

Differentially Private Normalizing Flows for Privacy-Preserving Density Estimation

Chris Waites
Stanford University
waites@stanford.edu

Rachel Cumminngs
Columbia University
rac2239@columbia.edu

ABSTRACT

Normalizing flow models have risen as a popular solution to the problem of density estimation, enabling high-quality synthetic data generation as well as exact probability density evaluation. However, in contexts where individuals are directly associated with the training data, releasing such a model raises privacy concerns. In this work, we propose the use of normalizing flow models that provide explicit differential privacy guarantees as a novel approach to the problem of privacy-preserving density estimation. We evaluate the efficacy of our approach empirically using benchmark datasets, and we demonstrate that our method substantially outperforms previous state-of-the-art approaches. We additionally show how our algorithm can be applied to the task of differentially private anomaly detection.

CCS CONCEPTS

• Security and privacy → Privacy protections; Privacy protections; Data anonymization and sanitization; • Information systems → Database design and models; • Theory of computation → Machine learning theory; Machine learning theory; Theory of database privacy and security.

KEYWORDS

differential privacy, density estimation, normalizing flows, anomaly detection

ACM Reference Format:

Chris Waites and Rachel Cumminngs. 2021. Differentially Private Normalizing Flows for Privacy-Preserving Density Estimation. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21), May 19–21, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3461702.3462625>

1 INTRODUCTION

The task of density estimation requires constructing an estimate of an unknown probability density function, given observed data. This density estimate can then be used to perform a variety of relevant analysis tasks, including log likelihood evaluation and synthetic data generation. In settings involving sensitive data, the construction and subsequent release of such an estimate could

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIES '21, May 19–21, 2021, Virtual Event, USA
© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8473-5/21/05...\$15.00
<https://doi.org/10.1145/3461702.3462625>

potentially leak private information. Without a rigorous privacy guarantee, nothing prevents a model from memorizing a row in the training set, assigning disproportionate density to a point, or any other vulnerability due to arbitrary analysis of the learned parameters. Since density estimation remains a task of interest to the modeling community, continued attention is required to develop privacy-preserving methods for density estimation.

Differential privacy [16] has emerged as the predominant privacy notion in the context of statistical data analysis. At a high level, differentially private analyses limit the extent to which the distribution of outputs can change due to the inclusion or exclusion of any one individual from the analysis. Algorithms which adhere to this notion exhibit a number of desirable properties, including privacy guarantees which hold regardless of the auxiliary information an adversary may have and composition of privacy guarantees across multiple analyses. Hence differential privacy acts as a compelling gold standard in the design of privacy-preserving analyses.

Tools for density estimation have held longstanding interest due to their versatility. Their ability to address a wide range of distributional learning tasks is precisely why the existence of an accurate and privacy-preserving density estimation is surprising. For example, privately constructing such a model implicitly yields a differentially private approach to anomaly detection—a task of substantial investigation [3, 21, 41]—as an immediate application of likelihood inference. In addition, given that density estimators often enable efficient sampling, such a model would yield a method for privacy-preserving synthetic data generation. This task in particular has been of longstanding interest to the privacy community [47] as it addresses many of the limitations imposed by the query-release model [14] by allowing large numbers of arbitrary analyses. Privately generating a synthetic dataset only incurs a fixed privacy cost during the training process; all subsequent queries on the synthetic data are automatically differentially private due to the privacy notion's post-processing guarantee, so the privacy cost does not scale with the number of downstream analyses performed.

Normalizing flow models are an attractive approach to the task of density estimation due to their empirical ability to approximate arbitrary, high-dimensional distributions. These models approach the task of density estimation via a transformation on a chosen base density by a sequence of invertible, non-linear transformations, enabling density querying on the resulting distribution via an application of the change-of-variables formula. Approaches to density estimation in this manner include: Non-linear Independent Components Estimation (NICE) [9], Real NVP [10], Glow [33], and Masked Autoregressive Flows (MAF) [44]. Until this work, it was an open question whether normalizing flow models could be constructed in a differentially private manner to handle the task of privacy-preserving density estimation, combining the rigorous

guarantees of differential privacy with the strong empirical performance exhibited by normalizing flows.

In this work we propose the use of normalizing flow models trained in a differentially private manner as a novel approach to the task of privacy-preserving density estimation. We provide an algorithm (DP-NF, Algorithm 1 in Section 3) that privately optimizes the model parameters via gradient descent using DP-SGD [1], which adds Gaussian noise to clipped gradient updates ensure differential privacy. Additionally, we achieve tighter privacy guarantees than established in previous work [1] via composition with the recently introduced notion of Gaussian differential privacy [11]. We apply this optimization to the parameters of a Masked Autoregressive Flow [44], our primary architecture of consideration, and achieve empirical results (Section 4) which convincingly outperform previous approaches. Further, we show that our algorithm can be applied to solve the problem of differentially private anomaly detection (Section 5), and show that it leads to better true/false positive rates than existing private methods.

1.1 Related Work

Gaussian mixture models (GMMs) are known to be a particularly strong density estimation tool [43] since they are a *universal approximator of densities* – that is, they are able to approximate any density function arbitrarily well given a sufficient number of components [36]. They approach the task of density estimation by modeling the data distribution as a weighted sum of Gaussian distributions. The first differentially private algorithm for learning the parameters of a Gaussian mixture model comes from the work of [40], which uses their *sample-and-aggregate* framework to convert non-private algorithms into private algorithms, applied to the task of learning mixtures of Gaussians. However, their approach exhibits strong assumptions on the range of the parameter space and assumes a uniform mixture of spherical Gaussians. Follow-up work of [31] proposes a modernized approach which improves upon the sample complexity of the aforementioned work and removes the strong a priori bounds on the parameters of the mixture components, although it makes the assumption that the components of the mixture are well-separated.

There has also been work in learning the parameters of a Gaussian mixture model through differentially private variants of the expectation maximization (EM) algorithm. One notable instance of this is DPGMM [49], which achieves a privacy guarantee at each iteration of EM through the application of calibrated Laplace noise to the estimated model parameters following each maximization step. These individual privacy guarantees are then combined into an overall privacy guarantee via sequential composition, i.e., by taking the sum of privacy parameters in each iteration. The work of [45] introduces DP-EM, a general framework for privacy-preserving optimization via expectation maximization. Their approach follows a conceptually similar idea of applying either calibrated Laplace or Gaussian noise to the model parameters at the end of each EM iteration. They apply this method to learning mixtures of Gaussians, henceforth referred to as DP-MoG, and they demonstrate significantly better privacy guarantees through composition via the moments accountant and zero-concentrated differential privacy (zCDP) [6]. Given that their work makes no notable assumptions

about the task and provides an empirical evaluation of their method, this is the most comparable approach to our own. As such, it is used as a baseline in our experimental results.

In addition, we take note of more classical approaches to the task of privacy-preserving density estimation. One of the simplest yet most widely used methods for density estimation is through the use of histograms, and previous work [8?] has investigated their private estimation. Unfortunately, such an approach scales poorly with the dimension and complexity of the distribution while asserting an unrealistic discretization of the space. Kernel density estimation is another closely related approach, often characterized as the smooth analog to the classical discrete histogram. The work of [26] proposes a method for privately querying the density of such an estimator through the addition of calibrated Gaussian noise. As a non-parametric approach, it has the drawback that it requires storage of the entire dataset at test time to enable querying (proving impractical for large-scale datasets) while still degrading similarly with dimension.

There have also been a number of deep learning based approaches to generative modeling which vary in their relevance. Although work of this nature technically allows for both sampling and likelihood evaluation, it does not allow for *exact* likelihood inference as is the case for mixtures of Gaussians and normalizing flows. There is also expansive literature concerning differentially private approaches to training Generative Adversarial Networks, yet these methods are strictly limited to sampling and do not provide a straightforward approach to likelihood inference.

Finally, we include a brief overview of the extensive literature concerning density estimation via normalizing flows. One important subset are those characterized by *coupling layers*: transformations which partition the dimensions of its input and map them in a way that retains invertibility and a tractable Jacobian. This includes Non-linear Independent Components Estimation (NICE) [9], as well as its subsequent generalization Real NVP [10]. Another notable approach, Glow [33], makes use of such coupling layers while also proposing the use of an invertible weight matrix decomposition to generalize the notion of permutation layers. Alternatively, some make use of *autoregressive transformations*, which are transformations that utilize the chain rule of probability to represent a joint distribution as a product of its conditionals. Such models include Masked Autoregressive Flow (MAF) [44], a generalization of Real NVP optimized for density estimation, as well as its closely related Inverse Autoregressive Flow [34] optimized for variational inference, among others [20, 28, 42].

2 PRELIMINARIES

2.1 Normalizing Flows

Let $p(\cdot)$ be the probability density function characterizing an unobservable distribution of interest, and let $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$ be n observed i.i.d. samples from this distribution. The task of density estimation is to find an approximation of $p(\cdot)$ via some model $p_{\theta}(\cdot)$ given \mathbf{X} . In the context of normalizing flows, this model is characterized by a prior distribution $q(\cdot)$, chosen to exhibit a simple and tractable density (e.g., the spherical multivariate Gaussian distribution), and a sequence of K bijective functions $f_{\theta} = f_1 \circ f_2 \circ \dots \circ f_K$, parameterized fully by θ . The function f_{θ} acts as a transformation

between the prior distribution $q(\cdot)$ and the approximated distribution $p_\theta(\cdot)$.

Given such a model, it can be used to efficiently sample $\mathbf{x} \sim p_\theta$ by first sampling $\mathbf{z} \sim q$ and then transforming the sample as $\mathbf{x} = f_\theta(\mathbf{z})$. If p_θ is a good approximation of p , then this generative process gives an efficient (approximate) oracle for sampling from the unknown distribution.

Since f_θ is invertible, one can also perform exact likelihood evaluation on observed points from the data distribution via the change of variables formula, as follows:

$$\begin{aligned}\log p_\theta(\mathbf{x}) &= \log q(f_\theta^{-1}(\mathbf{x})) + \log \left| \det \left(\frac{\partial f_\theta^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right| \\ &= \log q(f_\theta^{-1}(\mathbf{x})) + \sum_{i=1}^K \log \left| \det \left(\frac{\partial f_i^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right) \right|.\end{aligned}$$

Finding a good approximation p_θ is achieved through optimization of θ to minimize the negative log likelihood of the observed dataset:

$$\mathcal{L}(\theta) = -\frac{1}{n} \sum_{i=1}^n \log p_\theta(\mathbf{x}^{(i)}). \quad (1)$$

In practice, one will typically find the MLE $\theta^* = \arg \min_\theta \mathcal{L}(\theta)$ using some non-convex optimization method, such as stochastic gradient descent.

2.2 Differential Privacy

Differential privacy [16] has become the gold standard for ensuring the privacy of statistical analyses applied to sensitive databases. At a high level, it ensures that changing a single entry in the database will have only a small effect on the distribution of analysis results.

DEFINITION 1 ([16]). *A randomized algorithm $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ satisfies (ϵ, δ) -differential privacy (DP) if for any two input database $D, D' \in \mathcal{D}$ that differ in a single entry and for any subset of outputs $\mathcal{S} \subseteq \mathcal{R}$, it satisfies,*

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

One common algorithmic approach for achieving differential privacy is adding noise that scales with the *sensitivity* of the function being evaluated, which is the maximum change in the function's value that can result from changing a single data point. Differentially private algorithms are robust to *post-processing*, meaning that any data-independent function of a differentially private output retains the same privacy guarantee, and they enjoy *composition*, meaning that the privacy parameters degrade gracefully as additional analyses are performed on the dataset. The simplest version of composition is that the privacy parameters ϵ and δ "add up" over multiple analyses, although stronger versions of composition are also used.

Differentially Private Stochastic Gradient Descent (DP-SGD, presented formally in Algorithm 3 in Appendix A.4) was introduced by [1] as a method for private non-convex optimization. At each step t , DP-SGD subsamples¹ a small set of data points and uses this batch to compute a gradient update. To achieve a differential privacy guarantee, DP-SGD adds mean-zero Gaussian noise to the

¹The original algorithm of [1] does this via Poisson subsampling, but it can also be done via uniform subsampling while retaining a privacy guarantee [48].

average of the per-example gradients. The standard deviation of this noise is scaled with the sensitivity of the gradient estimation. Since this is unbounded, the per-example gradients are first clipped to ensure that the ℓ_2 -norm is at most some input parameter C , thus bounding the sensitivity, and then adds noise which scales with C .

[1] also introduced the *moments accountant*, which provides tight privacy composition across multiple gradient update steps in DP-SGD. To describe the moments accountant, given an algorithm \mathcal{M} and two neighboring datasets D, D' , first we denote the privacy loss of a particular outcome o as $L^{(o)} = \log(\Pr(\mathcal{M}_D = o) / \Pr(\mathcal{M}_{D'} = o))$. The moments accountant calculates a privacy budget by bounding the moments of the privacy loss random variable $L^{(o)}$. That is, if we consider the log of the moment generating function (MGF) of the privacy loss random variable evaluated at λ , i.e., $\alpha_{\mathcal{M}}(\lambda; \mathcal{D}, \mathcal{D}') = \log \mathbb{E}_{o \sim \mathcal{M}_D} [e^{\lambda L^{(o)}}]$, the worst case over all neighboring databases $\max_{\mathcal{D}, \mathcal{D}'} \alpha_{\mathcal{M}}(\lambda; \mathcal{D}, \mathcal{D}')$ composes linearly across multiple mechanisms (Theorem 2.1 [1]) and allows for conversion to an associated (ϵ, δ) -differential privacy guarantee through the relation $\delta = \min_\lambda \exp[\alpha_{\mathcal{M}}(\lambda) - \lambda \epsilon]$. Follow up work of [5] introduced NoisySGD, which followed the same algorithmic structure but analyzed privacy composition under Gaussian differential privacy [11]. For the purpose of this work it is sufficient to simply note the associated benefits of analysis under Gaussian differential privacy: it naturally lends itself to composition under subsampling, allows for analytically tractable expressions of the privacy guarantees of NoisySGD, while providing a slightly tighter overall privacy bound than that achieved by the moments accountant. Further details are provided in Appendix A.

3 DIFFERENTIALLY PRIVATE NORMALIZING FLOWS

In this section we introduce our algorithm for differentially private density estimation via normalizing flows, DP-NF, presented in Algorithm 1. It is based on the DP-SGD algorithm of [1], which is a differentially private method for performing stochastic gradient descent. We also briefly discuss performance improvements using data-dependent initialization of normalization layers and using a differentially private estimate of the distribution to act as a prior, both of which are explored further Appendix B. We emphasize that our primary technical contribution is not in the design of these algorithms, but rather the novel application of these tools to the problem of differentially private density estimation in a way that yields substantial performance over prior work, as demonstrated by our empirical results in Section 4.

3.1 DP-NF Algorithm

Training a normalizing flow model corresponds to minimizing the loss function in Equation (1): $\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^N \log p_\theta(\mathbf{x}^{(i)})$. This loss function is non-convex when applied to the optimization of a non-linear normalizing flow model, and hence optimization is typically performed via gradient descent on θ . To make this training private in Algorithm 1, we update θ using the DP-SGD algorithm of [1] described in Section 2.2, with some subtle yet important augmentations to the standard minibatch gradient descent process to allow for an explicit privacy guarantee, in accordance with DP-SGD.

Algorithm 1 DP-NF, Differentially private density estimation via normalizing flows

```

1: Input: Dataset  $X = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}\}$ , initialized parameters  $\theta$ , learning rate  $\eta$ , batch size  $b$ , noise scale  $\sigma$ , upper-bound on  $\ell_2$  norm of per-example gradient  $C$ , training privacy budget  $\epsilon$ , training privacy tolerance  $\delta$ , privacy accountant  $P$ .
2:  $t \leftarrow 1$ 
3: while  $P(t, b/n, \sigma, C, \delta) < \epsilon$  do
4:   Take a uniformly random subsample  $I_t \subseteq \{1, \dots, n\}$  with batch size  $b$ .
5:   for  $i \in I_t$  do
6:      $\mathbf{g}_t^{(i)} \leftarrow \nabla_{\theta} - \log p_{\theta}(\mathbf{x}^{(i)})$ 
7:      $\bar{\mathbf{g}}_t^{(i)} \leftarrow \mathbf{g}_t^{(i)} / \max\{1, \|\mathbf{g}_t^{(i)}\|_2/C\}$ 
8:   end for
9:    $\theta \leftarrow \theta - \eta \cdot \frac{1}{b} (\sum_i \bar{\mathbf{g}}_t^{(i)} + \mathcal{N}(\mathbf{0}, \sigma^2 C^2 \mathbf{I}))$ 
10:   $t \leftarrow t + 1$ 
11: end while
12: Output  $\theta$ 

```

First, batches are sampled via uniform subsampling (Line 4). That is, each possible batch of size b has equal likelihood of being chosen (as opposed to repeatedly shuffling the dataset and taking equally sized partitions of the dataset, which is often preferred in practice). Second, rather than computing the gradient with respect to the entire batch, the gradient with respect to each individual data point is calculated, clipped to have maximum ℓ_2 norm C , averaged, then added with a randomly sampled Gaussian noise vector (Lines 6-9).

Algorithm 1 also requires a *privacy accountant* to be specified as input. This privacy accountant will dynamically track the ϵ privacy loss incurred by composition over all gradient update steps as a function of the training parameters, and will halt the algorithm once a pre-specified budget is reached. A privacy accountant $P(t, b/n, \sigma, C, \delta)$ takes in the round t of training, the sampling probability b/n of a single point (here a batch of size b is sampled uniformly from a set of n data points), the noise scale σ that is added to preserve privacy, the bound C on the ℓ_2 norm of each gradient, and the privacy parameter δ . At every time step, the privacy accountant maintains the current ϵ privacy budget that has been expended until round t given the input parameters. Common choices for this accountant include the moments accountant (MA) [1] or composition via Gaussian differential privacy (GDP) [11]. In our experiments in Section 4, we yield preferable results using a GDP privacy accountant.

In summary, DP-NF in Algorithm 1 is a modified version of DP-SGD, instantiated to train a normalizing flow model with the analyst's choice of privacy accountant.

The privacy guarantees of DP-NF follow as an immediate corollary from those of DP-SGD [1] when instantiated with the moments accountant, and from NoisySGD [5] when instantiated with the Gaussian differential privacy accountant.

THEOREM 1. *DP-NF is (ϵ, δ) -differentially private.*

3.2 DP-NF Extensions

In practice, one will find that many deep learning models (including the normalizing flow models used in our experiments) are much better optimized using adaptive learning rate optimization schemes. Given this, we found significant benefit in using a direct extension to DP-SGD which applies noisy gradients to the model according to the Adam [32] optimizer. Both methods achieve identical privacy guarantees given that computation of the first and second moments of the noisy gradients are merely deterministic data-independent functions of them. Thus they differ only in the post-processing of the noisy gradients, and the privacy guarantees are unchanged.

Two further extensions of Algorithm 1 are proposed below, which may provide substantial improvements to empirical performance.

Data-Dependent Initialization of Normalization Layers. Intermediate normalization layers such as activation normalization [33] have been proposed as a means to improve the stability of normalizing flow models. Activation normalization is characterized by a feature-wise offset and scaling of inputs by a learned set of parameters \mathbf{b} and \mathbf{w} , i.e., $(\mathbf{x}^{(i)} - \mathbf{b})/\mathbf{w}$. In practice, these parameters are typically set via data-dependent initialization [46] by setting \mathbf{b} and \mathbf{w} as the per-feature means and standard deviations observed throughout a forward pass of a sampled batch of data. These parameters can also be estimated privately, e.g., by applying the Laplace Mechanism [16] to the clipped mean and standard deviation, thus allowing for data-dependent initialization of these normalization layers. For more details, see Appendix B.1.

Differentially Private Data-Dependent Priors. Section 2.1 suggested the analyst choose a data-independent prior q , such as the multivariate spherical Gaussian. However, recent work suggests that modest improvements in empirical results can be achieved through the use of more complex priors, such as a mixture of Gaussians [44], or by fitting a Gaussian mixture model to the data [30]. A natural privacy-preserving approach would be to first use DP-MoG [45] with privacy budget (ϵ_1, δ_1) to estimate a prior, and then refine the prior using DP-NF with privacy budget (ϵ_2, δ_2) to yield an encompassing normalizing flow model. This process would be $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differentially private, and may yield preferable results in settings where the distribution is highly discontinuous, but also locally non-linear. For more details, see Appendix B.2.

4 EXPERIMENTAL RESULTS

In this section we present experimental results demonstrating the empirical performance of our approach, evaluating our algorithm on a variety of real and synthetic datasets on varying tasks. In the main body we focus our evaluation on a single dataset (the Life Science dataset [12], described next), although refer to Appendix C for all additional results on other real and synthetic datasets.

In all our experiments on the Life Science dataset, we used $\delta = 1.52 \times 10^{-5}$. Our baseline method for comparison [45] used $\delta = 1.00 \times 10^{-4}$. However, this corresponds to $\delta = O(1/n)$ on this dataset, which is typically deemed unacceptably large in the privacy community. Instead, our choice of $\delta = 1.52 \times 10^{-5} = 1/n^{1.1}$, which is sublinear in the size of the database. Smaller values of δ would not change our qualitative results, nor would they substantially change our quantitative results.

Table 1: Average test log likelihood for varying privacy budgets ϵ . Error bars denote standard deviation over ten independent cross-validation splits. Bolded results denote best performing model for a given ϵ .

| Life Science | | $\delta = 1.52 \times 10^{-5}$ | $\epsilon = 0.50$ | $\epsilon = 1.00$ | $\epsilon = 2.00$ | $\epsilon = 4.00$ |
|---------------|-----------------------------------|--------------------------------|-----------------------------------|------------------------------------|------------------------------------|-------------------|
| DP-NF (GDP) | 8.90 ± 0.18 | | 9.41 ± 0.12 | 10.20 ± 0.09 | 10.77 ± 0.24 | |
| DP-NF (MA) | 7.37 ± 0.17 | | 8.67 ± 0.12 | 9.32 ± 0.09 | 10.09 ± 0.18 | |
| DP-MoG (MA) | | 2.30 ± 0.27 | 5.09 ± 0.20 | 8.87 ± 0.06 | 9.10 ± 0.06 | |
| DP-MoG (zCDP) | | -8.93 ± 0.49 | -0.16 ± 0.37 | 2.95 ± 0.28 | 5.48 ± 0.18 | |

4.1 Datasets, Implementation, & Setup

The Life Science dataset is a standard density estimation benchmark dataset from the UCI machine learning repository [12] containing 26,733 real-valued records of dimension 10. This dataset was used in the original evaluation of our baseline model [45]. Results using additional datasets are presented in Appendix C.

Experiments were run on a machine with 2 CPUs, 13 GB RAM, and a single NVIDIA Tesla K80 GPU, and took on the order of half an hour to five hours to run in wall-time, depending on the number of iterations and the dimensionality of the dataset. Models were implemented in the Jax [4] deep learning framework, and used privacy accounting implementations from TensorFlow Privacy [23].

Hyperparameter Search and Model Selection. Reported privacy budgets in our results correspond only to the training of each model, and does not include privacy loss from hyperparameter search and model selection. We chose not to select hyperparameters in a privacy-preserving manner because this was not the focus of our contribution and because it was not done in our baseline method.² It was generally observed that choices in network structure itself had relatively negligible impacts on results. We found that training parameters such as the gradient clipping bound and batch size had a much more substantial impact on model performance, which is consistent with observations made in [1].

Model Architecture. The architecture of the model used in our experiments was a variant of a Masked Autoregressive Flow (MAF) [44] composed of a repeated sequence of five blocks, each containing a MADE [22] layer, a reversal layer, and an optional activation normalization layer. Models were optimized via Adam, with default parameters of $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Further details of training parameters and procedures are given in Appendix C.3.

4.2 Empirical Performance of DP-NF

We implemented our algorithm for differentially private normalizing flows on the Life Science dataset (and other datasets as described in Appendix C), and evaluated our performance against the baseline of DP-MoG [45] for a variety of quantitative and qualitative metrics related to density estimation tasks.

4.2.1 Quantitative Evaluation: Expected Log Likelihood. Arguably the most foundational metric for density estimators is the expected

²These can be done privately. For example, [25] provides discrete optimization methods that can be used for private hyperparameter search over discrete model architectures. [2] uses ReportNoisyMax [19] for private model selection. Some work has also been done to account for high-performance models without having to spend a significant privacy budget [7, 35].

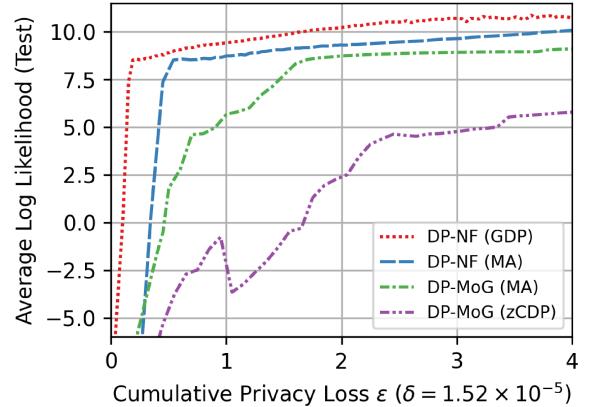


Figure 1: Average log likelihood (higher is better) across ten independent cross-validation splits as a function of the cumulative privacy loss ϵ . DP-MoG was configured to use the Gaussian mechanism with 3 components, as per the original work. DP-NF composed with GDP (as well as MA for fair comparison).

log likelihood they assign to held out test points. Figure 1 presents average log likelihood assigned to a held out test set under DP-NF and the baseline method DP-MoG [45] as a function of ϵ . We divided the dataset into 10 pairs of training (90%) and test sets (10%), and reported the average test log likelihood per data point across the 10 independent trials. Better methods should assign higher log likelihood for points in the held out test set since these points were indeed sampled from the underlying distribution of interest. We found that DP-NF reliably assigned much higher likelihoods to holdout data than that of DP-MoG for identical privacy budgets, across a variety of privacy accountant methods.

The privacy guarantees of DP-NF proved quite practical, providing substantial privacy improvements over DP-MoG for the same model performance. For example, DP-NF matched the peak performance of DP-MoG (achieved around $\epsilon \approx 4$) for only an expenditure of $\epsilon \approx 0.5$. These results are also listed in Table 1 with error bars showing standard deviation across 10 independent runs.

[45] showed performance of DP-MoG under several different privacy accountant methods, with the moments accountant of [1] providing the best performance. We compared DP-NF using the moments accountant for fair comparison, and using the novel Gaussian

differential privacy (GDP) accountant of [5]. Figure 1 and Table 1 show that DP-NF outperforms DP-MoG for all privacy accountant methods considered for either model, emphasizing that while the GDP accountant does provide some benefit, the vast majority of the performance improvements come from the DP-NF method itself. The benefits of using the GDP accountant are further explored in the appendix.

4.2.2 Quantitative Evaluation: Downstream Machine Learning Tasks. Next we further evaluate the quality of our model by measuring the performance of downstream machine learning models trained on its generated synthetic data. A proper method for evaluating the strength of density estimation approaches is through the quality of their synthetic data, as measured by the ability to train a machine learning model that performs well on future, real data.

To perform this evaluation, we trained DP-NF and DP-MoG (along with their non-private variants for reference) to learn the distribution of the training data, and then used these models to generate a synthetic dataset. We then trained a simple regressor— k -nearest neighbors with default library settings ($k = 3$)—on each synthetically generated dataset and evaluated their in predicting a target value on real, held out Life Science data. The Life Science dataset does not have an immediately associated prediction task, as it is primarily used as solely a density estimation benchmark. To artificially construct a prediction task, we simply chose to isolate the last column to act as a label and treated the remaining nine columns as features.

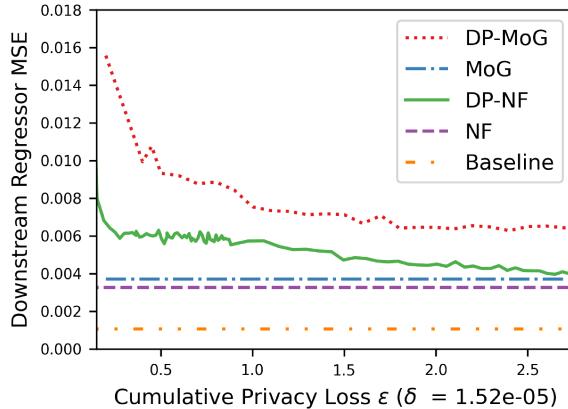


Figure 2: Mean squared error of regressor (k -nearest neighbors with $k = 3$) on real, held out test data when trained on synthetically generated data by various approaches. Baseline refers to training the regressor on the real data.

Figure 2 shows the mean squared error attained by each regressor as a function of the privacy expenditure of the data generation approach. Horizontal lines denote that the approach is $\epsilon = \infty$, i.e., non-private. “Baseline” refers to training the regressor directly on the training dataset provided to the density estimator. Upon inspection, we find that DP-NF generates data of convincingly higher quality than that of our comparison method DP-MoG for

all values of ϵ . In addition, for higher values of ϵ , DP-NF converges to the quality achievable by a non-private MoG, whereas DP-MoG hits an apparent plateau well before this point.

4.2.3 Qualitative Evaluations. Figure 3 shows that DP-NF provides a qualitative increase in sample quality under visualization. It presents dimension-wise histograms of synthetically generated features for three features of the Life Science dataset, using DP-NF (left column) and DP-MoG (right column) for comparison. (See Figure 11 in Appendix C for dimension-wise histograms of all 10 features.) Both methods used $\epsilon = 0.6$ and $\delta = 1.52 \times 10^{-5}$. In every plot, the synthetic data in orange is superimposed over the real data in blue. We qualitatively see that for nearly all ten features, the distribution of data generated by DP-NF closely matches that of the real data, while DP-MoG was relatively unable to replicate regions of concentrated density for certain dimensions. This could be due to the fact that for a fixed number of components, the DP-MoG model is constrained to cover the support of the distribution and must ignore nuanced details. Normalizing flow models, on the other hand, have heightened expressiveness over traditional statistical methods like Gaussian mixture models, and we see that they are able to capture these nuances more readily.

As another qualitative evaluation of sample visualization, Figure 4 shows the density of synthetic data generated by each model when projected to two dimensional space via PCA, for varying ϵ values. The top row shows DP-NF, the bottom row shows DP-MoG, and the right figure shows the real data. In all plots, lighter pixels correspond to regions of higher density, and dark pixels indicate lower density. We see that DP-NF is better able to capture some of the observable qualities exhibited in the real data, for example the gradual compression of density to the left of the distribution.

5 APPLICATION: DIFFERENTIALLY PRIVATE ANOMALY DETECTION

Our DP-NF algorithm can be used as a tool for differentially private anomaly detection. Given a density estimator, a straightforward approach to anomaly detection is through a simple likelihood thresholding mechanism. For a given point, to determine whether it is in-distribution or out-of-distribution, we can simply return a binary value which denotes in-distribution if the log likelihood assigned to the point by the model is above some empirically derived threshold T , and out-of-distribution otherwise. For the purposes of our experiments, we assume that such a threshold is easily estimated, i.e., by selecting the value for T which optimizes anomaly detection performance on a public test set. In the private setting, we can approach this task in a privacy-preserving manner by training either DP-NF or DP-MoG on the dataset. By the post-processing property of differential privacy, we can make arbitrarily many anomaly detection queries to the privately trained model while incurring no privacy loss beyond what is incurred during training.

In Figure 5, we illustrate the efficacy of our approach in performing anomaly detection under this likelihood thresholding mechanism. We randomly generated points that were uniformly distributed around the tails of the test dataset, i.e., between the 5th and 30th percentiles and the 70th and 95th percentiles dimension-wise. The total number of synthetically generated anomalies was equal to the total number of test points. Figure 5 shows ROC curves of both

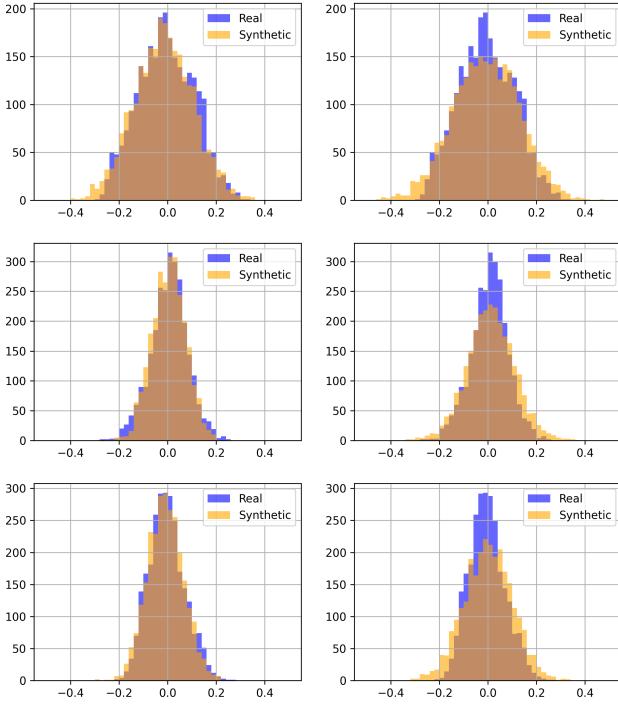


Figure 3: Dimension-wise histograms of synthetically generated Life Science data, superimposed over real data, for $\epsilon = 0.6$ and $\delta = 1.52 \times 10^{-5}$. Left Column: DP-NF. Right Column: DP-MoG. Note DP-NF’s ability to capture regions of concentrated density, whereas DP-MoG struggles in this respect.

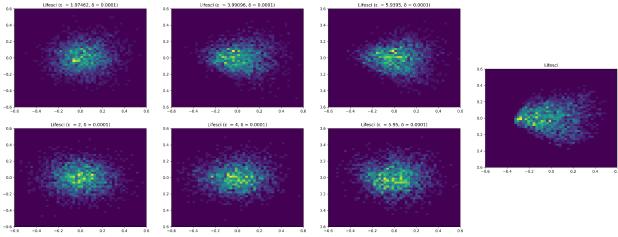


Figure 4: Synthetically generated Life Science data for $\epsilon = 2, 4$, and 6 , projected to two dimensions via PCA. Top row: DP-NF. Bottom row: DP-MoG. Right: Real data. Note the compression to the left of the distribution of real data that is captured by DP-NF as ϵ increases, but not present in the synthetic data generated by DP-MoG.

private and non-private methods for this binary prediction problem. That is, it shows the tradeoff in the true positive rate and the false positive rate in predicting in-distribution or out-of-distribution correctly for varying selections of the likelihood threshold. We observe that DP-NF outperformed the other private method DP-MoG for the same privacy guarantee. Our approach performed comparably to

non-private MoG, and is of course upper-bounded by a non-private normalizing flow.

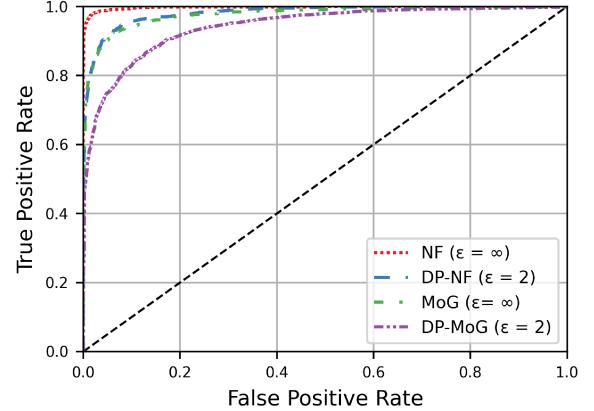


Figure 5: ROC curves displaying true positive rate and false positive rate for private and non-private likelihood threshold models. Privacy expenditure was calculated using the moments accountant with $\delta = 1.52 \times 10^{-5}$.

There are other possible methods for making anomaly detection algorithms differentially private, beyond the approach described above of using DP-NF directly to privately train the model. An alternative approach is to, rather than learning the parameters of the model in a differentially private manner, partition the dataset into k parts and train a non-private density estimator on each part. Then, given a new point of interest, each model casts a vote regarding its belief on the point being in-distribution or out-of-distribution by testing whether their density assigned to the point exceeds T . We then aggregate these votes privately, to ensure that the final prediction is differentially private with respect to the training set. In our algorithm for differentially private anomaly detection (Algorithm 2), we use the Exponential Mechanism [38] for this private aggregation. We note that our overall approach is an instantiation of the *sample-and-aggregate* framework of [40]. This approach is visualized in Figure 6 and presented formally in Algorithm 2.

The privacy guarantees of DP-AD follow as an immediate corollary from those of sample-and-aggregate [40] and the Exponential Mechanism [38].

THEOREM 2. *DP-AD is $(\epsilon, 0)$ -differentially private.*

Although Algorithm 2 is instantiated with the Exponential Mechanism [38], one can opt for any differentially private aggregation method. In contexts requiring answers to a large but bounded number of queries where the number of anomalies is expected to be small in comparison to the total number of queries (an appropriate assumption for many applications), an immediate extension of this approach using the Sparse Vector technique [17, 18] would be natural.

To evaluate this approach, we applied Algorithm 2 to the Half-Moons dataset, which is a synthetic dataset with 30,000 data points,

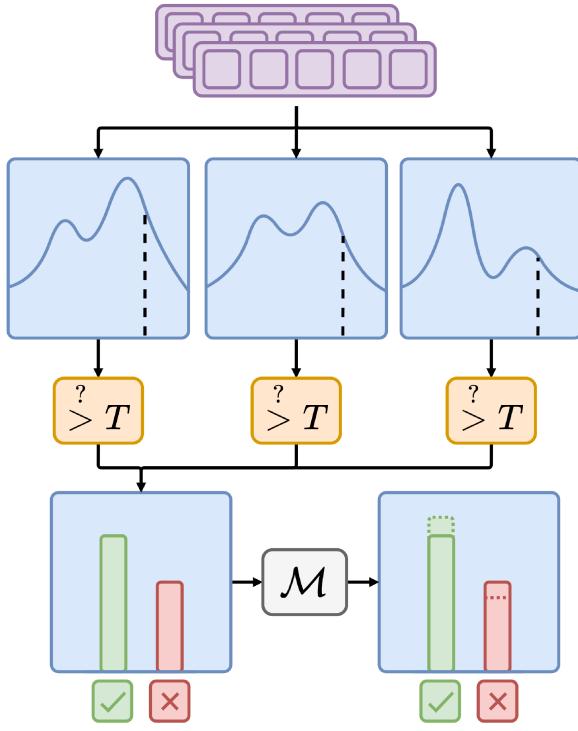


Figure 6: Ensembled anomaly detection framework. In this setting, k non-private density estimators are trained on separate partitions of the training dataset. At inference time new examples (purple) are fed to each density estimator, voting on whether the given example is “in-distribution” or “out-of-distribution” depending on whether it exceeds a known likelihood threshold. Then, we perform a noisy aggregation on these votes to preserve privacy.

Algorithm 2 DP-AD, Differentially private anomaly detection via an ensemble of density estimators

```

1: Input: Dataset  $X = \{x^{(1)}, \dots, x^{(n)}\}$ , example  $x$ , number of partitions  $k$ , likelihood threshold  $T$ , privacy budget  $\epsilon$ .
2:  $X_1, \dots, X_k \leftarrow \text{partition}(X, k)$ 
3:  $\theta_1, \dots, \theta_k \leftarrow \text{train}(X_1, \dots, X_k)$ 
4:  $c \leftarrow 0$ 
5: for  $i \in [k]$  do
6:   if  $p_{\theta_i}(x) > T$  then
7:      $c \leftarrow c + 1$ 
8:   end if
9: end for
10: Sample response as “in-distribution” with probability  $\frac{\exp(\epsilon c/2)}{\exp(\epsilon c/2) + \exp(\epsilon(k-c)/2)}$  and “out-of-distribution” with probability  $\frac{\exp(\epsilon(k-c)/2)}{\exp(\epsilon c/2) + \exp(\epsilon(k-c)/2)}$ 
11: Output response

```

each with two real-valued features. This dataset is visualized in Figure 7, and more details about this synthetic dataset are given in Appendix C.

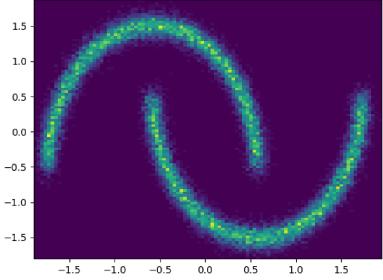


Figure 7: Half-Moons Dataset. Lighter regions correspond to higher density of data points.

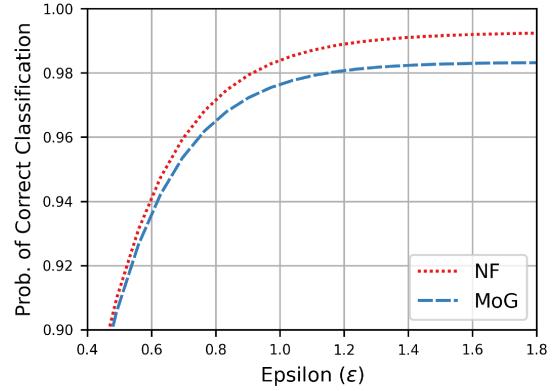
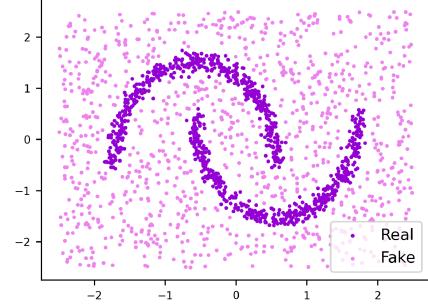


Figure 8: Anomaly detection through private aggregation of votes (either “in-distribution” or “out-of-distribution”) on synthetically generated data. Training set was partitioned into 10 pieces, then each piece was fit to a non-private model, either MoG or NF. Differentially private aggregation was performed via the exponential mechanism. Top: Training data with anomalies; Bottom: Anomaly detection performance.

Figure 8 illustrates the data used and the results of this evaluation. The training data (dark purple) was partitioned into 10 pieces and used to fit a set of independent non-private models (all either MoG or NF). Then, anomalies (light purple) were added to the dataset,

and both real and fake data were fed into the learned models. These data are illustrated in the top image of Figure 8. The Exponential Mechanism was used to perform differentially private aggregation to yield the final ensemble prediction when applied to held out test points. The bottom image of Figure 8 presents the probability of a correct classification between “in-sample” and “out-of-sample” data as a function of ϵ under this scheme. Both approaches were given the threshold that optimized their classification performance. Notice that using normalizing flows in this setting naturally yields a performance improvement over the existing method, which is due in part to the ability of normalizing flows to capture densities which do not adhere to a Gaussian distribution.

The benefit of this ensembled approach is that each individual model can be trained non-privately, which may substantially improve the quality of the learned models. For certain problems, this might be necessary when differentially private optimization is known to degrade performance substantially, for example in high dimensional problems with image data. The associated downside is that the analyst’s privacy budget would now degrades as a function of the number of queries, whereas before the privacy budget was independent of the number of queries. The analyst would have to determine the application-specific tradeoff between costs and benefits of this approach, based on practical constraints imposed by the problem context.

6 CONCLUSION

Privacy is a subject of increasing importance and growing concern. By our work adhering to the framework of differential privacy, one is able to make definitive statements regarding the privacy of participants involved in our analysis. Our results could also be used to enable differentially private synthetic data generation, which would allow data curators to provide privatized synthetic versions of their sensitive or protected datasets, thereby enabling broader access to these data.

In this work, we have demonstrated the efficacy of differentially private normalizing flow models as a novel approach to the task of privacy-preserving density estimation. We have shown the ability of these models to assign high likelihoods to holdout data and generate qualitatively realistic synthetic data, improving on existing state-of-the-art methods. Going forward, there exist several interesting directions for further development. For example, it remains to be seen how normalization layers such as activation normalization, whose parameters are likely disproportionately sensitive to perturbation during differentially private optimization, could be better adapted to such. Further, one might hypothesize that sampling via partitions of a shuffled dataset may yield improved results given more regular sampling of each data point, and associating such a sampling method with rigorous privacy guarantees if possible could yield empirical improvements. Finally, in this study we only considered a particular subset of normalizing flows in existence. Although, many alternative neural density estimators capable of expressing highly discontinuous distributions are in continuous development, including FFJORD [24], Neural Spline Flows [13], Neural Autoregressive Flows [28], and Transformation Autoregressive Networks [42].

ACKNOWLEDGMENTS

R.C. supported in part by NSF grants CNS-1850187 and CNS-1942772 (CAREER), a Mozilla Research Grant, and a JPMorgan Chase Faculty Research Award. Most of this work was completed while R.C. was at Georgia Institute of Technology.

REFERENCES

- [1] Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM Conference on Computer and Communications Security*, CCS ’16, 308–318.
- [2] Beaulieu-Jones, B. K.; Wu, Z. S.; Williams, C.; Lee, R.; Bhavnani, S. P.; Byrd, J. B.; and Greene, C. S. 2019. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes* 12(7): e005122.
- [3] Bittner, D.; Sarwate, A. D.; and Wright, R. 2018. Using Noisy Binary Search for Differentially Private Anomaly Detection. In *Proceedings of the 2nd International Symposium on Cyber Security Cryptography and Machine Learning (CSCML)*, volume 10879, 20–37.
- [4] Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; and Wanderman-Milne, S. 2018. JAX: composable transformations of Python+NumPy programs. URL <http://github.com/google/jax>.
- [5] Bu, Z.; Dong, J.; Long, Q.; and Su, W. J. 2020. Deep Learning with Gaussian Differential Privacy. *Harvard Data Science Review* 2(3).
- [6] Bun, M.; and Steinke, T. 2016. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Proceedings of the 13th Conference on Theory of Cryptography*, TCC ’16, 635–658.
- [7] Chaudhuri, K.; and Vinterbo, S. A. 2013. A Stability-based Validation Procedure for Differentially Private Machine Learning. In *Advances in Neural Information Processing Systems 26*, NIPS ’13, 2652–2660.
- [8] Chawla, S.; Dwork, C.; McSherry, F.; and Talwar, K. 2005. On the Utility of Privacy-Preserving Histograms. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, UAI ’05.
- [9] Dinh, L.; Krueger, D.; and Bengio, Y. 2014. NICE: Non-linear Independent Component Estimation. ArXiv preprint 1410.8516.
- [10] Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2017. Density estimation using Real NVP. In *Proceedings of the International Conference on Learning Representations*, ICLR ’17.
- [11] Dong, J.; Roth, A.; and Su, W. J. 2019. Gaussian Differential Privacy. ArXiv preprint 1905.02383.
- [12] Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. URL <http://archive.ics.uci.edu/ml>.
- [13] Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019. Neural Spline Flows. In *Advances in Neural Information Processing Systems 32*, NeurIPS ’19.
- [14] Dwork, C. 2008. Differential Privacy: A Survey of Results. In *Theory and Applications of Models of Computation (TAMC)*, volume 4978 of *Lecture Notes in Computer Science*, 1–19. Springer Verlag.
- [15] Dwork, C.; and Lei, J. 2009. Differential privacy and robust statistics. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, STOC ’09, 371–380.
- [16] Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC ’06, 265–284.
- [17] Dwork, C.; Naor, M.; Pitassi, T.; and Rothblum, G. N. 2010. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, STOC ’10, 715–724.
- [18] Dwork, C.; Naor, M.; Reingold, O.; Rothblum, G. N.; and Vadhan, S. P. 2009. On the complexity of differentially private data release: Efficient algorithms and hardness results. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, STOC ’09, 381–390.
- [19] Dwork, C.; and Roth, A. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* 9(3–4).
- [20] Fakoor, R.; Chaudhari, P.; Mueller, J.; and Smola, A. J. 2020. TraDE: Transformers for Density Estimation. ArXiv preprint 2004.02441.
- [21] Fan, L.; and Xiong, L. 2013. Differentially Private Anomaly Detection with a Case Study on Epidemic Outbreak Detection. In *Proceedings of the IEEE 13th International Conference on Data Mining Workshops*, 833–840.
- [22] Germain, M.; Gregor, K.; Murray, I.; and Larochelle, H. 2015. MADE: Masked Autoencoder for Distribution Estimation. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 881–889.
- [23] Google. 2018. TensorFlow Privacy. URL <https://github.com/tensorflow/privacy>.
- [24] Grathwohl, W.; Chen, R. T. Q.; Bettencourt, J.; Sutskever, I.; and Duvenaud, D. 2019. FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. In *Proceedings of International Conference on Learning Representations*, ICLR ’19.

[25] Gupta, A.; Ligett, K.; McSherry, F.; Roth, A.; and Talwar, K. 2010. Differentially Private Combinatorial Optimization. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, 1106–1125.

[26] Hall, R.; Rinaldo, A.; and Wasserman, L. 2013. Differential Privacy for Functions and Functional Data. *Journal of Machine Learning Research* 14(1): 703–727.

[27] He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, 1026–1034.

[28] Huang, C.-W.; Krueger, D.; Lacoste, A.; and Courville, A. 2018. Neural Autoregressive Flows. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 2078–2087.

[29] Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML '15, 448–456.

[30] Izmailov, P.; Kirichenko, P.; Finzi, M.; and Wilson, A. G. 2020. Semi-Supervised Learning with Normalizing Flows. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 4615–4630.

[31] Kamath, G.; Sheffet, O.; Singhal, V.; and Ullman, J. 2019. Differentially Private Algorithms for Learning Mixtures of Separated Gaussians. In *Advances in Neural Information Processing Systems 32*, NeurIPS '19.

[32] Kingma, D.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 2015 International Conference on Learning Representations*, ICLR '15.

[33] Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative Flow with Invertible 1x1 Convolutions. In *Advances in Neural Information Processing Systems 31*, NeurIPS '18.

[34] Kingma, D. P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; and Welling, M. 2016. Improved Variational Inference with Inverse Autoregressive Flow. In *Advances in Neural Information Processing Systems 29*, NIPS '16.

[35] Liu, J.; and Talwar, K. 2019. Private Selection from Private Candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC '19, 298–309.

[36] McLachlan, G.; and Basford, K. 1988. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker.

[37] McMahan, H. B.; Andrew, G.; Erlingsson, U.; Chien, S.; Mironov, I.; Papernot, N.; and Kairouz, P. 2018. A General Approach to Adding Differential Privacy to Iterative Training Procedures. ArXiv preprint 1812.06210.

[38] McSherry, F.; and Talwar, K. 2007. Mechanism Design via Differential Privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, 94–103.

[39] Mironov, I. 2017. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium*, CSF '17, 263–275.

[40] Nissim, K.; Raskhodnikova, S.; and Smith, A. 2007. Smooth Sensitivity and Sampling in Private Data Analysis. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, STOC '07, 75–84.

[41] Okada, R.; Fukuchi, K.; and Sakuma, J. 2015. Differentially Private Analysis of Outliers. In *Proceedings of the 2015 European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part II*, ECMLPKDD '15, 458–473.

[42] Oliva, J.; Dubey, A.; Zaheer, M.; Poczos, B.; Salakhutdinov, R.; Xing, E.; and Schneider, J. 2018. Transformation Autoregressive Networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 3898–3907.

[43] Papamakarios, G. 2019. Neural Density Estimation and Likelihood-free Inference. ArXiv preprint 1910.13233.

[44] Papamakarios, G.; Pavlakou, T.; and Murray, I. 2017. Masked Autoregressive Flow for Density Estimation. In *Advances in Neural Information Processing Systems 30*, NIPS '17, 2338–2347.

[45] Park, M.; Foulds, J.; Choudhary, K.; and Welling, M. 2017. DP-EM: Differentially Private Expectation Maximization. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, 896–904.

[46] Salimans, T.; and Kingma, D. P. 2016. Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks. In *Advances in Neural Information Processing Systems 29*, NIPS '16.

[47] Surendra, H. S.; and Mohan, H. S. 2017. A Review Of Synthetic Data Generation Methods For Privacy Preserving Data Publishing. *International Journal of Scientific & Technology Research* 6: 95–101.

[48] Wang, Y.-X.; Balle, B.; and Kasiviswanathan, S. 2019. Subsampled Rényi Differential Privacy and Analytical Moments Accountant. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, AISTATS '19, 1226–1235.

[49] Yuncheng Wu; Yao Wu; Hui Peng; Juru Zeng; Hong Chen; and Cuiping Li. 2016. Differentially private density estimation via Gaussian mixtures model. In *Proceedings of the IEEE/ACM 24th International Symposium on Quality of Service (IWQoS)*, 1–6.

A ADDITIONAL PRIVACY PRELIMINARIES

A.1 Differential Privacy

We achieve differential privacy in Algorithm 1 through the Gaussian Mechanism, which adds mean-zero Gaussian noise to the value of a function f evaluated on the data. The scale of the noise depends on the *sensitivity* of the function. The ℓ_2 -sensitivity of a function f is denoted $\Delta_2 f$, and is the maximum change in the ℓ_2 norm of f if one entry in the database were to be changed. Formally, for $f : \mathbb{N}^{|X|} \rightarrow \mathbb{R}^k$, $\Delta_2 f = \max_{D, D'} \text{neighbors} \|f(D) - f(D')\|_2$. In our case, this function is the computation of the gradient given a sampled batch of data.

THEOREM 3 (GAUSSIAN MECHANISM [19]). *Let $f : \mathbb{N}^{|X|} \rightarrow \mathbb{R}^k$ and let $\mathcal{M}(x, f, \epsilon) = f(x) + (Z_1, \dots, Z_k)$ where Z_i are i.i.d. random variables drawn from $\mathcal{N}(0, \sigma^2)$ and $\sigma > \sqrt{2 \ln(1.25/\delta) \Delta_2 f / \epsilon}$. Then \mathcal{M} is (ϵ, δ) -differentially private.*

Differential privacy *composes*, meaning that the privacy guarantees degrade smoothly as more analyses are performed on the same dataset. The simplest version of privacy composition is that the ϵs and δs "add up" across analyses. Tighter composition bounds are possible, including the approaches outlined in the following two subsections.

A.2 Moments Accountant

The moments accountant [1] was proposed initially as a means for tight composition of the privacy guarantees of DP-SGD. To characterize this analysis, we note the privacy loss associated with a given outcome o , given as $L^{(o)} = \log \frac{\Pr(\mathcal{M}_D=o)}{\Pr(\mathcal{M}_{D'}=o)}$. Further, for two given datasets we define the MGF of this random variable evaluated at some value λ as $\alpha_{\mathcal{M}}(\lambda; \mathcal{D}, \mathcal{D}') = \log \mathbb{E}_{o \sim \mathcal{M}_D} [e^{\lambda L^{(o)}}]$. Finally, the "worst case" upper bound across all possible pairs of datasets is given as $\alpha_{\mathcal{M}}(\lambda) \triangleq \max_{\mathcal{D}, \mathcal{D}'} \alpha_{\mathcal{M}}(\lambda; \mathcal{D}, \mathcal{D}')$.

This notation can be used to give composition guarantees for the privacy parameters across multiple algorithms run on the same dataset.

THEOREM 4 ([1]). *Suppose that an algorithm \mathcal{M} consists of a sequence of adaptive algorithms $\mathcal{M}_1, \dots, \mathcal{M}_k$ where $\mathcal{M}_i : \prod_{j=1}^{i-1} \mathcal{R}_j \times \mathcal{D} \rightarrow \mathcal{R}_i$. Then for any λ , $\alpha_{\mathcal{M}}(\lambda) \leq \sum_{i=1}^k \alpha_{\mathcal{M}_i}(\lambda)$. For any $\epsilon > 0$, \mathcal{M} is (ϵ, δ) -differentially private for $\delta = \min_{\lambda} \exp(\alpha_{\mathcal{M}}(\lambda) - \lambda \epsilon)$.*

This analysis was then later characterized under the framework of Rényi differential privacy [39], a relaxation of (ϵ, δ) -differential privacy which is defined in a manner closely resembling the moments accountant privacy analysis.

DEFINITION 2 $((\alpha, \epsilon)\text{-RDP}$ [39]). *A randomized algorithm $f : \mathcal{D} \rightarrow \mathcal{R}$ is said to have ϵ -Rényi differential privacy of order α , or $(\alpha, \epsilon)\text{-RDP}$ for short, if for any adjacent $D, D' \in \mathcal{D}$ it holds that $D_{\alpha}(f(D) \| f(D')) \leq \epsilon$, where $D_{\alpha}(P \| Q) \triangleq \frac{1}{\alpha-1} \log \mathbb{E}_{x \sim Q} \left(\frac{P(x)}{Q(x)} \right)^{\alpha}$.*

Finally, the privacy analysis performed in [1] regarding DP-SGD assumes the sampling is performed via Poisson subsampling, i.e., each individual example has independent sampling probability q of being included in the batch at each iteration. In some contexts, fixed-sized batches can enable a variety of performance improvements by allowing for compilation. Privacy analysis under uniform

subsampling, where each batch is sampled uniformly across all possible batches of size b , was considered in [48] :

THEOREM 5 (RDP-DP CONVERSION [48]). *For all integers $\alpha \geq 2$, if \mathcal{M} is $(\alpha, \epsilon(\alpha))\text{-RDP}$, then the randomized algorithm applied to a subsampled batch of data without replacement is $(\alpha, \epsilon'(\alpha))\text{-RDP}$ where $\epsilon'(\alpha) \leq \frac{1}{\alpha-1} \log(1 + \gamma^2 \binom{\alpha}{2} \min\{4(e^{\epsilon(2)} - 1), e^{\epsilon(2)} \min\{2, (e^{\epsilon(\infty)} - 1)^2\}\} + \sum_{j=3}^{\alpha} \gamma^j \binom{\alpha}{j} e^{(j-1)\epsilon(j)} \min\{2, (e^{\epsilon(\infty)} - 1)^j\})$.*

A.3 Gaussian Differential Privacy

Gaussian differential privacy (GDP) is a recently proposed relaxation of (ϵ, δ) -differential privacy established in [11], and further expanded upon in the context of deep learning in [5]. This definition exhibits several appealing properties, including simplified analysis under composition and subsampling, derivation of analytically tractable expressions for the privacy guarantees of NoisySGD, while providing a slightly tighter privacy bound than that which is achieved through analysis via the moments accountant. The framework of μ -Gaussian differential privacy acts as the basis for our analysis.

We note that the Gaussian mechanism (Definition 3) $\mathcal{M}(D) = f(D) + \mathcal{N}(0, (c\Delta_2 f / \epsilon)^2)$ is $(\epsilon/c)\text{-GDP}$ [11].³

The overall privacy guarantee corresponding to the k -fold adaptive composition of k mechanisms each satisfying μ_i -GDP will be $\sqrt{\mu_1^2 + \mu_2^2 + \dots + \mu_k^2}$ -GDP. Finally, μ -GDP allows for a conversion to a corresponding (ϵ, δ) -differential privacy guarantee using the fact that an algorithm is μ -GDP if and only if it is $(\epsilon, \delta(\epsilon))$ -differentially private for all $\epsilon \geq 0$, where $\delta(\epsilon) = \Phi(-\frac{\epsilon}{\mu} + \frac{\mu}{2}) - e^{\epsilon} \Phi(-\frac{\epsilon}{\mu} - \frac{\mu}{2})$ and $\Phi(\cdot)$ is the cumulative density function of the Normal distribution.

Figure 9 shows that GDP privacy accounting gives substantial improvements over the moments accountant method when used as the privacy accountant method in DP-NF. For each number of iterations, GDP accounting yields a lower ϵ privacy value.

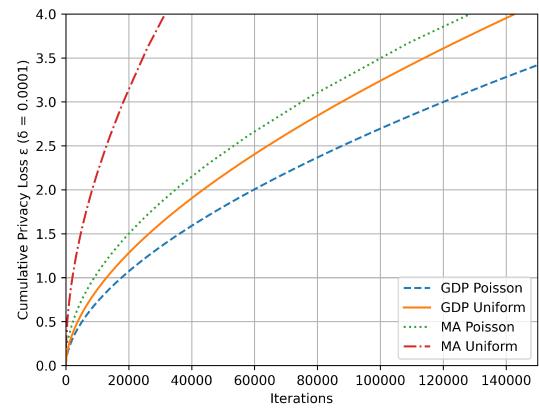


Figure 9: Cumulative privacy loss ϵ given Life Science training parameters ($q = b/n = 100/21384 = 4.676 \times 10^{-3}$, $\sigma = 2.1$, $\delta = 1.00 \times 10^{-4}$) as a function of training iterations.

³Further detail concerning the privacy guarantees achieved when batches are subsampled is given in Section 2 of [5].

Algorithm 4 DP-NF-INIT, data dependent initialization of activation normalization layers

- 1: **Input:** Dataset $X = \{x^{(1)}, \dots, x^{(n)}\}$, transformation f (e.g. MADE [22]), number of layers K , initialization privacy budget ϵ , initialization privacy tolerance δ , data-independent parameter initialization method R (e.g. He initialization [27]).
- 2: $\{\theta_1, \dots, \theta_K\} \leftarrow R()$
- 3: **for** $k = 1, \dots, K$ **do**
- 4: $X \leftarrow \text{clip}(f_{\theta(k)}(X), \tilde{c})$
- 5: $\mathbf{b}^{(k)} \leftarrow \mu(X) + \text{Lap}(\frac{2\sqrt{4K \ln(1/\delta)} \Delta \hat{\mu}}{\epsilon})$
- 6: $\mathbf{w}^{(k)} \leftarrow \sigma(X) + \text{Lap}(\frac{2\sqrt{4K \ln(1/\delta)} \Delta \hat{\sigma}}{\epsilon})$
- 7: $X \leftarrow (X - \mathbf{b}^{(k)}) / \mathbf{w}^{(k)}$
- 8: **end for**
- 9: **Output** concatenation of $\theta^{(1)}, \mathbf{b}^{(1)}, \mathbf{w}^{(1)}, \dots, \theta^{(K)}, \mathbf{b}^{(K)}, \mathbf{w}^{(K)}$

A.4 DP-SGD

Algorithm 3 is a formal description of DP-SGD of [1], from which our approach is based. Specifically, Algorithm 1 is an instantiation of DP-SGD, substituting negative log likelihood for the generic loss function in Algorithm 3.

Algorithm 3 DP-SGD, differentially private stochastic gradient descent [1]

- 1: **Input:** Dataset $X = \{x^{(1)}, \dots, x^{(n)}\}$, loss function $\mathcal{L}(\theta)$, learning rate η_t , noise multiplier σ , batch size L , gradient norm bound C .
- 2: **Initialize** θ_0 randomly
- 3: **for** $t \in [T]$ **do**
- 4: Take a Poisson random subsample L_t with per-example probability L/N .
- 5: **Compute gradient**
- 6: for each $i \in L_t$, compute $\mathbf{g}_t(x_i) \leftarrow \nabla_{\theta_t} \mathcal{L}(\theta_t, x_i)$
- 7: **Clip gradient**
- 8: $\tilde{\mathbf{g}}_t(x_i) \leftarrow \mathbf{g}_t(x_i) / \max(1, \|\mathbf{g}_t(x_i)\|_2 / C)$
- 9: **Add noise**
- 10: $\tilde{\mathbf{g}}_t(x_i) \leftarrow \frac{1}{L} (\sum_i \tilde{\mathbf{g}}_t(x_i) + \mathcal{N}(\mathbf{0}, \sigma^2 C^2 \mathbf{I}))$
- 11: **Descend**
- 12: $\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{\mathbf{g}}_t(x_i)$
- 13: **end for**
- 14: **Output** θ_T and compute overall (ϵ, δ) using a privacy accounting method.

B DP-NF EXTENSIONS

B.1 Data-Dependent Initialization of Normalization Layers

Intermediate normalization layers such as batch normalization [29] and activation normalization [33] have been shown to improve the stability of normalizing flow models. In our context, batch normalization is incompatible with our approach since that batch statistics are shared when computing the forward pass of the layer, precluding the ability to calculate truly independent per-example gradients

as required by NoisySGD. Activation normalization is more appropriate in our setting since no such batch statistics are calculated. Activation normalization is characterized by an offset and scaling of its inputs feature-wise by a learned set of parameters \mathbf{b} and \mathbf{w} , i.e. $\mathbf{y}^{(i)} \leftarrow (x^{(i)} - \mathbf{b}) / \mathbf{w}$. In practice, typically these parameters are set via data-dependent initialization [46] by computing a forward pass on a sampled batch of data and setting \mathbf{b} and \mathbf{w} to be the per-feature means and standard deviations of the inputs it had observed respectively. Since these statistics are calculated directly from the data, this approach is not privacy-preserving.

One potential approach to making differentially private activation normalization is to privatize these statistics using the Laplace Mechanism [16]. This approach is outlined in Algorithm 4, where $\text{clip}(X, \tilde{c})$ clips the values of X to be in the range $[-\tilde{c}/2, \tilde{c}/2]$, $\mu(X)$ computes the feature-wise mean of X , $\sigma(X)$ computes the feature-wise standard deviation of X , and R is some data-independent parameter initialization method which maps standardized inputs to standardized outputs in expectation, e.g., He initialization [27].

We note that Algorithm 4 is far from the only approach for differentially private activation normalization. If the analyst has some domain knowledge about appropriate ranges of these parameters, she could use the differentially private Propose-Test-Release framework [15] to first normalize X and then add noise proportionally to her proposed (and tested) sensitivity. While this approach seems practical for the outer layers, it is unlikely that an analyst would have numerical intuition for appropriate parameter values in all layers (especially when K is large). Thus even though the noise addition scheme described in Algorithm 4 may seem naive, it is likely the most practical approach.

Although the utility of activation normalization layers is quite evident, the original work [33] proposing such layers provided little evidence to support the idea that data-dependent initialization yielded statistically significant improvements over a default initialization scheme, i.e., $\mathbf{b} \leftarrow \mathbf{0}$ and $\mathbf{w} \leftarrow \mathbf{1}$. In our experiments, we observed little distinction in contexts where the input data was assumed to be standardized and parameters were initialized to maintain variance between layers. Despite this, we include the approach for completeness for potential future contexts where data-dependent initialization of such parameters deems necessary.

B.2 DP-MoG as a Prior

Thus far it has been assumed that we use the spherical multivariate Gaussian distribution to act as a prior for our model. Although, naturally any distribution could act as such a prior as long as it exhibits a tractable density function. For example, simply extending the single standardized Gaussian to a mixture of Gaussians has been shown [44] to exhibit modest performance improvements. This mixture could be fit to the data [30] a priori as well.

Hence, a natural extension to our proposed approach would be to fit DP-MoG first with privacy budget (ϵ_1, δ_1) to act as a prior, and then to refine this prior by training a sequence of nonlinear bijective functions with privacy budget (ϵ_2, δ_2) to yield an encompassing normalizing flow model. This yields a worst-case $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differential private guarantee by sequential composition. Although, this guarantee is easily improved by composing these privacy guarantee under some alternative privacy definition (e.g. RDP or GDP)

Table 2: Test negative log likelihood (lower is better) on Pinwheel dataset (Figure 10) for varying privacy budgets ϵ , composed via the moments accountant. From top to bottom: Standard DP-NF with a single spherical Gaussian as a prior, DP-NF with a GMM prior of 5 components (fit non-privately), DP-NF with a GMM prior of 5 components (learned privately for $\epsilon = 0.2, \delta = 0$), DP-MoG with 5 components, DP-MoG with 10 components. When a privacy budget is expended on the prior, this was included when calculating the overall cost indicated by the column headings.

| Pinwheel $\delta = 3.70 \times 10^{-5}$ | $\epsilon = 1.50$ | $\epsilon = 2.50$ | $\epsilon = 3.50$ | $\epsilon = 4.50$ |
|---|-------------------|-------------------|-------------------|-------------------|
| DP-NF _{spherical Gaussian prior} | 2.80 | 2.78 | 2.74 | 2.63 |
| DP-NF _{non-private GMM prior} | 3.05 | 2.40 | 1.91 | 1.87 |
| DP-NF _{private GMM prior, $\epsilon = 0.2$} | 2.77 | 2.70 | 2.62 | 2.47 |
| DP-MoG ₅ | 2.75 | 2.76 | 2.76 | 2.76 |
| DP-MoG ₁₀ | 2.76 | 2.88 | 2.77 | 2.77 |

before subsequently converting to a corresponding (ϵ, δ) -DP guarantee. One might hypothesize that this approach would yield preferable results in contexts where the distribution at hand is composed of several discontinuous components, while exhibiting locally non-linear density within each component.

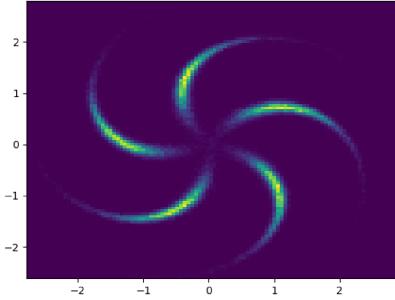


Figure 10: Pinwheel Dataset

To capture this context, we evaluate the efficacy of this approach on the Pinwheel dataset, as illustrated in Figure 10 (further details of this dataset are given in Appendix C). The Pinwheel dataset is a common density estimation benchmark consisting of a number of disconnected components with nonlinear boundary. A Gaussian mixture model would naturally have difficulty approximating such a distribution for a small number of components, while classical normalizing flow models with a single standardized Gaussian as a prior might have difficulty expressing its discontinuous density.

As shown in Table 2, using a trained GMM to act as a prior can aid in performance. First, we note that both DP-NF and DP-MoG demonstrate difficulty in achieving negative log likelihoods lower than 2.65-2.70, even when the number of components for DP-MoG is increased. DP-NF with DP-MoG as a prior modestly outperforms both alternatives used in isolation. Additionally, if one assumes that a GMM prior of five components fit to the population can be assumed to be public, one achieves dramatic performance improvements over all methods.

C ADDITIONAL RESULTS AND TRAINING DETAILS

In this section, we provide additional results on the performance of DP-NF, compared to the baseline mechanism of DP-MoG, as evaluated on several real and synthetic datasets. These datasets are summarized in Table 3. Details of the real and synthetic datasets are respectively given in Sections C.1 and C.2. Further details of hyperparameter training for all experiments are given in Section C.3.

C.1 Additional Results on Real Datasets

Our analysis aims to cover a range of datasets composed of both synthetic and real-world datasets. In this section, we provide short descriptions of each dataset used for evaluation, and provide results of DP-NF evaluated on these datasets.

Table 3: Dimensionality D and number of examples n for each dataset.

| Dataset | Real | D | n |
|--------------|-----------|-----|---------|
| Life Science | Real | 10 | 26,733 |
| Gowalla | Real | 2 | 100,000 |
| Power | Real | 6 | 100,000 |
| Gaussians | Synthetic | 2 | 30,000 |
| Half-Moons | Synthetic | 2 | 30,000 |
| Pinwheel | Synthetic | 2 | 30,000 |

Life Science. The Life Science dataset is a density estimation benchmark dataset from the UCI machine learning repository [12] used in our baseline [45] in their evaluation of DP-MoG. It contains 26,733 real-valued records of dimension 10 characterizing the principle components of measurements made in a variety of chemical and biological experiments.

Power. The Power dataset is a density estimation benchmark dataset from the UCI machine learning repository [12] used in much of the normalizing flow literature [24, 44]. It contains measurements of electric power consumption in a household over a period of 47 months, and was preprocessed according to the description given in [44].

Gowalla. The Gowalla dataset contains the locations in terms of longitude and latitude of the social network’s users’ check-ins.

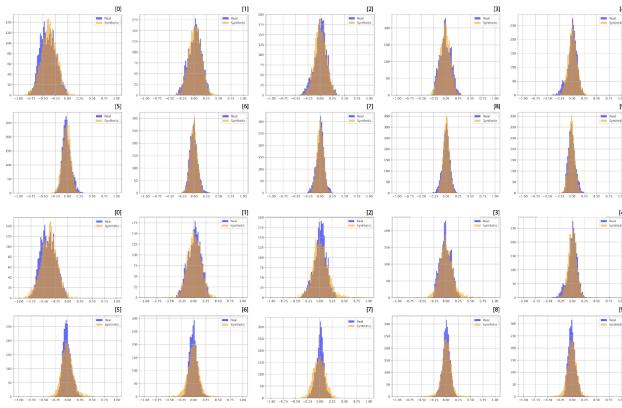


Figure 11: Dimension-wise histograms of synthetically generated Life Science data, superimposed over real data, for $\epsilon = 0.5$ and $\delta = 1.52 \times 10^{-5}$. Dimensions are given in order from 0 to 9 left-to-right, top-to-bottom. Top two rows: DP-NF. Bottom two rows: DP-MoG. Note that synthetic data from DP-NF represents the real data well, while DP-MoG is relatively unable to capture concentrated regions of density in the real data.

The total number of points is 1,256,384, which was reduced to 100,000 via a random sample. It was used in the evaluation of our baseline [45], but applied to the task of k -means clustering rather than learning the components of a Gaussian mixture model.

We evaluated the performance of DP-NF and the baseline DP-MoG for comparison on these real-world datasets. Results on the Life Science dataset are given in the body of the paper in Table 1. Results on the Power and Gowalla datasets are given in Table 4 below.

Table 4: Average test log likelihood for varying privacy budgets ϵ . Error bars denote standard deviation over ten independent cross-validation splits. DP-MoG_x refers to a Gaussian mixture of x components. Omitted entries occur when model training was early-stopped because it had already converged at lower ϵ values.

| Power | | | |
|--------------------------------|-------------------|-------------------|-------------------|
| $\delta = 1.11 \times 10^{-5}$ | $\epsilon = 1.25$ | $\epsilon = 1.75$ | $\epsilon = 2.25$ |
| DP-NF (GDP) | -2.12 ± 0.29 | - | - |
| DP-NF (MA) | -3.43 ± 0.19 | -2.47 ± 0.13 | -2.16 ± 0.25 |
| DP-MoG ₁₀ (MA) | -2.88 ± 0.04 | -2.87 ± 0.04 | -2.86 ± 0.04 |
| DP-MoG ₁₀ (zCDP) | -3.29 ± 0.04 | -3.10 ± 0.04 | -3.04 ± 0.04 |
| Gowalla | | | |
| $\delta = 1.11 \times 10^{-5}$ | $\epsilon = 2.25$ | $\epsilon = 3.50$ | $\epsilon = 4.00$ |
| DP-NF (GDP) | -0.44 ± 0.35 | - | - |
| DP-NF (MA) | -1.81 ± 0.11 | -0.77 ± 0.30 | 0.47 ± 0.32 |
| DP-MoG ₁₁ (MA) | -0.92 ± 0.01 | -0.92 ± 0.01 | -0.91 ± 0.01 |
| DP-MoG ₁₁ (zCDP) | -0.87 ± 0.01 | -0.93 ± 0.01 | -0.93 ± 0.01 |

We also provide Figure 11, which is the full version of Figure 3 in Section 4. Figure 11 provides dimension-wise histograms of the synthetically generated data for all ten dimensions of the Life Science dataset, presented in numerical order by axis index from left to right. The top two rows correspond to our baseline, and the bottom two rows correspond to our approach. For each image, we visualize the synthetically generated data (given in orange) and superimpose it over real data (given in blue) for comparison. We observe that for nearly all ten features, the distribution of data generated by DP-NF closely resembled that of the real data while DP-MoG was unable to replicate regions of concentrated density for certain dimensions.

C.2 Additional Results on Synthetic Datasets

We also evaluate our DP-NF method on several synthetic datasets. We perform this on the Half-Moons dataset (Figure 7), as well as a synthetically constructed dataset of a mixture of 8 Gaussians (Figure 12). This is done to demonstrate the heightened expressiveness of our approach as compared to a Gaussian mixture approach, alongside a worst-case scenario where the data is truly generated by a mixture of Gaussians, where DP-MoG would be expected to outperform our method. Results are presented in Table 5.

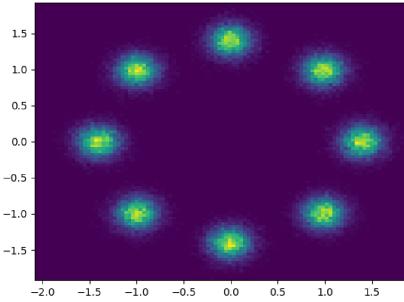


Figure 12: Gaussian Dataset

C.3 Hyperparameters and Training details

In this subsection we detail decisions about hyperparameter section in training, including the gradient clipping parameter C , regularization of the loss function, and choice of privacy accountants.

For the gradient clipping parameter C , prior work [1] suggested that a reasonable heuristic for setting this clipping parameter was to set C equal to the median of the ℓ_2 norms of the unclipped gradients observed over the course of a non-private training execution. In the context of normalizing flows, we found that much larger values for C yielded significantly preferable results. A natural explanation for this arises under consideration of log likelihood as an objective. In cases where a given point is assigned near-zero density, a large gradient update would be incurred to prevent further deterioration. When this gradient update is clipped, the resulting update may be insufficient to avoid associated numerical instability if this point is assigned density further approaching zero. Despite excess noise

Table 5: Average test log likelihood for varying privacy budgets ϵ . Error bars denote standard deviation over ten independent cross-validation splits. DP-MoG $_x$ refers to a Gaussian mixture of x components. Omitted entries occur when model training was early-stopped because it had already converged at lower ϵ values.

| Half-Moons | | | | |
|--------------------------------|-------------------|-------------------|-------------------|--|
| $\delta = 3.70 \times 10^{-5}$ | $\epsilon = 1.50$ | $\epsilon = 2.25$ | $\epsilon = 3.00$ | |
| DP-NF (GDP) | -2.23 ± 0.06 | - | - | |
| DP-NF (MA) | -2.57 ± 0.04 | -2.43 ± 0.06 | -2.22 ± 0.06 | |
| DP-MoG ₃ (MA) | -2.59 ± 0.01 | -2.58 ± 0.01 | -2.58 ± 0.01 | |
| DP-MoG ₃ (zCDP) | -2.73 ± 0.01 | -2.60 ± 0.01 | -2.60 ± 0.01 | |
| Gaussians | | | | |
| $\delta = 3.70 \times 10^{-5}$ | $\epsilon = 1.50$ | $\epsilon = 2.25$ | $\epsilon = 3.00$ | |
| DP-NF (GDP) | -2.45 ± 0.11 | - | - | |
| DP-NF (MA) | -2.73 ± 0.02 | -2.61 ± 0.08 | -2.45 ± 0.11 | |
| DP-MoG ₈ (MA) | -2.44 ± 0.01 | -2.45 ± 0.01 | -2.45 ± 0.01 | |
| DP-MoG ₈ (zCDP) | -2.42 ± 0.01 | -2.41 ± 0.01 | -2.43 ± 0.01 | |

being applied to updates with larger C , we found this to merely prolong training without a significant degradation in resulting model quality.

With respect to regularization, [44] suggested a modest amount of ℓ_2 regularization (i.e., a coefficient of 10^{-6}) in the context of non-private normalizing flows. We found this regularization approach to substantially degrade the quality of the resulting models and was generally omitted in our training. This makes intuitive sense, as the suggested regularization of [44] serves to decrease model weights over the course of training, and differentially private optimization applies Gaussian noise vectors of constant variance to gradients throughout training. As model weights tend toward zero, one would expect that the noise injection from privacy eventually dominates the learned criteria of the model.

For privacy accountants, we found that composition under Gaussian differential privacy (GDP) [11] consistently yielded the tightest privacy bounds throughout our experiments. See Figure 9 in

Appendix A.3 for an illustration of the improvements in privacy composition that can be achieved by GDP, over the moments accountant of [1]. Since our baseline method DP-MoG yielded the best performance under the moments accountant in [45], we included both GDP composition and moments accountant for fair comparison. We found that DP-NF consistently outperformed DP-MoG even when both methods used the moments accountant, and that further performance improvements could be achieved by DP-NF using GDP composition as a privacy accountant method.

D LIMITATIONS OF PROPOSED APPROACH

Normalizing flow models are trained in a manner which minimizes the average negative log likelihood of the observed data. As such, it is not uncommon when training such models that a given point is assigned near-zero density, provoking a loss approaching infinity. We observed this issue was somewhat exacerbated due to differentially private optimization, in part due to noise injection, but primarily as a result of subsampling. In the case of uniform subsampling, as required in our privacy analysis, there is no strict guarantee that any individual point is regularly sampled. This acts in contrast to typical sampling methodology in which the dataset is repeatedly shuffled and partitioned into equal sized batches over the course of an epoch. Rigorous privacy guarantees associated with equally sized and disjoint sampling is, to the best of our knowledge, a currently unsettled issue [37] and a potential avenue for improvement given its theoretical convergence guarantees. We did not find this limitation to be ultimately confounding in any way.

Additionally, given that DP-NF is ultimately a deep learning based approach to density estimation, it naturally involves the optimization of a large number of parameters and a high resource expenditure in terms of time and space complexity. This is especially highlighted in comparison to DP-MoG as a baseline, which takes on the order of seconds to run on CPU as compared to our method which can take on the order of an hour on GPU. Although, we find this tradeoff between resource expenditure and distribution quality to be justified, particularly in the context of differentially private data analysis and social science applications which rarely demand strict resource constraints within reason.