The Philosophy of Differential Privacy

Claire McKay Bowen and Simson Garfinkel

1. Differential Privacy at 15

Differential privacy is under attack.

In 2006, a group of computer scientists invented differential privacy as an approach to provide mathematically provable privacy guarantees for statistical data publications based on confidential data. Systems that implement differential privacy inject privacy-protecting noise into statistics that are computed on the confidential data. To date, differential privacy is the only data privacy system for which the privacy guarantee does not depend on assumptions about attackers who seek to undo the privacy guarantee—assumptions such as how much auxiliary data the attacker might have, how much computer power is at their disposal, or how today's data publications might be combined with future data releases.

But differential privacy has its critics. They say that differential privacy injects too much noise into published statistics, making them unusable for any practical purpose [RFMS19]. As a result, academic journals and the US courts are debating and questioning whether data practitioners should use differential privacy for official and commercial statistics.

Differential privacy emerged from the computer science community as a rigorous definition for the privacy loss, or

Claire McKay Bowen is the lead data scientist for privacy and data security at the Urban Institute. Her email address is cbowen@urban.org.

Simson Garfinkel is a senior data scientist at the US Department of Homeland Security and part-time faculty at George Washington University. His email address is simsong@acm.org.

Simson Garfinkel's affiliation with the US Department of Homeland Security is provided solely for the purpose of identification. Claire McKay Bowen and Simson Garfinkel are both former employees of the US Census Bureau. Any views or opinions expressed in this article are the authors' and do not reflect views or policy of the US government, the Department of Commerce, the US Census Bureau, or the Department of Homeland Security.

Communicated by Notices *Associate Editor Richard Levine.*

DOI: https://doi.org/10.1090/noti2363

amount of disclosure risk, associated with data publishing. The community refers to these rigorous and quantifiable definitions as formal privacy loss measures. Since then, many data privacy experts have come to regard differential privacy as the "gold standard" for privacy protection. It's called a *formally private* method because statisticians can mathematically prove the worst-case, or maximum, privacy loss that will result from a data publication that uses differentially private methods.

Differential privacy was transformative. For the first time, the scientific community had a quantifiable definition of privacy—or more specifically, a mathematical definition for the incremental *privacy loss* from a publication based on confidential data. Previous techniques for privacy-preserving data publishing, such as *k-anonymity* [Swe02], were based on operational, empiric definitions: these methods described *how* a dataset would be protected, but they made no claims as to whether the resulting published data actually achieved the goal of protecting the data subjects' privacy. As a result, mathematicians using those techniques could not make assertions about the data protection without making sweeping assumptions as to how an attacker might exploit the published data, harness computational capabilities, and leverage available auxiliary datasets.

A second major breakthrough of differential privacy was that it did away with the (incorrect) belief that published data were either privacy protecting or privacy eroding. Instead, the definition of differential privacy recognizes that the privacy loss associated with each data publication or statistic is incremental and cumulative. Differential privacy then gives data practitioners an adjustable knob that can be tuned to control the maximum amount of potential privacy loss in any specific statistic or data publication.

We will walk through an example to understand the fundamental difference between differential privacy and traditional methods for *statistical disclosure limitation* or *control*. Imagine an organization that is required to publish aggregate statistics regarding the ages of people who live on a block, but is simultaneously prohibited from publishing any person's individual age. In one publication, the organization notes that there are three people who live on the block, and that their median age is 30. Of course, this immediately reveals that one person on the block is 30. Oh well! In another publication, the organization states that the average age is 44. At this point, there are only 30 possible combinations of integer ages that produce those official statistics. Hmm...

Now suppose that a local newspaper publishes an article the following month with a photo of the oldest person in the city, who just happens to live on that block and just celebrated their 80th birthday. Suddenly the cat is out of the bag: the ages of the block's inhabitants are 22, 30, and 80 [GAM19].

With differential privacy, the organization is prohibited from publishing exact statistics about the block. Instead of stating that $n = 3$, $\bar{x} = 44$, and $\hat{x} = 30$, the organization might publish that $n = 3.5$, $\bar{x} = 42.4$, $\hat{x} = 39$ ($\epsilon = 2.0$), where ϵ is the so-called *privacy loss budget*.

With this improved publication strategy, the newspaper article no longer undoes the privacy protection. Revealing the age of the oldest person doesn't compromise the ages of the other two, because the statistics organization published *noisy measurements*. Even better, because noise was intentionally added to the organization's official statistics, there is no way to validate the newspaper's article: perhaps the person really is 80, but perhaps the person is 79 or 83. This simple example shows why adding intentional inaccuracy to published statistics is a powerful approach for protecting privacy.

Differential privacy acknowledges that *every* data release based on confidential data fundamentally achieves two competing results: it benefits the public by making available new statistics, and it causes privacy loss to the individuals on whose data the publication is based. The inventors of differential privacy realized this balance between public benefit and private cost is inherently a public policy decision that is best left to policymakers and not something that can be reasoned about and decided by mathematicians. The privacy loss budget, ϵ, is the knob that differential privacy gives policymakers to negotiate this trade-off. Differential privacy allows privacy *researchers* to experiment with a wide range of ϵ values, but only the *practitioners* who use differential privacy to publish statistics need to be concerned with the final value that policymakers dictate.

Balancing the privacy loss against the social good in a data publication is not new: the US government has made this trade-off since 1840 when it first started protecting information collected about establishments and businesses for the decennial census. Statistical protection was extended to data about individuals in 1930.[1] The US Census Bureau has traditionally protected this information using suppression or generalization. However, such techniques do not compose, making them brittle when there are multiple data releases based on the same dataset, or when the attacker has auxiliary information.

Differential privacy does compose: the math makes it possible to compute the total potential privacy loss from multiple individual releases. Differential privacy also gives policymakers fine-grain control over that trade-off between accuracy and privacy loss because ϵ can enjoy any value between 0 and ∞. Furthermore, policymakers can split a single ϵ into many individual parameters, $\epsilon_0, \ldots, \epsilon_k$, giving many knobs for the statistics an organization publishes.

[1] *See* `https://www.census.gov/history` *for a history of privacy-preserving techniques used by the US Census Bureau.*

Differential privacy is a *definition* of the maximum privacy loss that can result from a data publication or statistic. As such, there is no single differentially private mechanism, algorithm, method, or technique (we will use these terms interchangeably). Instead, privacy researchers have created various mathematical *mechanisms* that can produce a wide range of data products—including statistical tables, public use microdata (or individual level data), and statistical machine learning classifiers. Somewhat confusingly, researchers have also created variants of differential privacy itself, each with definitions that are typically relaxed from the original.

The version of differential privacy created in 2006 is known as ϵ-differential privacy (or sometimes "pure" differential privacy). When $\epsilon = 0$, the data published have no relationship to the confidential data that are being protected. Although covered under the definition, this is not useful in practice. When $\epsilon = \infty$, no noise is added to the published data products. This is not privacy protecting. But it is acceptable under differential privacy to publish the confidential data when $\epsilon = \infty$ and should lay to rest the misconception that differential privacy cannot produce data with sufficient accuracy. If $\epsilon = \infty$ has insufficient accuracy, then the underlying data have insufficient accuracy or the differentially private mechanism is poorly constructed! Typically $0 < \epsilon < \infty$, and policymakers have the tough job of deciding the value.

Since 2006, members of the formal privacy research community have developed a variety of differential privacy variants, each with a different definition, some with additional parameters. The most common definitions are (ϵ, δ)-differential privacy, concentrated differential privacy, and Rényi differential privacy. Most of these mechanisms use Laplace or Gaussian noise, while some use less known distributions, such as Wishart. Each have different *composition* properties and different error distributions that result from their application.

In the following sections, we continue to explain the philosophy of differential privacy, provide the basic terminology associated with differential privacy, and show how mechanisms can be applied to protect simple statistical products. We discuss the two fundamental models of using differential privacy: the *trusted curator model* and the *local model*. We also review some of the challenges, both technical and political, that organizations have encountered when applying differential privacy. However, this article will refrain from a detailed presentation of differentially private algorithms and proofs of their correctness. For that information, we refer interested readers to our references.

We also refer to various groups within the formal privacy community. This community consists of researchers (e.g., privacy experts), practitioners, data curators, data users, and public policymakers. Researchers are individuals who specialize in developing data privacy methods. Practitioners are those who implement differentially private methods on real systems and release data, often to the public. Data curators are individuals or institutions that are responsible for the collection and storage of the confidential data. Data users are analysts and other researchers (e.g., economists) of the publicly released data.

2. It's About the Uncertainty

The existential reason that Dwork, McSherry, Nissim, and Smith created differential privacy was to provide privacy for individuals whose data were incorporated into *statistical* databases. "Intuitively, if the database is a representative sample of an underlying population, the goal of a privacy-preserving statistical database is to enable the user to learn properties of the population as a whole while protecting the privacy of the individual contributors" [DMNS06].

With this intuition, the authors first imagined that the data would be held by a "trusted server," and processed with an arbitrary "query function" (f) to produce a "true answer." Today we call the server a *trusted curator*. Previously Nissim and Dinur had shown that if the server answers too many questions with sufficient accuracy, then all of the private data in the database can be compromised [DN03].

But what if some specific individual's data weren't part of that collection? In that case, that person's privacy couldn't be impacted by the published statistical product.

To understand this intuition, let's go back to our city block example with three people on it. Imagine that Alice is an elderly woman who lives across town. When the statistical agency publishes that the average age of the block's residents is 44, clearly Alice's age cannot be compromised, because her age was not included in the statistic! But if Alice moves to the block and the city publishes that the average age is now 55.5, an attacker trying to determine Alice's age could reasonably infer that she is 90.

The literature of statistical disclosure limitation uses the term "data intruder" to describe an attacker who takes a data publication and attempts to learn something about one or more of the respondents that the trusted curator has pledged to keep confidential. The statistical agency applies a *query function* to its confidential database to produce the published statistics, while the data intruder tries to infer the *preimage* of the query function—the underlying confidential data—that produced the statistical publication.

Differential privacy uses noise to produce *uncertainty* about the preimage. With differential privacy, a statistical agency might start consulting with data users to learn what they wanted to do with the published statistics and the level of accuracy they required. The agency might learn

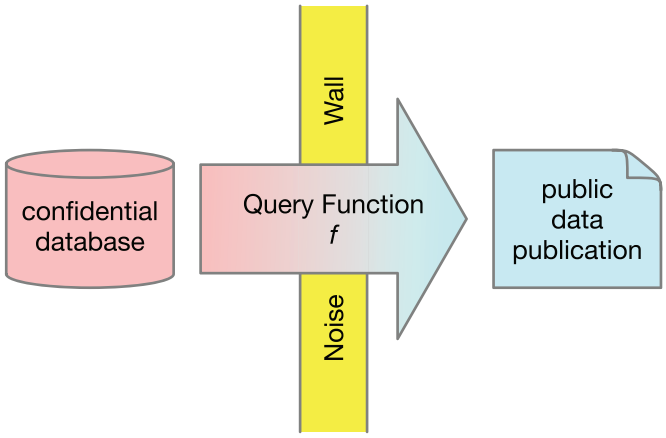

Figure 1. Differential privacy constructs a noise wall around the confidential database. Any query crossing the noise wall is changed unpredictably, creating uncertainty in the mind of the attacker regarding function f's preimage, the confidential database.

that data users want the count of people under and over a specific age instead of the count of people and the average and median ages. For the previous year, the agency might then publish that there were two people living on the block under the age of 65 and two people over the age of 65 ($\epsilon = 2$). In the current year, the agency might publish that there are three people under 65 and two people over 65 ($\epsilon = 2$).

In other words, differential privacy creates a wall of noise that surrounds and protects the confidential data: any query results that cross the noise wall must be modified in a manner that is unpredictable to the attacker (Figure 1).

Earlier we stated that differentially private methods are fundamentally different from another popular privacy mechanism, k-anonymity [Swe02]. Differentially private methods are based on the relationship between the confidential database and the public data publication, whereas k-anonymity is based on the mechanism applied to the confidential database itself. Here k-anonymity relies on the assumption that a data intruder cannot reliably single out a person's data if there are at least k individuals in a particular group who share the same identifying characteristics. k-anonymity had the aspirational hope that applying the k-anonymity mechanism produces a dataset that protects privacy. Researchers soon discovered that it did not, because *any* characteristic is potentially identifying when linked with an appropriate external database [Gar15].

Another difference between differentially private methods and k-anonymity mechanisms is that differential privacy does not provide absolute private guarantees. Differential privacy does not say that certain values of ϵ are "safe" and others are not. It also does not say that a single dataset benefits all attackers equally. Attackers with auxiliary information will learn more about the confidential database than attackers without such information: differential privacy makes it possible to compute *how much better* those attackers will be able to do so.

3. Differential Privacy's Math

Although this article refrains from covering the proofs, we present the basic differential privacy terminology. We provide both a high-level explanation and the mathematical definitions.

3.1. Differential privacy.

From the previous section, we understand that differential privacy protects a data publication or statistic with noise to create uncertainty, but how does it determine the amount of noise to add? At a high level, differential privacy links the potential for privacy loss to how much the estimate for a unique statistic (or query) from the underlying confidential data changes with the absence or presence of *any individual record that could possibly be in the database*.

More specifically, differential privacy quantifies the privacy loss for each statistic by the parameter ϵ. When adding or removing a single user from the data, a differentially private algorithm's output distribution changes by an amount limited by ϵ. This framework allows for a formal guarantee of the amount of information released about a confidential database over an arbitrary number of analyses, and it does not require assumptions concerning how a data intruder would attack the data or the amount of information they possess compared to other privacy loss measures.

However, the definition we present assumes the records are disjoint from one another. Records are often not in practice, requiring privacy researchers to make adjustments when developing their differentially private methods.

Mathematically, we define differential privacy as follows. Let $X \in \mathbb{R}$ represent the original database with dimension $n \times q$. We define n as the number of records in the original database and a statistical query as $f : \mathbb{R}^{n \times q} \to \mathbb{R}^k$, where the f maps the possible datasets of X to k real numbers.

Differential privacy ([DMNS06]). A sanitization algorithm, $\mathcal{M}$, satisfies ϵ-differential privacy if for all subsets $S \subseteq Range(\mathcal{M})$ and for all X, X' such that $d(X, X') = 1$,

$$\frac{\Pr(\mathcal{M}(X) \in S)}{\Pr(\mathcal{M}(X') \in S)} \leq \exp(\epsilon), \tag{1}$$

where $\epsilon > 0$ is the privacy loss budget and $d(X, X') = 1$ represents the possible ways that X' differs from X by one record.

This distance can be alternatively interpreted as the presence or absence of a record, or as a change in a record, where X and X' have the same dimensions. Both are valid

interpretations, but have different impacts mathematically on how the differentially private methods compose.

When adding Laplace noise, one of the most popular ϵ-differential privacy mechanisms, many data users find the extremes of Laplace noise to be inconvenient in some applications. A popular variant of differential privacy known as (ϵ, δ)-differential privacy offers similar theoretical levels of privacy protection with less noise, but with a small probability that, in some cases, the privacy definition bound does not hold.

(ϵ, δ)-Differential privacy ([DKM+06]). A sanitization algorithm, $\mathcal{M}$, satisfies (ϵ, δ)-differential privacy if for all X, X' that are $d(X, X') = 1$,

$$\Pr(\mathcal{M}(X) \in S) \leq \exp(\epsilon) \Pr(\mathcal{M}(X') \in S) + \delta, \qquad (2)$$

where $\delta \in [0, 1]$. ϵ-differential privacy is a special case of (ϵ, δ)-differential privacy when $\delta = 0$.

Most recent applications have used $10^{-7} \leq \delta \leq 10^{-10}$. This additional privacy parameter can produce significant improvements in query accuracy with very little real-world impact on privacy loss *most* of the time. As with ϵ, the value for δ is a policy decision and more research is needed to gain a better sense of what an appropriate δ should be.

3.2. **Composition theorems.** An advantage of differential privacy over other attempts to create mathematical definitions for privacy is that any method or algorithm that satisfies differential privacy will compose. This composition allows differential privacy practitioners to adjust an overall ϵ (sometimes referred to as the global privacy loss budget) for any use of the confidential data, and then divvy out smaller values of ϵ (i.e., $\epsilon = \epsilon_1 + \epsilon_2 + \epsilon_3 + \cdots$) for each data publication or statistic that makes use of the confidential data.

Mathematically, we define this distinction through two composition theorems.

Composition theorems ([DR13]). *Suppose a mechanism, $\mathcal{M}$, provides (ϵ_j, δ_j)-differential privacy for $j = 1, \ldots, k$.*

(a) ***Sequential composition**. The sequence of $\mathcal{M}_j(X)$ applied on the same X provides $(\sum_j \epsilon_j, \sum_j \delta_j)$-differential privacy.*

(b) ***Parallel composition**. Let D_j be disjoint subsets of the input domain D. The sequence of $\mathcal{M}_j(X \cap D_j)$ provides $(\max(\epsilon_j), \max(\delta_j))$-differential privacy.*

For k-many statistical queries on X, the composition theorems state that the practitioner may allocate a portion of the overall desired level of ϵ to each statistic by sequential composition. As an example, a common practice is splitting ϵ equally by k, such as a practitioner wanting to release k-many summary statistics and allocating ϵ/k to each statistic.

Unlike sequential composition, parallel composition does not require the privacy loss budget to be divided because the differentially private noise is applied to disjoint subsets of the input domain. A popular example is injecting noise to a histogram, where the bins are disjoint subsets of the data. This means a differentially private mechanism can add noise to each bin independently, without needing to split ϵ. Most differentially private methods try to leverage parallel composition instead of sequential to avoid splitting the privacy loss budget, because smaller ϵ means less accuracy. As we noted, this assumes that each person can only occupy a single bin, which may not be the case for some histograms.

3.3. **Post-processing.** Another important theorem within the differential privacy framework is the post-processing theorem, which states that any function applied to a differentially private output produces a new output that is also differentially private.

Post-processing theorem ([DMNS06]). *Let $\mathcal{M}$ be a mechanism that satisfies ϵ-differential privacy, and let g be any function. Then $g(\mathcal{M}(X))$ also satisfies ϵ-differential privacy.*

In practice, most applications of differential privacy use the post-processing theorem to enforce physical or structural data constraints, to avoid splitting the privacy loss budget further, and/or to generate synthetic data (i.e., pseudo records that are statistical representative of original data based on an underlying model). For instance, some differentially private methods sanitize the underlying distribution of the confidential data and sample from that noisy distribution to create synthetic data.

3.4. **Sensitivity.** Up to this point, we have focused on how the privacy loss parameter ϵ may be allocated to adjust the amount of noise being injected to an output. However, the query function's sensitivity to a single person's data contribution also determines how much noise should be added. Broadly, this sensitivity encodes how robust or resistant the differentially private mechanism is to the presence of outliers. Differential privacy quantifies this sensitivity by measuring how much the output changes in the confidential data given the absence or presence of the most extreme possible record that *could exist* in the population, but might not be observed in the data. The reasoning behind this measure is that if we do not know how the data intruder will attack the data or what external information they may have to help them, then we should protect for every *possible version of the data that could exist*. If the output crossing the noise barrier is very sensitive to the presence of a single outlier, then more noise is added to decrease the chance that an attacker will determine whether the outlier is present or absent.

To help understand this concept better, imagine the data we want to protect contains demographic and

financial information. The query we want answered is, "what is the median net worth of people on this block?" According to differential privacy, we must consider the change of the most extreme possible record that could exist or be in any given data that has socioeconomic information. For our example, that person is Jeff Bezos.[2] As it turns out, the median function is not very sensitive to the addition or removal of a single outlier in a population. If Bezos is altered or removed in the data, the median net worth will not change very much. This means we can provide a more accurate answer by adding less noise to the median net worth result, because the median net worth query is less sensitive to outliers like Bezos. Not all functions have this property! Consider the query, "what is the average net worth of people living on this block?" Unlike the previous query, the answer would significantly change if Bezos is or is not living on the block. In order to protect Bezos, a differentially private algorithm would then need to provide a significantly less accurate answer by adding more noise for the same privacy loss parameter.

We mathematically define this sensitivity as follows:

l_1-Global sensitivity ([DMNS06]). The global sensitivity of a function f is

$$\Delta_1 f = \sup_{d(X,X')=1} \|f(X) - f(X')\|_1. \tag{3}$$

Some differentially private mechanisms calculate the sensitivity under different norms, such as l_2 distance, referred to as l_2-global sensitivity. Choosing a norm depends on the specific differential privacy definition and the noise distribution being used.

While the definition is straightforward, calculating the global sensitivity can be difficult. For instance, we cannot directly calculate the global sensitivity of one of the most common statistical analyses, regression analysis, if the coefficients are unbounded. To address this issue, privacy researchers creatively figured out ways to add noise within the regression analyses, such as adding noise on the objective function.

3.5. Fundamental mechanisms.

Most differentially private algorithms build from the basic or fundamental differentially private mechanisms that add noise from various probability distributions. The Laplace Mechanism is the most basic mechanism that satisfies ϵ-differential privacy. This mechanism adds noise drawn from a Laplace distribution with a mean of zero and a scale parameter based on the sensitivity of the output and the privacy loss parameter.

Laplace Mechanism. The Laplace Mechanism satisfies ϵ-differential privacy by adding noise to f that are drawn from a Laplace distribution with the location parameter at 0 and scale parameter of $\Delta_1 f \epsilon^{-1}$ such that:

$$f^*(X) = f(X) + Laplace\left(0, \Delta_1 f \epsilon^{-1}\right). \tag{4}$$

Another popular mechanism is the Gaussian Mechanism that satisfies (ϵ, δ)-differential privacy that relies on l_2-global sensitivity.

Gaussian Mechanism. The Gaussian Mechanism satisfies (ϵ, δ)-differential privacy by adding Gaussian noise with zero mean and variance, σ^2, such that:

$$f^*(X) = f(X) + N\left(0, \sigma^2 I\right), \tag{5}$$

where $\sigma = \Delta_2 f \epsilon^{-1} \sqrt{2 \log(1.25/\delta)}$.

Both of these fundamental mechanisms are easy to apply in practice, but can only be readily applied to numerical outputs. A more general mechanism is the Exponential Mechanism, that satisfies ϵ-differential privacy and samples from the possible outputs rather than adding noise directly. In practice, the Exponential Mechanism can be computationally expensive without adjustments to the code or limiting the possible outputs for the target statistic.

Exponential Mechanism. The Exponential Mechanism releases values with a probability proportional to

$$\exp\left(\frac{\epsilon f(X, \theta)}{2\Delta_1 f}\right) \tag{6}$$

and satisfies ϵ-differential privacy, where $f(X, \theta)$ is the score or quality function that determines the values for each possible output, θ, on X.

Although there are many other differentially private mechanisms that use different noise distributions, the Laplace Mechanism and Gaussian Mechanism are the most common. The advantages and disadvantages of each varies based on the data and the use case. For instance, the Laplace Mechanism has smaller variation than the Gaussian Mechanism. However, the Gaussian Mechanism performs better over multiple queries, because multiple Gaussian distributions are still Gaussian. The Laplace Mechanism has only one tuning parameter, ϵ, whereas users of the Gaussian Mechanism must balance ϵ and δ.

3.6. Setting epsilon.

Setting the value of ϵ is up to the policymakers, who will ultimately shoulder the responsibility for selecting the budget. Yet, the decision should also be informed by the privacy researcher, who can explain to policymakers how to interpret the privacy and utility trade-off, and the participants in the data, who will have their own sense of personal privacy versus the common good. Considering all these perspectives *still* leaves the important decision of selecting a privacy loss budget open to debate.

Early in the development of differential privacy, researchers said that setting ϵ was a policy decision. At the same time, they stated that an ϵ less than or equal to one

[2] *According to* Forbes, *Jeff Bezos is the richest person in the world by net worth for 2020.*

was ideal, whereas an epsilon of two or three would release too much information. The researchers also noted that $\epsilon < 1.0$ allowed for certain algorithmic simplifications. Today we can look back at these statements and explain them as reflecting both the immaturity of the field and the fact that early research was almost entirely theoretical. Researchers had little to no experience balancing privacy loss against the societal benefit of a data release.

In contrast, many practitioners who have used differential privacy for actual data releases have required larger values of ϵ to achieve query results that are statistically "fit for purpose." In 2008, for example, researchers applied the (ϵ, δ)-differential privacy method with values at $(8.6, 10^{-5})$ to release a synthetic version of the OnTheMap data, a dataset based on individual US commuters [MKA+08]. More recently, Google's 2020 COVID-19 Mobility Reports provided movement trends over time by geography (e.g., county level in the United States) for different categories, such as transit stations and workplaces; Google used $\epsilon = 2.64$-differential privacy for the daily reports. Of course, the total privacy loss steadily increases as these daily reports are combined. In the same year, LinkedIn revealed their LinkedIn's Audience Engagement API that protected LinkedIn members' content engagement data, which used (ϵ, δ)-differential privacy with daily values of $(0.15, 10^{-10})$ [ABC+20, RSP+20].

As for how data participants perceive an appropriate value of ϵ, there is little published research on that topic. To date, very few federal agencies have adopted differential privacy for their data products. For the initial data releases of the 2020 Census, the Census Bureau decided on $\epsilon = 17.14$ for the persons file and $\epsilon = 2.47$ for the housing unit data, with $\delta = 10^{-10}$ for each.[3]

These examples (and the lack thereof) suggest there are many factors that affect the choice of ϵ and δ, including the type of information released to the public, social perception of privacy protection, statistical accuracy of the release data, among others. Essentially, more research is necessary to address the open question of how to select ϵ. More on this later.

3.7. Models of operation.

Over the years, differential privacy has evolved into two distinct approaches, or *models*: the "trusted curator" model and "local differential privacy."

3.7.1. Trusted curator model. Also called the *central model*, the trusted curator model is what we have been discussing for most of this article. A *trusted curator* receives confidential data, performs the queries and applies the differentially private noise, and releases the results. If the trusted curator is asked two similar queries, then noise must be applied twice. This results in expending more of the total privacy loss budget. If the privacy loss is capped, at some point the trusted curator must stop answering queries. If the privacy loss budget is not capped, then every time a query is answered, an attacker will be able to make more precise inferences about the nature of the confidential data.

3.7.2. Local differential privacy. Local differential privacy does not have a trusted curator. It instead focuses on how individual data are collected, where each respondent applies differential privacy *locally*, to their own data, before submitting the protected data to the curator. As a result, the curator no longer needs to be trusted. Local differential privacy is conceptually similar to the survey technique known as *randomized response* [War65]. In fact, the two approaches are mathematically equivalent under many conditions [WZFY20].

When employing randomized response, each individual receives a privacy budget (ϵ) where the log-difference in the probability of generating the same noisy response from two different individual responses is bounded by $(-\epsilon, \epsilon)$. This framework contrasts from the other differential privacy definitions, where ϵ is applied to the entire confidential data.

Local differential privacy can be modeled as a single complex query on a dataset, where the query asks for the value of every row. Alternatively, it can be modeled as n queries on a database of n rows, in which the first query asks for the first row, the second query asks for the second row, and so on. Local differential privacy thus enjoys parallel composition, because each row (or each of the n queries) is not related to any other row (or query).

At first, local differential privacy seems quite attractive to data curators considering approaches for privacy-preserving data analysis. After all, the resulting dataset is differentially private, and the data are microdata, allowing repeated reanalysis without additional privacy loss. The problem with this approach, as Kifer and Machanavajjhala eloquently phrased it, is that there is "no free lunch in data privacy" [KM11]. In the case of local differential privacy, substantially more noise is added to queries computed with locally noised microdata than on queries that are computed over true data by a trusted curator.

3.7.3. Hybrid models. Google's experience with integrating differential privacy into the Chrome web browser shows both the promise and real-world limitations of local differential privacy. In 2014, Google integrated Randomized Aggregatable Privacy-Preserving Ordinal Response (RAPPOR) into the Chrome web browser to collect sensitive statistics from users without jeopardizing their privacy [EPK14]. The system implemented local differential privacy, adding noise to microdata about a browser's

[3] *See* https://www2.census.gov/programs-surveys/decennial/2020/program-management/data-product-planning/2010-demonstration-data-products/ppmf20210428/2021-04-28-PPMF-rho-allocation-tables.xlsx.

home page, visited sites, and the processes running on end-user machines. One of the purposes of RAPPOR was to assist in the collection of security metrics. The idea was that even with the added noise, if Google detected that malware running on a user's machine was correlated with having visited a particular URL, Google would have additional data that could be used to improve its malware filtering systems. Unfortunately, Google no longer uses RAPPOR in practice. "There's just too much noise, even though the epsilon was pretty high."[4]

Following the release and use of RAPPOR, Google developed a hybrid approach that used secure aggregation to first combine statistics from multiple users, after which the differentially private noise is applied. "Thus, secure aggregation over just 1024-user subgroups holds the promise of 32× improvement in differentially private estimate precision," reported researchers from Google and Cornell [BIK+17]. Google noted that this extra precision is critical for machine learning applications.

4. Challenges for Adoption

Differential privacy was invented in 2006, so it is only 15 years old as of this writing. To put this into context, it's instructive to look at another breakthrough mathematical privacy-protecting technique, public key cryptography.

The basics of public key cryptography were researched in 1976 (the Diffie-Helman key exchange), 1977 (the RSA algorithm), and 1978 (certificates). But it wasn't until 17 years later that the *ssh* and *SSL* protocols were invented for sending information securely over the Internet. Bugs in these algorithms and their implementations were still being worked out as recently as the past decade. And today, despite billions of users deploying end-to-end encryption through various messaging platforms (e.g., Signal), a significant part of the global internet traffic still occurs over unencrypted HTTP.

Likewise, the invention of differential privacy inspired an explosion of new data privacy research, applying differential privacy to Bayesian inference, deep learning, facial recognition, generative adversarial network, genome-wide association studies, location privacy, multiparty computation, principal component analysis, recommender systems, social network analysis, and SQL queries, to name a few.

But while the number of publications grow, deep challenges persist in transitioning these mathematics from research to production. We group these challenges into *scientific challenges* owing to the immaturity of the field, *transition challenges* of moving theory to practice, and *social challenges* due to the radical differences between differential privacy and the previous techniques for privacy-preserving data analysis and publication.

4.1. Scientific challenges.

One of the major persisting problems is that much of the differential privacy research is highly theoretical. Some proposed algorithms can only be used with data that meet specific conditions not typically met in real-world data, or have unrealistic assumptions on what information is publicly available to inform or improve the method's results. Although the science is correct, it doesn't apply to real-world applications.

As an example, Google's first release of RAPPOR only worked well for the small number of cases in the original paper. More importantly, the privacy researchers initially did not track the cumulative privacy loss resulting from repeated use of the system to make measurements on the same individuals. Later, privacy researchers who did not work for Google showed that RAPPOR with a 30-minute reporting period would result with a privacy loss budget of $\epsilon = 25.63$ per day or $\epsilon = 769$ over the course of a month [RSP+20].

Another example is Uber's differentially private SQL queries system [JNS18]. In 2018, Johnson et al. developed a differentially private approach for 8.1 million queries for data gathered on rider and driver information, trip logs, and customer support from March 2013 to August 2016. While their method was a great step towards creating practical differentially private applications, at least one of differential privacy's creators argued that the method suffered from errors that led to the unintentional overreporting of the quality of their results.[5]

With any new methodology or technology, testing and experimenting is a necessary part of the process. Researchers are learning from their mistakes and continue to improve their methods. For instance, Google's previous research into differential privacy allowed the company to quickly develop "mobility reports" to provide access to invaluable data during the COVID-19 global pandemic [ABC+20].

4.1.1. Evaluating mechanisms. Another concern facing researchers is the correctness of approaches for evaluating differential privacy mechanisms.

From November 2018 to May 2019, the National Institute of Standards and Technology Public Safety Communications Research (NIST PSCR) Division hosted the first "Differential Privacy Synthetic Data Challenge." The data challenge called for researchers to develop practical and viable differentially private data synthesis methods that were scored based on specific metrics that were given ahead of time. Three "Marathon Matches" provided participants

[4] *Comment by Ulfar Erlingsson at the 2018 Theory and Practice of Differential Privacy.* `https://twitter.com/TedOnPrivacy/status/1051848416416976896?s=19`. *Mentions of the deprecation are also evident in the Chromium source code at* `https://bit.ly/370fr45`.

[5] *See footnote 2 on page 4 of* `https://arxiv.org/pdf/1909.01917.pdf`, *where the authors found that the open-source software version of Flex released the full list of all user identifiers with simple queries.*

with training data that were identical in structure and variables to the real-world test data used for final scoring. Although the challenge inspired new differentially private synthetic generation methods, many of the top scoring methods heavily relied on the training data to improve the method's results for the scoring metrics. This consequently caused some of these methods to perform worse on other utility metrics not used in the competition [BS21].

4.1.2. Applications to non-tabular data. One of the largest scientific challenges facing the differential privacy research community is applying the framework to non-tabular data, such as blocks of text longer than a few words, qualitative information, images, and video. To date, most of these efforts are experimental and not ready to be applied in practice. This is why many differentially private applications to date have been applied to tabular data.

Both Google's RAPPOR and COVID-19 Mobility Reports added noise to the counts of bits or counts of people within certain location categories, respectively. The OnTheMap and 2020 Census data products are counts of the United States population within specific geographic regions. The NIST PSCR Data Challenge used the San Francisco Fire Department's Call for Service data and the state Public Use Microdata Sample, where many of the competitors considered the data as tabular to apply their methods and then later post-processed the synthetic data to resemble the confidential data structure.

The few non-tabular differentially private applications tend to oversimplify the problem or significantly sacrifice privacy to improve utility. One such data type is social network data that represents a relational network among individuals or groups. In social network analysis, researchers define the people or entities in the network as *nodes* with various attributes (e.g., demographic information) and *edges* that represent a paired interaction between two nodes. Facebook data could be considered social network data, where the individuals are the nodes with attributes like gender and age while the edges are if the individuals are "Facebook friends" or not.

This structure is why social network data are very challenging to protect with differential privacy, because differential privacy must protect the worst-case scenario between two neighboring networks. An example scenario is comparing an empty network (i.e., all isolated or non-connecting nodes) and a star network (i.e., all nodes are only connected to one node). If the confidential data is a star network, then the worst-case scenario is to remove the central node, which makes the network into an empty network.

Given this difficulty, privacy researchers initially focused on protecting the privacy of the edges rather than the node and edges together to increase the statistical inference. This form of differential privacy is known as *edge differential privacy*, which examines the absence or presence of a particular edge, instead of *node differential privacy*, which examines the absence or presence of a node and all the adjacent edges. However, under the edge differential privacy framework, the privacy research assumes that node attributes are known or not private, which is not always the case in real life.

Researchers have developed models to generate differentially private synthetic data that satisfy node differential privacy. Modeling certain features in the data is how some researchers are trying to tackle text and image data, but there are still both technical and philosophical challenges to overcome.

For example, consider a photo of a person in a suburban neighborhood. If we want to protect the image using differential privacy, simply applying noise to every pixel produces unusable gibberish for any amount of noise that is privacy protecting. Instead, we could extract a model that contains the person's hair color, their facial features, and so on. We could then computationally remove the person from the image and insert a new, synthetic person, with noise added to each of the image parameters.

While this approach might obscure a person's face, an attacker might infer the person's identity from other information in the image. If a house address is clearly visible in the picture background, a data intruder might use that to learn the person's identity. If we continue down this path, we might end up trying to synthesize the entire image, possibly protecting more information than really needs to be protected. To date, efforts along these lines are starting to appear in research venues, but they do not appear to be ready for use in production.

4.1.3. Making use of public data. Although the differential privacy model assumes that an attacker might have all of the information in the world, few differentially private mechanisms make use of publicly available data to improve their accuracy without causing additional privacy loss. This is somewhat frustrating, given that the amount of noise being added may be substantial to achieve specific levels of privacy protection. Approaches for making use of public data are thus a largely unrealized opportunity for researchers to develop new mechanisms that would improve accuracy and corresponding impact on privacy.

4.2. **Transition challenges.** As our previous example of public key cryptography demonstrates, the time to transition new mathematics technology from the academic to practice can be quite lengthy. Often this is because the clean, idealized conditions of the lab do not map well to the messy realities of the world.

4.2.1. Human resource limitations. Probably the largest barrier to the deployment of differential privacy is the lack of practitioners who are simultaneously knowledgeable of the math and skilled in building production

systems. Although there are a growing number of production-ready differential privacy libraries, such as Google's Privacy on Beam,[6] the process of designing, building, deploying, and maintaining a working differential privacy system requires far more than a vetted differentially private algorithms library. On the technical side, the practitioner must identify which algorithms to use and where to insert them into the statistical pipeline. Frequently, the practitioner must rework statistical computations to allow the ϵ to be used more efficiently.

Knowledgeable data practitioners must also defend the very decision to use differential privacy in the first place. This includes explaining why the release of exact aggregate statistics inherently compromises the privacy of individuals, and defending the decision to intentionally add error (e.g., noise) to data products after a substantial amount of funds have been spent to make the data products as accurate as possible.

4.2.2. Setting the privacy loss budget. One of the differential privacy's characteristics is that the protection is *tunable*, allowing the tuning for any given data release to be set by a policymaker, rather than a technician or dictated by the privacy mechanism. This approach is advantageous for differential privacy researchers, because it frees them from the need to make a policy choice. The mathematics are clear: someone's ox is going to get gored. For any setting of ϵ, it is always the case that the data could be more accurate by sacrificing just a little more privacy, and the release could be a little more privacy preserving by making the data a little less accurate. So how should this choice be made?

While setting the value of ϵ is clearly a policy decision, it is a policy decision for which little has been researched or otherwise written to aid policymakers. This lack of guidance results in many differentially private applications that have privacy loss budgets that *seem* unjustifiably high or low.

As an example of too high, Apple announced in 2016 that they would deploy a local differentially private approach on their iPhone to gain information about suggested emojis based on keyboard strokes (in iOS 10 and 11). Apple did not divulge many details, so external privacy researchers conducted experiments to better understand the methodology [TKB+17]. They discovered the privacy loss budget used to collect users' data on mobile devices to be a daily value of 4 and thus a monthly value of 120. The scientific community viewed these values as unreasonably high. For differential private systems that use high values of ϵ, a concern is data intruders could reconstruct the entire confidential data, especially when the number of queries are unbounded.

Conversely, many data practitioners criticize differentially private applications that set ϵ too low. For example, in 2018, the US Census Bureau announced their new Disclosure Avoidance System, a differentially private method called the TopDown Algorithm, for the 2020 Census. In preparation, they applied the TopDown Algorithm on the 1940 and 2010 Census data to allow data users to compare against the unaltered 1940 Census data[7] and the original 2010 Census data release. For these demonstration datasets, the Census Bureau set the privacy loss budget at 0.25, 0.5, 0.75, 1, 2, 4, 6, and 8. For another data demonstration release, the Census Bureau set the privacy loss budget to 4.5.

After analyzing the data, many census data users and researchers expressed concerns over the accuracy and usefulness of the altered data [RFMS19]. These concerns resulted in the National Academies of Science, Engineering, and Medicine hosting a workshop in December 2019, where many of the speakers presented specific data quality issues.[8] Based on feedback from the scientific community, the US Census Bureau set the privacy loss budget for the next 2020 Decennial Census demonstration data to 10.3 for the persons file and 1.9 for the housing units data, with $\delta = 10^{-10}$.[9] On June 9, 2021, the Census Bureau committed to 17.14 for the persons file and 2.47 for the housing units data.[10]

Although there are a growing number of differentially private applications to provide further context on setting an appropriate ϵ, the scientific community needs more use cases and participation from public policymakers to decide what the right balance between privacy and utility is.

4.2.3. Machine resource limitations. Another problem is that many data practitioners or users have difficulties implementing differentially private methods, because of the insufficient computational resources. These limitations hinder the accessibility for the average data user who might not have the proper computing equipment to run the methods or the background to hand-code them.

In the comparative study of the NIST PSCR Data Challenge, for instance, Bowen and Snoke evaluated the implementation ease and the computational feasibility of the

[6] *To learn more about Privacy on Beam, check out their website at* https://github.com/google/differential-privacy/tree/main/privacy-on-beam.

[7] *The 1940 Census data are publicly available due to Title 44, where census records that contained information about individuals could be released for public use after 72 years.*

[8] *Details for the "Workshop on 2020 Census Data Products: Data Needs and Privacy Considerations" can be found at* https://www.nationalacademies.org/event/12-11-2019/workshop-on-2020-census-data-products-data-needs-and-privacy-considerations.

[9] *The US Census Bureau announced this privacy loss budget in their newsletter on April 19, 2021.* https://content.govdelivery.com/accounts/USCENSUS/bulletins/2cdb999.

[10] *The US Census Bureau set the privacy loss budget on June 9, 2021.* https://www.census.gov/newsroom/press-releases/2021/2020-census-key-parameters.html.

top competitors' methods [BS21]. All but one competitor provided open-source code via GitHub to earn additional prize money, lessening the burden on users to create their own code. As for computational demand, some of the differentially private synthetic data methods were computationally complex, such as one approach that required significant memory and high-performance graphics processing units (GPUs) to run in a reasonable amount of time. The other methods that were less computationally demanding instead need additional hand-coding to properly pre-process the confidential data based on public data (if available).

Realizing the lack of accessible differential privacy code, a research group out of Harvard University started OpenDP[11] in 2019. OpenDP aims "...to build trustworthy, open-source software tools for statistical analysis of sensitive private data...[that] will offer the rigorous protections of differential privacy for the individuals who may be represented in confidential data and statistically valid methods of analysis for researchers who study the data." As of the publication of this article, OpenDP has partnered with Microsoft to create SmartNoise,[12] but the platform is still under development.

4.2.4. Accessing the confidential data. From experience and the privacy literature, we know that synthetic data, including both traditional synthetic data and new synthetic data created using differentially private methods, often do not provide sufficient accuracy to create accurate complex statistical models. To address this concern, the US government created the Federal Statistical Research Data Centers that allow appropriately vetted researchers to access the confidential data and release results after those results undergo an appropriate disclosure review process at a federal statistical agency. This process is often slow, laborious, and underutilized.

In an effort to alleviate some of the burdens on researchers, the US Census Bureau provides researchers with access to two experimental synthetic databases (which are not differentially private) via the Synthetic Data Server (SDS) at Cornell University: the Synthetic Longitudinal Business Database and the SIPP Synthetic Beta Data Product. The SDS provides a validation server that allows researchers to submit their statistical programs to run on the underlying administrative data after testing it on the synthetic data. However, the SDS has two disadvantages. First, because it is not automated, the process consumes limited staff time. Also, the demand often exceeds available staff time, causing long delays for approval. Second, reviews may be inconsistent since they are manually evaluated by humans, and they do not adhere to formal notions of privacy that constrain the allowable output.

To automate the process, the Urban Institute in collaboration with the Internal Revenue Service Statistics of Income Division is now developing an automated prototype validation server using differentially private methods. The proposed server will allow data users to submit queries to gain differentially private summary statistics and regression models. This validation server will expand access to administrative tax data, which is currently limited to those with Internal Revenue Service clearance. However, the effort will take many years to perfect given many of the challenges we already discussed.

4.3. Social challenges.

The adoption of differential privacy has encountered considerable resistance from various groups. The pushback mostly stems from the lack of communication and data users already having access to data that is now being restricted.

4.3.1. Communicating privacy realities. One of the greatest challenges in deploying differential privacy is communicating to policymakers and the public the underlying mathematics of privacy itself from what we have learned over the past two decades. Such communications are difficult because the protection of private facts behaves differently than our intuition has been trained to think. This is perhaps similar to how the underlying physics of mass and energy at the quantum level does not match the intuitive understanding for most of us living in the macroscopic world.

Modern conceptions of privacy are overwhelmingly legal and philosophical, rather than mathematical. For much of the twentieth century, Western legal traditions typically viewed information as either *private* or *public*. Since the 1960s, there has been an increasing recognition by policymakers that this is not the case. There is steady effort to impose a variety of use-based controls to limit the impact on individuals from various kinds of information.

In the US today, privacy is frequently viewed in terms of the confidentiality of communications, transactions, medical information, and other kinds of facts. Separately, another branch of privacy law deals with intrusion into a person's seclusion.[13] These conceptions drive the policy discussion and the actions of consumers and businesses. However, they are not rooted in mathematics. In other words, these conceptions are based on how we imagine and want personal information to behave and not how it actually behaves. Just as a spendthrift might actually wish to live on credit and never actually pay the bills when they come due, so might organizations wish to make use of personal information while not paying the price of the inherent privacy loss. The math of differential privacy tells us there is a real cost to every data release. There is a running bill, even if we do not choose to acknowledge it.

[11] *OpenDP's website can be found at* `https://opendp.org/`.

[12] *To learn more about SmartNoise, check out their website at* `https://smartnoise.org/`.

[13] *Daniel Solove's "A taxonomy of privacy" summarizes the range of privacy concerns in the United States.* `https://www.jstor.org/stable/40041279`.

4.3.2. Changing data access. Many of those who oppose differential privacy come from organizations that have either enjoyed unfettered access to high-quality data for decades or believed that past public data releases were largely unaltered (but somehow protected). Data users will have to adjust how they conduct their analyses when given access to new source data that are intentionally noisy, as a result of being protected.

Currently many data users do not know how to account for the differentially private changes due to resource limitations or because they have not yet engaged in the conversation. This problem comes hand in hand with the need for better communication among the various groups to best convey user needs to continue their analyses.

For example, while some data users need access to finer-grained ethnic groups, such as Japanese American versus Asian American, others might only need the general East Asian American grouping. The former group could use the SDS validation server to query the specific counts, whereas the latter could use the official statistical tables.

While some users perceive differential privacy as removing access to data, differential privacy offers the opposite. Several federal statistical agencies, such as the Internal Revenue Service, the Bureau of Labor Statistics, and the Bureau of Economic Analysis, are exploring differentially private algorithms to release data that were previously not available due to privacy concerns.

In order to efficiently use the privacy loss budget to expand data access, data custodians and data users must cooperate and coordinate deploying differentially private methods. No longer is it sufficient for data custodians to perform some kind of "light touch" and provide "safe data" to well-meaning academics or marketers, with the hope that the data will be used for safe projects, by safe people, in safe settings, and with safe outputs. That is, the so-called "five safes" framework is fundamentally flawed because mathematically there is no such thing as safe data [Rit17]. This mathematical truth is the core justification for the cost and difficulty associated with the adoption of differential privacy.

5. Differential Privacy's Future

Fifteen years of differential privacy have brought a generation of mathematicians difficult truths about the nature of privacy. For businesses and statisticians, one of the most troubling truths is that it is not possible to use personal information in a data product without leaking some of that information. For policymakers, it means that there is no magic set of data attributes or level of aggregation that allows for data releases that have no privacy impact. In practice, this means that policymakers can no longer pass off to technologists the requirement to balance the benefit of public data with the cost to individual privacy.

What does this mean for the future? Similar to the emergence of ssh and SSL to ensure a secure internet, we must keep testing and using differential privacy to better understand and tackle all the challenges and limitations we discussed earlier. If we do not, we risk losing access to invaluable data, as more organizations realize that traditional statistical disclosure limitation methods do not address modern data privacy concerns.

However, the bulk of differential privacy research is *still* theoretical. We need to encourage and value more usable applications paired with educational materials on how to explain and implement differentially private methods. By doing so, we will have a more constructive conversation with data users on adopting differential privacy for practical applications. We also need continuing education of policymakers and journalists, who to date have portrayed the fight between traditional disclosure avoidance systems and differential privacy as a battle between old and new systems, rather than as a battle between systems that use *ad hoc* methods and have no formal proof of their validity, and those based on principled mathematical proof.

As we have discussed throughout this article, differential privacy is not a "silver bullet" that makes data releases safe, nor is it a one-size-fits-all method. Differential privacy instead accounts for the privacy loss that individuals suffer when their private data are used in computations that are later released to the public. Differential privacy provides policymakers, data custodians, and data users a powerful tool for balancing the competing concerns of public benefit with individual privacy.

ACKNOWLEDGMENTS. The authors would like to thank John Abowd and Phil Leclerc at the US Census Bureau for their work developing, implementing, and promoting differential privacy for the publication of official statistics. Thank you to Daniel Goroff and the Alfred P. Sloan Foundation for their work in supporting the transition of academia to a regime in which data can be analyzed in a more privacy-centric manner such as the validation server project. We also thank Damien Desfontaines for providing valuable comments and suggestions that improved the paper. Thank you to the inventors of differential privacy, Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith, for many, many discussions about the math, meaning, and history of differential privacy.

References

[ABC+20] Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, John Davis, Damien Desfontaines, Alex Fabrikant, Evgeniy Gabrilovich, Krishna Gadepalli, Bryant Gipson, Miguel Guevara, Chaitanya Kamath, Mansi Kansal, Ali Lange, Chinmoy Mandayam, Andrew Oplinger, Christopher Pluntke, Thomas Roessler, Arran Schlosberg, Tomer Shekel, Swapnil Vispute, Mia Vu, Gregory Wellenius, Brian Williams, and Royce J. Wilson, *Google covid-19 community mobility reports: Anonymization process description (version 1.0)*, arXiv:2004.04145, 2020.

[BIK+17] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth, *Practical secure aggregation for privacy-preserving machine learning*, Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175–1191.

[BS21] Claire McKay Bowen and Joshua Snoke, *Comparative study of differentially private synthetic data algorithms from the nist pscr differential privacy synthetic data challenge*, J. Privacy and Confidentiality **11** (2021), no. 1.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor, *Our data, ourselves: privacy via distributed noise generation*, Advances in cryptology—EUROCRYPT 2006, Lecture Notes in Comput. Sci., vol. 4004, Springer, Berlin, 2006, pp. 486–503, DOI 10.1007/11761679_29. MR2423560

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith, *Calibrating noise to sensitivity in private data analysis*, Theory of cryptography, Lecture Notes in Comput. Sci., vol. 3876, Springer, Berlin, 2006, pp. 265–284, DOI 10.1007/11681878_14. MR2241676

[DN03] Irit Dinur and Kobbi Nissim, *Revealing information while preserving privacy*, Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, 2003, pp. 202–210.

[DR13] Cynthia Dwork and Aaron Roth, *The algorithmic foundations of differential privacy*, Found. Trends Theor. Comput. Sci. **9** (2013), no. 3-4, 211–487, DOI 10.1561/0400000042. MR3254020

[EPK14] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova, *Rappor: Randomized aggregatable privacy-preserving ordinal response*, Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, 2014, pp. 1054–1067.

[GAM19] Simson Garfinkel, John M. Abowd, and Christian Martindale, *Understanding database reconstruction attacks on public data*, Comm. ACM **62** (2019), no. 3, 46–53.

[Gar15] Simson Garfinkel, *De-identification of personally identifiable information*, Technical Report NIST IR 8053, National Institute of Science and Technology, 2015.

[JNS18] Noah Johnson, Joseph P. Near, and Dawn Song, *Towards practical differential privacy for SQL queries*, Proceedings of the VLDB Endowment **11** (2018), no. 5, 526–539.

[KM11] Daniel Kifer and Ashwin Machanavajjhala, *No free lunch in data privacy*, Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data, 2011, pp. 193–204.

[MKA+08] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber, *Privacy: Theory meets practice on the map*, Proceedings of the 2008 IEEE 24th International Conference on Data Engineering, 2008, pp. 277–286.

[RFMS19] Steven Ruggles, Catherine Fitch, Diana Magnuson, and Jonathan Schroeder, *Differential privacy and census data: Implications for social and economic research*, Aea papers and proceedings, 2019, pp. 403–408.

[Rit17] Felix Ritchie, *The 'Five Safes': a framework for planning, designing and evaluating data access solutions*, Zenodo, 2017.

[RSP+20] Ryan Rogers, Subbu Subramaniam, Sean Peng, David Durfee, Seunghyun Lee, Santosh Kumar Kancha, Shraddha Sahay, and Parvez Ahammad, *Linkedin's audience engagements api: A privacy preserving data analytics system at scale*, arXiv:2002.05839, 2020.

[Swe02] Latanya Sweeney, *k-anonymity: A model for protecting privacy*, Int. J. Uncertain. Fuzziness Knowl.-Based Syst. **10** (2002), no. 5, 557–570.

[TKB+17] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang, *Privacy loss in apple's implementation of differential privacy on macos 10.12*, arXiv:1709.02753, 2017.

[War65] S. L. Warner, *Randomised response: a survey technique for eliminating evasive answer bias*, J. American Statistical Association **60** (1965), no. 309, 63–69.

[WZFY20] Teng Wang, Xuefeng Zhang, Jingyu Feng, and Xinyu Yang, *A comprehensive survey on local differential privacy toward data statistics and analysis*, Sensors (Basel, Switzerland) **20** (2020Dec), no. 24, 7030 (eng).

Claire McKay Bowen

Simson Garfinkel

Credits

Opening image is courtesy of metamorworks via Getty.

Figure 1 is courtesy of Simson Garfinkel.

Photo of Claire McKay Bowen is courtesy of Donna Marion.

Photo of Simson Garfinkel is courtesy of the US Census Bureau.