A General Approach to Adding Differential Privacy to Iterative Training Procedures

H. Brendan McMahan
mcmahan@google.com

Galen Andrew
galenandrew@google.com

Úlfar Erlingsson
ulfar@google.com

Steve Chien
schien@google.com

Ilya Mironov
mironov@google.com

Nicolas Papernot
papernot@google.com

Peter Kairouz
kairouz@google.com

Abstract

In this work we address the practical challenges of training machine learning models on privacy-sensitive datasets by introducing a modular approach that minimizes changes to training algorithms, provides a variety of configuration strategies for the privacy mechanism, and then isolates and simplifies the critical logic that computes the final privacy guarantees. A key challenge is that training algorithms often require estimating many different quantities (vectors) from the same set of examples — for example, gradients of different layers in a deep learning architecture, as well as metrics and batch normalization parameters. Each of these may have different properties like dimensionality, magnitude, and tolerance to noise. By extending previous work on the Moments Accountant for the subsampled Gaussian mechanism, we can provide privacy for such heterogeneous sets of vectors, while also structuring the approach to minimize software engineering challenges.

1 Introduction

There has been much work recently on integrating differential privacy (DP) techniques into iterative training procedures like stochastic gradient descent [Chaudhuri et al., 2011, Bassily et al., 2014, Abadi et al., 2016, Wu et al., 2017, Papernot et al., 2017]; for completeness we provide a formal definition of DP in Appendix A. Although these works differ in the granularity of privacy guarantees offered and the method of privacy accounting, most proposed approaches share the general idea of iteratively computing a model update from training data and then applying the Gaussian mechanism for differential privacy to the update before incorporating it into the model. Our goal in this work is to decouple, to the extent possible, three aspects of integrating a privacy mechanism with the training procedure:

a) the specification of the training procedure itself (e.g., stochastic gradient descent with batch normalization and simultaneous collection of accuracy metrics and training data statistics),

b) the selection and configuration of the privacy mechanisms to apply to each of the aggregates collected (model gradients, batch normalization weight updates, and metrics), and

c) the accounting procedure used to compute a final (ε, δ)-DP guarantee.

	Per-example SGD	Microbatch SGD	Federated learning (user-level DP)
record	gradient on one example	average gradient on one *microbatch* (~10 examples)	model update from one user
sample	minibatch (~100 examples)	minibatch (~10 microbatches for a total of 100 examples)	set of participating user devices for the round

Table 1: Defining *record* and *sample* in different training contexts.

This separation is critical: the person implementing a) is likely not a DP expert, and this code typically already exists; there are many configuration options for b), which will likely require experimentation, and this configuration logic may become complex; thus isolating the key privacy calculations in c) and keeping them as simple (and well tested) as possible prevents bugs in a) or b) from introducing errors in the calculation of the actual privacy achieved.

While model training is our primary motivation, the approach is applicable to any iterative procedure that fits the following template. We have a database with n records. A *record* might correspond to a single training example, a "microbatch" of examples, or all of the data from a particular user or entity (e.g., to achieve user-level DP as in McMahan et al. [2018]). On each round, a random subset of records (a *sample*) is selected and the training procedure consumes the results of a number of vector queries over that sample; see Table 1. Such vector queries may include the average gradient for each layer, updates to batch-normalization parameters, or the average value for different training accuracy metrics. We describe a general approach to allocating a privacy budget across each of these queries and analyzing the privacy cost of the complete mechanism, all respecting the decoupling of concerns described earlier. Our analysis builds on the Moments Accountant approach of Abadi et al. [2016], which applies to a single vector query per round, and generalizes the extension of McMahan et al. [2018] to multi-vector queries.

We focus on the following basic building block for a single vector. Suppose we have a database X with n records consisting of vectors $x^i \in \mathbb{R}^D$ and we are interested in estimating the average[1] $\frac{1}{n}\sum_i x^i$. Given a selection probability q, clipping threshold S, and noise multiplier z, the procedure is:

1. Select a subset of the records $R \subseteq [1, \ldots, n]$ by choosing each record with probability q.
2. Clip each x^i for $i \in R$ to have maximum L_2 norm S using $\pi_S(x) = x \cdot \min(1, S/\|x\|_2)$.
3. Output $\hat{x} = \frac{1}{qn}\left(\sum_{i \in R} \pi_S(x^i) + \mathcal{N}(0; \sigma^2 I)\right)$ where $\sigma = zS$.

The quantity $\sum_{i \in R} \pi_S(x^i) + \mathcal{N}(0; \sigma^2 I)$ is the output of the Gaussian mechanism for sums. As $\mathbb{E}[|R|] = qn$, scaling it by $1/qn$ produces an unbiased estimate of the average. The *noise multiplier* $z \equiv \sigma/S$ (the ratio of the noise to the L_2-sensitivity of the query) acts as a knob to trade off privacy vs. utility. If we choose $z = \frac{1}{\varepsilon}\sqrt{2 \ln 1.25/\delta}$, the mechanism is $(q\varepsilon, q\delta)$-differentially private with respect to the full database [Beimel et al., 2014, Dwork and Roth, 2014]. Importantly, the privacy cost of this mechanism is fully specified by q together with the *privacy tuple* (S, σ), where S is an upper bound on the L_2 norm of the vectors being summed, and σ is the standard deviation of the noise added to the sum.

We generalize the above procedure to the case where each record corresponds to a collection of vectors. We still do the sampling step (1) only once, but we estimate the average of each of the vectors separately, potentially with different clipping thresholds and noise multipliers. Let $v = (v_1, \ldots, v_m)$ be the total set of vectors for which averages are to be estimated privately. (Note, we will include the superscript i on the i^{th} record v^i only when necessary in equations that sum over records.) In general, we may partition this set of m vectors into multiple groups, e.g., fully connected layers vs. convolutional layers vs. metrics. We assume the user (that is, the person using the privacy tools defined here) has identified the relevant set of groups whose averages are needed

---

[1]For simplicity, we focus on unweighted average queries, for example to compute average gradient on a batch of examples; the generalization to weighted average and vector sum queries is straightforward. We also restrict attention to the fixed expected denominator f_f of McMahan et al. [2018]; extension to other estimators for averages like their f_c is straightforward.

in the training procedure. For each of these, they need to specify a privacy mechanism together with some hyperparameters. We first describe the privacy mechanisms that can be applied to individual vectors or groups of vectors, then show how the privacy cost of the full collection of mechanisms can be calculated, and finally propose strategies for choosing the parameters to achieve the desired privacy versus utility tradeoff.

Implementations of techniques in this paper may be found in the open-source TensorFlow Privacy framework [Google et al., 2018] for TensorFlow [Abadi et al., 2015], as described in Section 7.

2 Privacy mechanisms for a group of vectors

In this section, we describe two strategies that can be applied to a single group of vectors, *WLOG* the first k, $(v_1, \ldots, v_k)$, for $k \leq m$; when $k = 1$, the two mechanisms described are identical. Both mechanisms allow individual noise standard deviation parameters to be used for the separate groups. While this might at first seem to preclude the use of the Moments Accountant, which requires spherical noise, we will show how to resolve this issue in the next section.[2]

Separate clipping and noise parameters. This strategy essentially treats the whole group as a single concatenated vector $v = (v_1, \ldots, v_k)$. The user provides S_g, a clipping parameter, and σ_g, a noise parameter. For now, assume both of these parameters are simply chosen so as to provide reasonable utility for the resulting average; we will discuss strategies for choosing these parameters in detail in Section 4. The output of the mechanism is

$$\hat{v}_j = \frac{1}{qn} \sum_{i \in R} \pi_{S_g}\big(v^i\big)_j + \mathcal{N}(0; \sigma_g^2 I) = \frac{1}{qn} \left(\sum_{i \in R} \pi_{S_g}\big(v^i\big) + \mathcal{N}(0; \tilde{\sigma}_g^2 I) \right)_j,$$

where $\tilde{\sigma}_g = qn\sigma_g$ and $j \in [k]$. The final expression shows that the mechanism is equivalent to the Gaussian mechanism for sums with privacy tuple $(S_g, \tilde{\sigma}_g)$. Applying this mechanism with $S_g = S^*$ to all m vectors recovers the "flat clipping" approach of McMahan et al. [2018], and applying this mechanism separately to each of the vectors with $S_g = S^*/\sqrt{m}$ recovers their "per-layer clipping" approach (where S^* is the total L_2 bound). Another reasonable strategy that takes into account dimensionality is to apply the mechanism separately with $S_g = S^*/\sqrt{d_g/D}$ where d_g is the dimensionality of v_g.

Joint clipping. Here we introduce a new mechanism that allows us to clip less aggressively than applying the previous strategy to each vector individually, while still letting different vectors live on different multiplicative scales. The user supplies as input scale parameters $\alpha_1, \ldots, \alpha_k$, which may be thought of as bounds or reasonable L_2 norm clip parameters on the individual v_i, were they to be clipped individually. The strategy first does a pre-processing step via the scaling operator $\mathbf{s}(v; \alpha_{1:k}) = (v_1/\alpha_1, \ldots, v_k/\alpha_k)$. If $\|v_j\|_2 \leq \alpha_j$ for all $j \in [k]$, then the joint norm $\|\mathbf{s}(v; \alpha_{1:k})\|_2 \leq \sqrt{k}$, however it may typically be much less. Then joint clipping and noising is performed using a total clipping parameter $S_g \in [0, \sqrt{k}]$ and noise with the standard deviation of $\alpha_j \sigma_g$. The mechanism's output then scales the vectors back by the α_j factor in post-processing:

$$\hat{v}_j = \frac{1}{qn} \sum_{i \in R} \alpha_j \pi_{S_g}\big(\mathbf{s}(v^i; \alpha_{1:k})\big)_j + \mathcal{N}\big(0; (\alpha_j \sigma_g)^2 I\big)$$

$$= \frac{\alpha_j}{qn} \left(\sum_{i \in R} \pi_{S_g}\big(\mathbf{s}(v^i; \alpha_{1:k})\big) + \mathcal{N}\big(0; \tilde{\sigma}_g^2 I\big) \right)_j,$$

where again $\tilde{\sigma}_g = qn\sigma_g$. The final expression shows the output can be written as a post-processing of the subsampled Gaussian mechanism for sums with privacy tuple $(S_g, \tilde{\sigma}_g)$. Note that if no clipping happens then $\alpha_j \pi_{S_g}(\mathbf{s}(v; \alpha_{1:k}))_j = v_j$ for all j.

To see where this mechanism might be superior to the first, suppose v_1 and v_2 have $\|v_1\|_2 \leq 1$ and $\|v_2\|_2 \leq 100$, and suppose they can tolerate noise standard deviations of 0.01 and 1 respectively.

---

[2]Privacy mechanisms for groups can be used within TensorFlow Privacy [Google et al., 2018] by employing the `NestedQuery` class, which evaluates an arbitrary nested structure of queries where each leaf query would be a `GaussianAverageQuery` corresponding to one group of vectors.

Additionally, assume it is known that either v_1^i or v_2^i will be zero for any record i. We could clip these separately, but this ignores the (useful) side information that one of the vectors is always zero. On the other hand, if we treat them as a single group, we cannot take into the account the fact they are on very different scales; in particular, we must pick a single noise value which will either be insufficient to add privacy for v_2, or will completely obscure the signal in v_1. The joint mechanism proposed here lets us directly handle this situation using $\alpha_1 = 1$, $\alpha_2 = 100$, $\sigma_g = 0.01$, and $S_g = 1$.

3 Composing privacy guarantees for multiple vector groups

Now, suppose we have partitioned the m vectors into G groups, and selected a privacy mechanism for each one, producing privacy tuples $(S_g, \tilde{\sigma}_g)$ for $g \in \{1, \ldots, G\}$. From a privacy accounting point of view, each of these mechanisms is equivalent to running a Gaussian sum query on vectors w_g with $\|w_g\| \le S_g$ and then adding noise $\tilde{\sigma}_g$ to the final sum. We now demonstrate a transformation that lets us analyze this composite mechanism as a single Gaussian sum query on the sample for use with the privacy accountant.

First, we scale each vector $\mathbf{s}(w; \tilde{\sigma}_{1:G}) = (\frac{w_1}{\tilde{\sigma}_1}, \ldots, \frac{w_G}{\tilde{\sigma}_G})$, so $\|\mathbf{s}(w; \tilde{\sigma}_{1:G})\| \le S^* \equiv \sqrt{\sum_g (S_g/\tilde{\sigma}_g)^2}$. Now, we imagine a single Gaussian sum query with noise standard deviation $\sigma = 1$, and output the estimate after rescaling by the $\tilde{\sigma}_g$ factors. This is equivalent since

$$\hat{w}_g = \frac{1}{qn}\left(\sum_{i \in R} w_g^i + \mathcal{N}(0; \tilde{\sigma}_g^2 I)\right) = \frac{\tilde{\sigma}_g}{qn}\left(\sum_{i \in R} \mathbf{s}(w^i; \tilde{\sigma}_{1:G}) + \mathcal{N}(0; I)\right)_g. \tag{1}$$

The final expression is a simple post-processing on the output of a single Gaussian sum query with parameters $(S^*, \sigma = 1)$. Thus, we can apply the privacy accountant to bound the privacy loss of iterative applications of this mechanism.

4 Hyperparameter selection strategies

Here we consider selecting hyperparameters q, S_g, and σ_g to achieve a particular privacy vs. utility tradeoff. Recall for both mechanisms, $\tilde{\sigma}_g = qn\sigma_g$, so the key quantity is

$$z = \frac{1}{S^*} = \left(\sum_g (S_g/\tilde{\sigma}_g)^2\right)^{-1/2} = qn\left(\sum_g (S_g/\sigma_g)^2\right)^{-1/2}.$$

Typically, a value of $z \approx 1$ will provide a reasonable privacy guarantee. If z is too small for the desired level of privacy, the user has several knobs available: clip more aggressively by decreasing the S_g's; noise more aggressively by scaling up the σ_g's; or increasing q. When datasets are large and the additional computational cost of processing larger samples R is affordable, this last approach is generally preferable, as observed by McMahan et al. [2018]. If additionally the total number of iterations T is known, then since the privacy cost scales monotonically with any of these adjustments to z, a binary search can be performed using the privacy accountant repeatedly with different parameters to find e.g. the precise value of q needed to achieve a particular (ε, δ)-DP guarantee.

Choosing σ_g and S_g. Typical approaches to setting S_g include: 1) using an *a priori* upper bound on the L_2 norm; 2) choosing S_g so that "few" vectors are clipped; or 3) running parameter tuning grids to find a value of S_g that does not reduce utility (e.g., the accuracy of the model) by too much. If private data is used in 2) or 3), the privacy cost of this should be accounted for. Similar strategies can be used to choose σ_g, e.g., selecting a value that will introduce an *a priori* acceptable amount of error, or more likely for model training, running experiments to find the largest amount of noise that does not slow the training procedure.

In some cases one may have bounds S_g on the norms of G groups plus an overall target value of z, which needs to be distributed across multiple groups. To achieve *proportional* noise, where $\tilde{\sigma}_g \propto S_g$ for all g, we can use $\tilde{\sigma}_g = z\sqrt{G}S_g$. Another reasonable alternative, *dimensionality adjusted* noise assigns noise proportional to the maximum root mean squared value of the components of w_g given its bound and its dimensionality: $\tilde{\sigma}_g = z\sqrt{D/d_g}S_g$, where d_g is the dimensionality of group g and $D \equiv \sum_g d_g$.

5 Sampling policies

The basic update step of the SGD algorithm operates on a small subset of records (the *mini-batch*). Convergence guarantees of the standard optimization theory hold under the assumption that each minibatch is an i.i.d. sample of the training dataset, and the original Moments Accountant by Abadi et al. [2016] supported privacy analysis in this regime.

In practice, there are valid reasons for using alternative policies for sampling minibatches, with implications for privacy analysis. We list three of the most common sampling policies below.

Minibatches are i.i.d. samples. Privacy of this sampling procedure is analyzed by Abadi et al. [2016] and it is used by the federated learning framework where decisions of whether to participate in a particular update step are made locally [McMahan et al., 2018]. If the privacy accountant is dependent on the secrecy of the sample (as in the case of the Moments Accountant), then the size of the sample cannot be released without applying a privacy-preserving mechanism, which can be as simple as additive noise. The variability of the sample's size makes this sampling policy a poor fit for hardware accelerators. It can be repaired by sampling subsets of a fixed size from the training set *without replacement*, which leads us to the next policy.

Minibatches are equally sized and independent. The basic SGD corresponds to this sampling policy and minibatches of cardinality 1. Recent works analyze composition of this sampling policy with a mechanism satisfying RDP [Wang et al., 2018] or tCDP [Bun et al., 2018]. Independence of minibatches makes analysis of multiple iterations of SGD straightforward via application of composition rules for differential privacy.

Minibatches are equally sized and disjoint. In practice, the most common manner of forming minibatches is permuting the training dataset and partitioning it into disjoint subsets of a fixed size. After a single pass (an *epoch*) the process is repeated. This sampling policy can be efficiently implemented, and has intuitive semantics: an epoch corresponds to a training cycle when all examples were visited exactly once. Quantitatively tight analysis of DP-SGD in this model is not known. (A related problem of analyzing *randomized response* followed by a random permutation is addressed by Erlingsson et al. [2019].)

6 Privacy ledger

In principle, privacy accounting (via, e.g. the moments accountant) could be done in tandem with calls to the mechanism to keep an online estimate of the (ϵ, δ) privacy guarantee. However we advocate a different approach which cleanly separates concerns b) and c) from the introduction. We maintain a *privacy ledger* and record two types of events: *sampling events*, which record that a set R of records has been drawn using parameters q and n, and *sum query* events, which record that a Gaussian sum query has been performed over some group of vectors with privacy tuple $(S_g, \tilde{\sigma}_g)$. Then the privacy accountant can process the ledger *post hoc* to produce a privacy guarantee, first converting each group of one sampling event plus some sum query events to an equivalent single sum query event with parameters $(S^*, \sigma = 1)$ using Equation (1).

There are two main advantages of this approach. First, bugs in the hyperparameter selection strategy code cannot affect the privacy estimate. Second, it allows the privacy accounting mechanism to be changed and the ledger reprocessed if, for example, a tighter bound on the privacy loss is discovered after the data has been processed.

7 TensorFlow Privacy

TensorFlow Privacy [Google et al., 2018][3] is a Python library that implements TensorFlow optimizers for training machine learning models with differential privacy. The library comes with tutorials and analysis tools for computing the privacy guarantees provided. From an engineering perspective, the implementation of differentially private optimizers found in the library leverages the decoupled structure outlined above to make it easier for developers to both (a) wrap most optimizers into their

---

[3]Available from `https://github.com/tensorflow/privacy` under Apache 2.0 license.

differentially private counterpart and (b) compare different privacy mechanisms and accounting procedures.

Perhaps the library is best illustrated by one of its main use cases: training a neural network with differentially-private stochastic gradient descent [Abadi et al., 2016]. Given the stochastic gradient descent optimizer class, `tf.train.GradientDescentOptimizer`, implemented in the main TensorFlow library, one first wraps it into a new optimizer that implements logic for both the clipping and noising of gradients needed to obtain privacy. This is done by having the optimizer estimate the gradients via an instance of a class implementing the `DPQuery` interface. A `DPQuery` is responsible for clipping gradients computed by the optimizer, accumulating them, and returning their noisy average to the optimizer. This introduces two additional hyperparameters to the optimizer: the *clipping norm* and the *noise multiplier*. The `PrivacyLedger` class maintains a record of the sum query events for each sampling event which can then be processed by the RDP accountant.

In addition, our implementation leverages microbatches, as defined in Table 1. This implies that gradients are computed over several examples before they are clipped, and once all microbatches in a minibatch have been processed, they are averaged and noised. This introduces a third additional hyperparameter to the optimizer: the *number of microbatches*. Increasing it often improves utility but typically slows down training.

TensorFlow Privacy is also designed to work with training in a federated context in the vein of McMahan et al. [2018]. In that case the "gradients" supplied to the `DPQuery` would in fact be the model updates supplied by the users in a given round.

Finally, to compute the (ε, δ) differential privacy guarantee for the model, an implementation of the RDP accountant is provided. Given the sampling fraction q and the noise multiplier z, the RDP is computed for a step. Summing the RDP over the steps, it can then estimate ε for a fixed δ.

8 Floating-point arithmetic and randomness source

The hallmark feature of the definition of differential privacy is that it is *uncoditional*, in other words, it makes no assumptions about the adversarial knowledge or capabilities. It also puts a high burden on a differentially private implementation: its output distribution must have effectively infinite entropy. In practice, the distribution is defined over only a finite domain (such as a vector of single-precision floating-point numbers) and the source of randomness is guaranteed (at best) to be computationally secure. We consider these issues in turn.

Floating-point arithmetic. The problem of achieving differential privacy by means of standard floating-point arithmetic has been addressed for the additive Laplace mechanism by Mironov [2012]. We leave open the task of developing a provable floating-point implementation of DP-SGD and integrating it into an ML library.

Sources of randomness. Most computational devices have access only to few sources of entropy and they tend to be very low rate (hardware interrupts, on-board sensors). It is standard—and theoretically well justified—to use the entropy to seed a cryptographically secure pseudo-random number generator (PRNG) and use the PRNG's output as needed. Robust and efficient PRNGs based on standard cryptographic primitives exist that have output rate of gigabytes per second on modern CPUs and require a seed as short as 128 bits [Salmon et al., 2011].

The output distribution of a randomized algorithm A with access to a PRNG is indistinguishable from the output distribution of A with access to a true source of entropy *as long as the distinguisher is computationally bounded*. Compare it with the guarantee of differential privacy which holds against any adversary, no matter how powerful. As such, virtually all implementations of differential privacy satisfy only (variants of) Computational Differential Privacy introduced by [Mironov et al., 2009]. On the positive side, a computationally-bounded adversary cannot tell the difference, which allows us to avoid being overly pedantic about this point.

A training procedure may have multiple sources of non-determinism (e.g., dropout layers or an input of a generative model) but only those that are reflected in the privacy ledger must come from a cryptographically secure PRNG. In particular, the minibatch sampling procedure and the additive

Gaussian noise must be drawn from a PRNG for the trained model to satisfy computational differential privacy. In contrast, microbatches need not be chosen using a randomized process.

9 Conclusion

We have shown how the Gaussian mechanism can be applied to vectors of different types with different norm bounds and noise standard deviations, enabling training over heterogeneous parameter vectors, as well as simultaneous privacy-preserving estimation of other statistics such as classifier accuracy, or the number of instances in each class. By implementing iterative training algorithms in terms of a series of Gaussian sum queries and then recording for each query privacy events to a ledger to be processed by a privacy accountant, we separate the three major concerns of implementing privacy-preserving iterative training procedures while allowing flexibility in the specification of clipping strategy and noise allocation. The techniques described in the paper can be easily implemented using the Tensorflow Privacy library.

References

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://tensorflow.org/.

Martín Abadi, Andy Chu, Ian Goodfellow, Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *23rd ACM Conference on Computer and Communications Security (ACM CCS)*, pages 308–318, 2016.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473, 2014.

Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine Learning*, 94(3):401–437, 2014.

Mark Bun, Cynthia Dwork, Guy N. Rothblum, and Thomas Steinke. Composable and versatile privacy via Truncated CDP. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 74–86, 2018.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *J. Mach. Learn. Res.*, 12(Mar):1069–1109, 2011.

Cynthia Dwork and Aaron Roth. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science. Now Publishers, 2014.

Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2468–2479, 2019.

Google et al. TensorFlow Privacy. https://github.com/tensorflow/privacy, 2018.

Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International Conference on Learning Representations (ICLR)*, 2018. URL https://openreview.net/pdf?id=BJ0hF1Z0b.

Ilya Mironov. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and Communications Security (CCS)*, pages 650–661, 2012.

Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil Vadhan. Computational differential privacy. In *Advances in Cryptology—CRYPTO*, pages 126–142, 2009.

Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. In *Proceedings of the International Conference on Learning Representations*, 2017. URL https://arxiv.org/abs/1610.05755.

John K Salmon, Mark A Moraes, Ron O Dror, and David E Shaw. Parallel random numbers: As easy as 1, 2, 3. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, page 16. ACM, 2011.

Yu-Xiang Wang, Borja Balle, and Shiva Kasiviswanathan. Subsampled Rényi differential privacy and analytical moments accountant. *CoRR*, abs/1808.00087, 2018. URL http://arxiv.org/abs/1808.00087.

Xi Wu, Fengan Li, Arun Kumar, Kamalika Chaudhuri, Somesh Jha, and Jeffrey F. Naughton. Bolt-on differential privacy for scalable stochastic gradient descent-based analytics. In *Proceedings of SIGMOD*, pages 1307–1322, 2017.

A Differential Privacy

The formal definition of (ε, δ)-differential privacy is provided here for reference:

Definition 1. *A randomized mechanism $\mathcal{M} \colon \mathcal{D} \mapsto \mathcal{R}$ satisfies (ε, δ)-**differential privacy** if for any two adjacent datasets $X, X' \in \mathcal{D}$ and for any measurable subset of outputs $\mathcal{Y} \subseteq \mathcal{R}$ it holds that $\Pr\left[\mathcal{M}(X) \in \mathcal{Y}\right] \leq e^{\varepsilon} \Pr\left[\mathcal{M}(X') \in \mathcal{Y}\right] + \delta$.*

The interpretation of *adjacent datasets* above determines the unit of information that is protected by the algorithm: a differentially private mechanism guarantees that two datasets differing only by addition or removal of a single unit produce outputs that are nearly indistinguishable. For machine learning applications the two most common cases are *example-level* privacy (e.g., Chaudhuri et al. [2011], Bassily et al. [2014], Abadi et al. [2016], Wu et al. [2017], Papernot et al. [2017]), in which an adversary cannot tell with high confidence from the learned model parameters whether a given example was present in the training set, or *user-level* privacy (e.g., McMahan et al. [2018]) in which adding or removing an entire user's data from the training set should not substantially impact the learned model. It is also possible to consider X and X' to be adjacent if they differ by *replacing* a training example (or an entire user's data) with another, which would increase the ε by a factor of two.