preserving publication of trajectory data," *IEEE Transactions on Mobile Computing*, vol. 18, no. 10, pp. 2315–2329, 2018.
- [119] M. E. Gursoy, L. Liu, S. Truex, L. Yu, and W. Wei, "Utility-aware synthesis of differentially private and attack-resilient location traces," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pp. 196–211, 2018.
- [120] M. E. Gursoy, V. Rajasekar, and L. Liu, "Utility-optimized synthesis of differentially private location traces," in *2020 Second IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pp. 30–39, IEEE, 2020.
- [121] D. Mir and R. N. Wright, "A differentially private estimator for the stochastic kronecker graph model," in *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, pp. 167–176, 2012.
- [122] A. Sala, X. Zhao, C. Wilson, H. Zheng, and B. Y. Zhao, "Sharing graphs using differentially private graph models," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pp. 81–98, 2011.
- [123] Y. Wang and X. Wu, "Preserving differential privacy in degree-correlation based graph generation," *Transactions on data privacy*, vol. 6, no. 2, p. 127, 2013.
- [124] W. Lu and G. Miklau, "Exponential random graph estimation under differential privacy," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 921–930, 2014.
- [125] Q. Xiao, R. Chen, and K.-L. Tan, "Differentially private network data release via structural inference," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 911–920, 2014.
- [126] Z. Jorgensen, T. Yu, and G. Cormode, "Publishing attributed social graphs with formal privacy guarantees," in *Proceedings of the 2016 international conference on management of data*, pp. 107–122, 2016.
- [127] X. Chen, S. Mauw, and Y. Ramírez-Cruz, "Publishing community-preserving attributed social graphs with a differential privacy guarantee," *Proceedings on Privacy Enhancing Technologies*, vol. 2020, no. 4, pp. 131–152, 2020.

- [128] D. Metzler and W. B. Croft, "A markov random field model for term dependencies," in *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 472–479, 2005.
- [129] V. Bindschaedler and R. Shokri, "Synthesizing plausible privacy-preserving location traces," in *2016 IEEE Symposium on Security and Privacy (SP)*, pp. 546–563, IEEE, 2016.
- [130] R. McKenna, D. Sheldon, and G. Miklau, "Graphical-model based estimation and inference for differential privacy," in *International Conference on Machine Learning*, pp. 4435–4444, PMLR, 2019.
- [131] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.
- [132] M. Sklar, "Fonctions de repartition a dimensions et leurs marges," *Publ. inst. statist. univ. Paris*, vol. 8, pp. 229–231, 1959.
- [133] E. Bouy , V. Durrleman, A. Nikeghbali, G. Riboulet, and T. Roncalli, "Copulas for finance—a reading guide and some applications," Available at SSRN 1032533, 2000.
- [134] T. Bedford and R. M. Cooke, "Probability density decomposition for conditionally dependent random variables modeled by vines," *Annals of Mathematics and Artificial Intelligence*, vol. 32, no. 1, pp. 245–268, 2001.
- [135] S. Demarta and A. J. McNeil, "The t copula and related copulas," *International statistical review*, vol. 73, no. 1, pp. 111–129, 2005.
- [136] D. MacKenzie and T. Spears, "'the formula that killed wall street': The gaussian copula and modelling practices in investment banking," *Social Studies of Science*, vol. 44, no. 3, pp. 393–417, 2014.
- [137] J. Dissmann, E. C. Brechmann, C. Czado, and D. Kurowicka, "Selecting and estimating regular vine copulae and application to financial returns," *Computational Statistics & Data Analysis*, vol. 59, pp. 52–69, 2013.
- [138] W. B. Johnson, "Extensions of lipschitz mappings into a hilbert space," *Contemp. Math.*, vol. 26, pp. 189–206, 1984.
- [139] E. J. Cand s, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on information theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [140] E. Meckes, "Projections of probability distributions: A measure-theoretic dvoretzky theorem," in *Geometric aspects of functional analysis*, pp. 317–326, Springer, 2012.
- [141] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?," *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.
- [142] X. Chu, I. F. Ilyas, and P. Papotti, "Discovering denial constraints," *Proceedings of the VLDB Endowment*, vol. 6, no. 13, pp. 1498–1509, 2013.
- [143] H. Yao and H. J. Hamilton, "Mining functional dependencies from data," *Data Mining and Knowledge Discovery*, vol. 16, no. 2, pp. 197–219, 2008.
- [144] J. Li, J. Liu, H. Toivonen, and J. Yong, "Effective pruning for the discovery of conditional functional dependencies," *The Computer Journal*, vol. 56, no. 3, pp. 378–392, 2013.
- [145] C. Ge, S. Mohapatra, X. He, and I. F. Ilyas, "Kamino: constraint-aware differentially private data synthesis," *Proceedings of the VLDB Endowment*, vol. 14, no. 10, pp. 1886–1899, 2021.
- [146] P. Vines, F. Roesner, and T. Kohno, "Exploring adint: using ad targeting for surveillance on a budget-or-how alice can buy ads to track bob," in *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society*, pp. 153–164, 2017.
- [147] S. Latif, M. Usman, S. Manzoor, W. Iqbal, G. Qadir, G. Tyson, I. Castro, A. Razi, M. N. K. Boulos, A. Weller, et al., "Leveraging data science to combat covid-19: A comprehensive review," *IEEE Transactions on Artificial Intelligence*, vol. 1, no. 1, pp. 85–103, 2020.
- [148] R. Begleiter, R. El-Yaniv, and G. Yona, "On prediction using variable order markov models," *Journal of Artificial Intelligence Research*, vol. 22, pp. 385–421, 2004.
- [149] D. Proserpio, S. Goldberg, and F. McSherry, "A workflow for differentially private graph synthesis," in *Proceedings of the 2012 ACM workshop on Workshop on online social networks*, pp. 13–18, 2012.
- [150] R. Chen, B. C. Fung, S. Y. Philip, and B. C. Desai, "Correlated network data publication via differential privacy," *The VLDB Journal*, vol. 23, no. 4, pp. 653–676, 2014.
- [151] Y. M lle, C. Clifton, and K. B hm, "Privacy-integrated graph clustering through differential privacy," in *EDBT/ICDT Workshops*, vol. 157, 2015.
- [152] H. H. Nguyen, A. Imine, and M. Rusinowitch, "Differentially private publication of social graphs at linear cost," in *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 596–599, IEEE, 2015.
- [153] J. Leskovec and C. Faloutsos, "Scalable modeling of real graphs using kronecker multiplication," in *Proceedings of the 24th international conference on Machine learning*, pp. 497–504, 2007.
- [154] D. F. Gleich and A. B. Owen, "Moment-based estimation of stochastic kronecker graph parameters," *Internet Mathematics*, vol. 8, no. 3, pp. 232–256, 2012.
- [155] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat, "Systematic topology analysis and generation using degree correlations," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, pp. 135–146, 2006.
- [156] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*, pp. 265–284, Springer, 2006.
- [157] K. Nissim, S. Raskhodnikova, and A. Smith, "Smooth sensitivity and sampling in private data analysis," in *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pp. 75–84, 2007.
- [158] T. A. Snijders, P. E. Pattison, G. L. Robins, and M. S. Handcock, "New specifications for exponential random graph models," *Sociological methodology*, vol. 36, no. 1, pp. 99–153, 2006.
- [159] A. Caimo and N. Friel, "Bayesian inference for exponential random graph models," *Social networks*, vol. 33, no. 1, pp. 41–55, 2011.
- [160] A. Clauset, C. Moore, and M. E. Newman, "Hierarchical structure and the prediction of missing links in networks," *Nature*, vol. 453, no. 7191, pp. 98–101, 2008.
- [161] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *The journal of chemical physics*, vol. 21, no. 6, pp. 1087–1092, 1953.
- [162] J. J. Pfeiffer III, S. Moreno, T. La Fond, J. Neville, and B. Gallagher, "Attributed graph models: Modeling network structure with correlated attributes," in *Proceedings of the 23rd international conference on World wide web*, pp. 831–842, 2014.
- [163] F. Chung and L. Lu, "The average distances in random graphs with given expected degrees," *Proceedings of the National Academy of Sciences*, vol. 99, no. 25, pp. 15879–15882, 2002.
- [164] J. J. Pfeiffer, T. L. Fond, S. Moreno, and J. Neville, "Fast generation of large scale social networks while incorporating transitive closures," *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, pp. 154–165, 2012.
- [165] G. Acs, L. Melis, C. Castelluccia, and E. De Cristofaro, "Differentially private mixture of generative neural networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 6, pp. 1109–1121, 2018.
- [166] Q. Chen, C. Xiang, M. Xue, B. Li, N. Borisov, D. Kaarfar, and H. Zhu, "Differentially private data generative models," *arXiv preprint arXiv:1812.02274*, 2018.
- [167] N. C. Abay, Y. Zhou, M. Kantarcioglu, B. Thuraisingham, and L. Sweeney, "Privacy preserving synthetic data release using deep learning," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 510–526, Springer, 2018.
- [168] D. Lee, H. Yu, X. Jiang, D. Rogith, M. Gudala, M. Tejani, Q. Zhang, and L. Xiong, "Generating sequential electronic health records using dual adversarial autoencoder," *Journal of the American Medical Informatics Association*, vol. 27, no. 9, pp. 1411–1419, 2020.
- [169] U. T. Tantipongpipat, C. Waites, D. Boob, A. A. Siva, and R. Cummings, "Differentially private synthetic mixed-type data generation for unsupervised learning," in *2021 12th International Conference on Information, Intelligence, Systems & Applications (IISA)*, pp. 1–9, IEEE, 2021.
- [170] A. Torfi, E. A. Fox, and C. K. Reddy, "Differentially private synthetic medical data generation using convolutional gans," *Information Sciences*, vol. 586, pp. 485–500, 2022.
- [171] L. Fan and A. Pokkunuru, "Dpnet: Differentially private network traffic synthesis with generative adversarial networks," in *IFIP Annual Conference on Data and Applications Security and Privacy*, pp. 3–21, Springer, 2021.
- [172] P.-H. Lu and C.-M. Yu, "Poster: A unified framework of differentially private synthetic data release with generative adversarial network," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 2547–2549, 2017.
- [173] Y. Liu, J. Peng, J. James, and Y. Wu, "Ppgan: Privacy-preserving generative adversarial network," in *2019 IEEE 25th international conference on parallel and distributed systems (ICPADS)*, pp. 985–989, IEEE, 2019.
- [174] L. Xie, K. Lin, S. Wang, F. Wang, and J. Zhou, "Differentially private generative adversarial network," *arXiv preprint arXiv:1802.06739*, 2018.
- [175] L. Frigerio, A. S. de Oliveira, L. Gomez, and P. Duverger, "Differentially private generative adversarial networks for time series, continuous, and discrete open data," in *IFIP International Conference on ICT Systems Security and Privacy Protection*, pp. 151–164, Springer, 2019.
- [176] B. K. Beaulieu-Jones, Z. S. Wu, C. Williams, R. Lee, S. P. Bhavnani, J. B. Byrd, and C. S. Greene, "Privacy-preserving generative deep neural networks support clinical data sharing," *Circulation: Cardiovascular Quality and Outcomes*, vol. 12, no. 7, p. e005122, 2019.
- [177] G. Astolfi and K. David, "Generating tabular data using generative adversarial networks with differential privacy," in *Conference of European Statisticians*, 2021.
- [178] S. Wang, C. Rudolph, S. Nepal, M. Grobler, and S. Chen, "Part-gan: privacy-preserving time-series sharing," in *International Conference on Artificial Neural Networks*, pp. 578–593, Springer, 2020.
- [179] X. Zhang, S. Ji, and T. Wang, "Differentially private releasing via deep generative model (technical report)," *arXiv preprint arXiv:1801.01594*, 2018.
- [180] J. Yoon, J. Jordon, and M. van der Schaar, "PATE-GAN: Generating synthetic data with differential privacy guarantees," in *International Conference on Learning Representations*, 2019.
- [181] Y. Long, B. Wang, Z. Yang, B. Kailkhura, A. Zhang, C. Gunter, and B. Li, "G-pate: Scalable differentially private data generator via private aggregation of teacher discriminators," in *Advances in Neural Information Processing Systems (M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, eds.)*, vol. 34, pp. 2965–2977, Curran Associates, Inc., 2021.
- [182] B. Wang, F. Wu, Y. Long, L. Rimanic, C. Zhang, and B. Li, "Datalens: Scalable privacy preserving training via gradient compression and aggregation," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pp. 2146–2168, 2021.
- [183] T. Cao, A. Bie, A. Vahdat, S. Fidler, and K. Kreis, "Don't generate me: Training differentially private generative models with sinkhorn divergence," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12480–12492, 2021.
- [184] M. Vinaroz, M.-A. Charusaie, F. Harder, K. Adamczewski, and M. J. Park, "Hermite polynomial features for private data generation," in *International Conference on Machine Learning*, pp. 22300–22324, PMLR, 2022.
- [185] F. Harder, M. J. Asadabadi, D. J. Sutherland, and M. Park, "Differentially private data generation needs better features," *arXiv preprint arXiv:2205.12900*, 2022.
- [186] J.-W. Chen, C.-M. Yu, C.-C. Kao, T.-W. Pang, and C.-S. Lu, "Dpgen: Differentially private generative energy-guided network for natural image synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8387–8396, 2022.
- [187] T. Dockhorn, T. Cao, A. Vahdat, and K. Kreis, "Differentially Private Diffusion Models," *arXiv:2210.09929*, 2022.
- [188] J. Ho, A. Jain, and P. Abbeel, "Denoisising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.

- [189] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *International Conference on Learning Representations*, 2021.
- [190] P.-H. Lu, P.-C. Wang, and C.-M. Yu, "Empirical evaluation on synthetic data generation with generative adversarial network," in *Proceedings of the 9th International Conference on Web Intelligence, Mining and Semantics*, pp. 1–6, 2019.
- [191] S.-C. Li, B.-C. Tai, and Y. Huang, "Evaluating variational autoencoder as a private data release mechanism for tabular data," in *2019 IEEE 24th Pacific Rim International Symposium on Dependable Computing (PRDC)*, pp. 198–198, IEEE, 2019.
- [192] D. Arpit, S. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. Courville, Y. Bengio, et al., "A closer look at memorization in deep networks," in *International conference on machine learning*, pp. 233–242, PMLR, 2017.
- [193] J. Hayes, L. Melis, G. Danezis, and E. D. Cristofaro, "Logan: Membership inference attacks against generative models," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, pp. 133–152, 2019.
- [194] B. Hilprecht, M. Härterich, and D. Bernau, "Monte carlo and reconstruction membership inference attacks against generative models," *Proc. Priv. Enhancing Technol.*, vol. 2019, no. 4, pp. 232–249, 2019.
- [195] Y. Xu, S. Mukherjee, X. Liu, S. Tople, R. M. Dodhia, and J. M. L. Ferres, "Mace: A flexible framework for membership privacy estimation in generative models," *Transactions on Machine Learning Research*, 2022.
- [196] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [197] M. Park, J. Foulds, K. Choudhary, and M. Welling, "Dp-em: Differentially private expectation maximization," in *Artificial Intelligence and Statistics*, pp. 896–904, PMLR, 2017.
- [198] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, "Modeling tabular data using conditional gan," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [199] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*, vol. 4. Springer, 2006.
- [200] G. Hinton, Ö. Vinyals, J. Dean, et al., "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, vol. 2, no. 7, 2015.
- [201] N. Papernot, M. Abadi, U. Erlingsson, I. Goodfellow, and K. Talwar, "Semi-supervised knowledge transfer for deep learning from private training data," in *International Conference on Learning Representations*, 2017.
- [202] N. Papernot, S. Song, I. Mironov, A. Raghunathan, K. Talwar, and U. Erlingsson, "Scalable private learning with PATE," in *International Conference on Learning Representations*, 2018.
- [203] T. Elsken, J. H. Metzen, and F. Hutter, "Neural architecture search: A survey," *J. Mach. Learn. Res.*, vol. 20, p. 1997–2017, jan 2019.
- [204] A. Genevay, G. Peyré, and M. Cuturi, "Learning generative models with sinkhorn divergences," in *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617, PMLR, 2018.
- [205] T. Salimans, H. Zhang, A. Radford, and D. Metaxas, "Improving GANs using optimal transport," in *International Conference on Learning Representations*, 2018.
- [206] G. Patrini, R. Van den Berg, P. Forre, M. Carioni, S. Bhargav, M. Welling, T. Genewein, and F. Nielsen, "Sinkhorn autoencoders," in *Uncertainty in Artificial Intelligence*, pp. 733–743, PMLR, 2020.
- [207] C. Villani and C. Villani, "The wasserstein distances," *Optimal Transport: Old and New*, pp. 93–111, 2009.
- [208] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the 34th International Conference on Machine Learning (D. Precup and Y. W. Teh, eds.)*, vol. 70 of *Proceedings of Machine Learning Research*, pp. 214–223, PMLR, 06–11 Aug 2017.
- [209] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trounev, and G. Peyré, "Interpolating between optimal transport and mmd using sinkhorn divergences," in *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690, PMLR, 2019.
- [210] A. Smola, A. Gretton, L. Song, and B. Schölkopf, "A hilbert space embedding for distributions," in *Algorithmic Learning Theory: 18th International Conference, ALT 2007, Sendai, Japan, October 1–4, 2007. Proceedings 18*, pp. 13–31, Springer, 2007.
- [211] A. Berlinet and C. Thomas-Agnan, *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- [212] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola, "A kernel method for the two-sample-problem," *Advances in neural information processing systems*, vol. 19, 2006.
- [213] B. Schölkopf, A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural computation*, vol. 10, no. 5, pp. 1299–1319, 1998.
- [214] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," *Advances in neural information processing systems*, vol. 20, 2007.
- [215] M. Hermite, *Sur un nouveau développement en série des fonctions*. Imprimerie de Gauthier-Villars, 1864.
- [216] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [217] C. N. d. Santos, Y. Mroueh, I. Padhi, and P. Dognin, "Learning implicit generative models by matching perceptual features," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4461–4470, 2019.
- [218] A. Krizhevsky, V. Nair, and G. Hinton, "Cifar-10 (canadian institute for advanced research),"
- [219] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [220] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in neural information processing systems*, vol. 32, 2019.
- [221] Y. Song and S. Ermon, "Improved techniques for training score-based generative models," *Advances in neural information processing systems*, vol. 33, pp. 12438–12448, 2020.
- [222] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.
- [223] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, Y. Shao, W. Zhang, B. Cui, and M.-H. Yang, "Diffusion models: A comprehensive survey of methods and applications," *arXiv preprint arXiv:2209.00796*, 2022.
- [224] A. Hyvärinen and P. Dayan, "Estimation of non-normalized statistical models by score matching," *Journal of Machine Learning Research*, vol. 6, no. 4, 2005.
- [225] U. Grenander and M. I. Miller, "Representations of knowledge in complex systems," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 56, no. 4, pp. 549–581, 1994.
- [226] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965.
- [227] F. Yu, Y. Zhang, S. Song, A. Seff, and J. Xiao, "Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop," *arXiv preprint arXiv:1506.03365*, 2015.
- [228] T. Karras, M. Aittala, T. Aila, and S. Laine, "Elucidating the design space of diffusion-based generative models," *arXiv preprint arXiv:2206.00364*, 2022.
- [229] B. D. Anderson, "Reverse-time diffusion equation models," *Stochastic Processes and their Applications*, vol. 12, no. 3, pp. 313–326, 1982.
- [230] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [231] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [232] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *Advances in neural information processing systems*, vol. 30, 2017.
- [233] C. Li, H. Liu, C. Chen, Y. Pu, L. Chen, R. Henaou, and L. Carin, "Alice: Towards understanding adversarial learning for joint distribution matching," *Advances in Neural Information Processing Systems*, vol. 30, pp. 5495–5503, 2017.
- [234] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [235] J. Lee and C. Clifton, "How much is enough? choosing ϵ for differential privacy," in *Information Security: 14th International Conference, ISC 2011, Xi'an, China, October 26–29, 2011. Proceedings 14*, pp. 325–340, Springer, 2011.
- [236] M. G. Kendall, "Rank correlation methods," 1948.
- [237] B. Zhang and P. Cao, "Classification of high dimensional biomedical data based on feature selection using redundant removal," *PLOS ONE*, vol. 14, pp. 1–19, 04 2019.
- [238] M. J. Smith, J. E. Geach, R. A. Jackson, N. Arora, C. Stone, and S. Courteau, "Realistic galaxy image simulation via score-based generative models," *Monthly Notices of the Royal Astronomical Society*, vol. 511, no. 2, pp. 1808–1818, 2022.
- [239] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *International Conference on Learning Representations*, 2019.
- [240] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro, "Video-to-video synthesis," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [241] H. Bo, S. H. Ding, B. Fung, and F. Iqbal, "Er-ae: Differentially private text generation for authorship anonymization," *arXiv preprint arXiv:1907.08736*, 2019.
- [242] X. Li, F. Tramer, P. Liang, and T. Hashimoto, "Large language models can be strong differentially private learners," in *International Conference on Learning Representations*, 2022.
- [243] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 3–18, IEEE, 2017.
- [244] M. Nasr, S. Song, A. Thakurta, N. Papernot, and N. Carlin, "Adversary instantiation: Lower bounds for differentially private machine learning," in *2021 IEEE Symposium on security and privacy (SP)*, pp. 866–882, IEEE, 2021.
- [245] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," *arXiv preprint arXiv:1812.00984*, 2018.
- [246] V. Cheng, V. M. Suriyakumar, N. Dullerud, S. Joshi, and M. Ghassemi, "Can you fake it until you make it? impacts of differentially private synthetic data on downstream classification fairness," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 149–160, 2021.
- [247] G. Ganey, B. Oprisnu, and E. De Cristofaro, "Robin hood and matthew effects: Differential privacy has disparate impact on synthetic data," in *International Conference on Machine Learning*, pp. 6944–6959, PMLR, 2022.
- [248] O. Giles, K. Hosseini, G. Mingas, O. Strickson, L. Bowler, C. R. Smith, H. Wilde, J. N. Lim, B. Mateen, K. Amarasinghe, et al., "Faking feature importance: A cautionary tale on the use of differentially-private synthetic data," *arXiv preprint arXiv:2203.01363*, 2022.

TABLE 4: A reference table for the candidate algorithms of PPDS divided by data and task. SM=statistical methods, DM=DL-based methods.

Data Type	Applicable Task	Method	Approaches	Most Cited		Most Recent	
				Reference	Methodology/ Principle	Reference	Methodology/ Principle
Tabular data	Task-independent	SM	[91, 97–102] [105, 106, 108–110]	PrivBayes [97]	Marginal-based	PrivMRF [102]	Marginal-based
		DM	[172–175, 177] [178, 180, 181] [165–167, 169] [48, 184]	PATE-GAN [180]	Minimax training	DP-MERF [48]	Feature alignment
	Statistical queries	SM	[90, 93, 103, 104]	MWEM [90]	Query-based	RAP [104]	Query-based
	Clustering/nearest neighbors	SM	[89, 107]	PrivateProjection [107]	Projection-based	Dppro [89]	Projection-based
	Network traffic synthesis	DM	[171]	DPNeT [171]	Minimax training	DPNeT [171]	Minimax training
Health record synthesis	DM	[168, 170, 176]	RDP-CGAN [170]	Minimax training	RDP-CGAN [170]	Minimax training	
Trajectory data	Task-independent	SM	[111–120]	PrefixTree [111]	Tree-based	OptaTrace [120]	Distribution-based
Graph data	Task-independent	SM	[121–127]	dK [122]	dK-Graph	C-AGM [127]	AGM
Image data	Task-independent	DM	[95, 96, 173, 179] [48, 166, 181–187]	DPGAN [174]	Minimax training	DPDM [187]	Stochastic simulation

Appendix A. Method Selection in Real-World Scenarios

Given the massive amount of literature, one natural question from *practitioners* is how to efficiently locate the desired algorithm in real-world scenarios. Here we provide some guidelines for method selection.

A real-world application of PPDS is mostly determined by the target *data type* and *task*. As such, these two factors can be used to filter algorithms and obtain the candidate solutions as a first step for practitioners. In light of this observation, we present a scenario-oriented reference table (Table 4), which classifies the papers that we cover in Section 3 and 4 by data types and applicable tasks. Moreover, for each sublist of approaches, we identify the paper that has received the highest number of citations according to Google Scholar as “Most-cited” and the latest paper as “Fresh”. We also identify the methodologies used in the most-cited and fresh papers. With Table 4, practitioners can easily locate the set of possible candidates as well as the more competent ones within based on their need in real-world scenarios. Moreover, practitioners can refer to Table 1 and Table 2 for the detailed properties of these candidates with the help of Table 4.

Appendix B. Meta-Review

B.1. Summary

This is an SoK paper on the topic of privacy-preserving data synthesis (PPDS). The paper provides a comprehensive literature review, classifying existing data synthesis methods as either statistical methods or deep methods. The paper also proposes a “master recipe” that abstracts the components of a method and experimentally compares existing deep methods.

B.2. Scientific Contributions

- Other: SoK paper
- Independent Confirmation of Important Results with Limited Prior Research

B.3. Reasons for Acceptance

- 1) The paper provides an overview of synthetic data generation and state of the art techniques for each type of data, which may be useful for newcomers to the area. The paper also describes open problems and future directions, which may be useful in shaping future research.
- 2) The paper evaluates six deep PPDS approaches on image synthesis using two standard datasets and identify DP-MERF as an all-purpose best approach. The paper also provides guidelines for practitioners who seek to use PPDS.
- 3) The reviewers ask that prior to publication, you update your main paper to motivate why there are no experiments on tabular data. For instance, you can indicate that the parts of the paper related to tabular data generation methods have already been studied and evaluated empirically in prior work, and provide the citation.