Differential Privacy in Practice: Lessons Learned From 10 Years of Real-World Applications

Luis Del Vasto-Terrientes and David Sánchez | Universitat Rovira i Virgili
Josep Domingo-Ferrer | Universitat Rovira i Virgili and LAAS-CNRS and Université de Toulouse

Differential privacy (DP) is a widespread data protection mechanism. However, its application in real-world scenarios has been challenging. To shed some light on this, we offer a critical analysis of 21 DP deployments by top-tier companies and institutions over the past decade.

The mathematical foundation of ϵ-differential privacy (DP) hinges on the principle that the presence or absence of any record in a dataset should not influence the protected outcome up to a factor exponential in a parameter ϵ. If each record corresponds to an individual and ϵ is small, this means that the individual's information stays confidential. This robust privacy guarantee has attracted so much interest that DP has become the *de facto* data protection mechanism in areas such as data analytics, statistics, and machine learning.

Introduction

DP was conceived to mask responses to interactive aggregate queries against a remote database. In this scenario, DP behaves relatively well because the perturbation needed to enforce a low enough ϵ on the aggregate query responses can be small, and therefore, the accuracy—or the analytical utility—of the responses can be reasonably preserved. However, since this interactive setting (IS) imposes severe restrictions on the type and number of queries that can be answered, the research community has put a lot of effort into adapting the notion of DP to (the much more useful) non-ISs, such as data collection (DC) and data release (DR). Unfortunately, in these scenarios, DP struggles to achieve its strong privacy guarantee while keeping the protected outcome useful for analysis.

For all its popularity among researchers, the adoption of DP in the real world has been challenging. Several studies have examined the difficulties of implementing DP on the ground. For example, Dwork et al.[1] interviewed practitioners about barriers such as privacy loss and DP algorithmic limitations. Most of the issues revolved around the complexity of implementation and misunderstandings about DP parameters. Similarly, Garrido et al.[2] interviewed nine major companies, highlighting the need for better guidance in selecting ϵ and improved communication about DP mechanisms to nonexperts. A comprehensive list of DP applications is given at https://desfontain.es/blog/real-world-differential-privacy.html but without critical analysis.

We examine the practical outlook of DP from a different angle, with the aim of shedding light on its applicability from the perspective of practitioners. To this end, we critically analyze relevant real-world applications of DP in the last decade and systematically report relevant (and sometimes not clearly exposed) implementation details, such as the type of DP enforced and its privacy parameters. Our study not only covers cutting-edge applications but also encompasses past

Digital Object Identifier 10.1109/MSEC.2025.3578104

Copublished by the IEEE Computer and Reliability Societies

experiences that enable us to analyze how DP implementations have evolved over the years.

To provide a meaningful analysis, we considered applications that have been deployed in production environments by renowned top-tier IT companies and public statistical agencies. The applications should deal with large-scale datasets comprising at least millions of data points derived from real subjects. Furthermore, applications should have an impact on the business model or strategy of the implementer or influence public policy or decision making. Accessibility and transparency in the application design are also required (for example, privacy parameters, DC policies, etc.). In this sense, all of the information we report has been gathered from the most recent documentary sources provided by the primary actors, including white papers, reports, and blogs. Although we prioritize recent deployments, we also consider previous experiences that had a business, innovation, or public impact. Because of this focus, we excluded research applications—in particular, those that are theoretical, experimental, or with narrow or specialized academic scopes—and generic DP frameworks—whose outputs vary significantly according to user configuration.

As an outcome, we extract lessons learned that may help practitioners tone down their expectations regarding DP and raise awareness of the challenges they may face when implementing DP in practice. For some of these challenges, we also point out (existing but often overlooked) alternative solutions that can be better suited than DP for certain scenarios.

DP

DP's privacy parameter ϵ, also called the *privacy budget*, defines an upper bound on the privacy loss. The smaller the ϵ, the less noticeable in the released output is the presence or absence of any single record in the underlying original data. According to Dwork et al.,[1] taking $\epsilon \leq 1$ provides meaningful ex ante privacy *guarantees*. As ϵ increases beyond the suggested range, the privacy guarantees become weaker, or in the words of the DP inventors, "DP delivers privacy [guarantees] mostly in name."[1]

DP is usually enforced by adding to the query results noise that is inversely proportional to the budget ϵ and directly proportional to the sensitivity of the query result to the modification of any single record. This is known as *global sensitivity*.

Two approaches to enforcing DP can be distinguished according to who is responsible for running the DP mechanism. In the *central DP* model, the data curator is trustworthy and can access and protect the original raw data, while in *local DP* (LDP), the data owners locally run the DP mechanism on their own data. LDP provides additional privacy because the original data do not leave the owner's premises.

One of the most interesting properties of DP is composability, which applies to a sequence of DP mechanisms, such as consecutive queries executed over the same database. On the one hand, sequential composition establishes that for ϵ_1-DP and ϵ_2-DP mechanisms that access *nonindependent* datasets, the combination of their outputs satisfies $(\epsilon_1 + \epsilon_2)$-DP. On the other hand, parallel composition establishes that for ϵ_1-DP and ϵ_2-DP mechanisms that access *disjoint and independent* datasets, the combination of their outputs satisfies $\max(\epsilon_1, \epsilon_2)$-DP.

The aggregate queries usually considered in ISs—such as averages—tend to have relatively low sensitivities, especially for attributes with bounded domains (for example, age). These can be made differentially private with relatively small noise and return reasonably accurate answers. However, DP limits useful answers to a fixed number of queries; if m queries are to be answered on the same subject's data, sequential composition requires the overall ϵ to be split among the m queries, which entails an m-fold noise increase and the subsequent utility decrease. Moreover, since the noise is calibrated to the sensitivity of a particular query, queries with higher sensitivity cannot be answered with that noise. This constrains the usefulness of ISs, where data uses (that is, query types) are difficult to anticipate, and unconstrained querying on the same data is usually desirable.

In contrast, the much more flexible and widespread non-ISs—such as DR and DC—allow unconstrained access to *microdata*, that is, detailed records that correspond to individual respondents. When applying DP to these settings, each respondent record can be viewed as the output of an *identity query*, that is, a query that returns the value of a certain record from the micro-dataset. Unfortunately, the global sensitivity of identity queries is typically enormous because the attribute values in the record can vary within their entire domains. Hence, a correspondingly enormous amount of noise is required to enforce DP with reasonably safe ϵ. In fact, since the attribute values within a record refer to the same individual and sequential composition applies among them, sticking to the selected ϵ requires even more perturbation. This is likely to imply an unacceptable utility loss.

Due to the impossibility of obtaining usable data with a safe/small enough ϵ in non-ISs or for nonaggregate queries, a variety of relaxations of the DP definition have been proposed. Relaxations balance privacy guarantees against data utility, but this comes at the cost of a potential failure of the privacy guarantees DP offers. The best-known and most widely employed relaxations are the following.

- (ϵ, δ)-DP is also known as *approximate DP*, where parameter δ represents the probability that the upper bound on the privacy loss represented by ϵ is actually surpassed. This probability is closely tied to the size of the database. For meaningful privacy guarantees, one must take $\delta \ll 1/n$, where n is the number of records in the database; with $\delta = 1/n$ or more, (ϵ, δ)-DP would be compatible with leaking one or more records. This relaxation is the one chosen by most of the applications evaluated in this work.
- *Rényi DP* (*RDP*) is a refined relaxation of DP that leverages Rényi divergence, parameterized by α, in addition to ϵ. Compared to (ϵ, δ)-DP, RDP offers greater flexibility and more fine-grained control over privacy guarantees across multiple differentially private mechanisms. The composition of two mechanisms satisfying (α, ϵ_1)-RDP and (α, ϵ_2)-RDP results in an overall privacy guarantee $(\alpha, \epsilon_1 + \epsilon_2)$-RDP.
- *Zero-Concentrated DP* (*zCDP*) is also defined in terms of the Rényi divergence. However, its privacy guarantee is characterized by a single parameter $\rho \in (1, \infty)$, which constrains the values of ϵ and α. Its composition rule is straightforward; for two mechanisms ρ_1-zCDP and ρ_2-zCDP, the overall privacy guarantee is given by $(\rho_1 + \rho_2)$-zCDP.

RDP and zCDP are also attractive because they can be achieved with Gaussian noise, which is better understood than Laplace noise and more natural in many application domains. On the other hand, in addition to tighter accounting of privacy guarantees when composing multiple mechanisms, RDP and zCDP offer simpler composition rules than (ϵ, δ)-DP. However, their parameterizations can be converted to (ϵ, δ)-DP when needed. Therefore, for simplicity, we express the privacy guarantees of the surveyed applications in terms of (ϵ, δ)-DP.

The relaxations mentioned previously are just three among the dozens of derivative definitions of DP proposed in the literature. Their variety and complexity make it challenging for practitioners to fully understand the (privacy loss) implications of using them instead of pure unrelaxed DP. The confusion is further increased because most relaxations of DP, even though they only partially satisfy the privacy guarantee of DP, are still labeled "DP."

DRs

We next discuss relevant DP deployments, categorized by their data workflows. In this first category, we include noninteractive applications that aim at releasing fine-grained data, such as very disaggregated count data.

U.S. 2020 Decennial Census (U.S. Census Bureau)

The U.S. Census Bureau (USCB) is a government agency whose primary task is to conduct the U.S. decennial censuses. To prevent the reidentification of census respondents, since 1990, the USCB has implemented statistical disclosure control techniques, such as data swapping, top and bottom coding, and cell suppression.[3] However, for the release of the 2020 Decennial Census, the USCB chose to apply DP with the argument that new advances in computing technology increased the risk of reidentification.[a] In particular, it used the zCDP relaxation.

The USCB had previously employed DP in the *OnTheMap* application (2008),[4] which is recognized as the first DP deployment. In this application, DP was leveraged to synthetically build a privacy-protected statistical model from the USCB's Longitudinal Employer-Household Dynamics Program.

The latest 2020 release (March 2023) resulted in a total privacy budget of $\epsilon = 52.83$, which accounts for the budget accumulated by the composition of the nonindependent DP-protected data products released.[b,c] This significantly exceeds the recommendation of the DP inventors (that is, $\epsilon \leq 1$). On the good side, $\delta = 10^{-10}$ was set within acceptable bounds, considering that the size of the U.S. population was around 330 million (that is, $\delta < 1/330$ million).

Even with such a high ϵ, studies point to data inaccuracies that affect smaller communities and minority populations during redistricting.[5] This occurs because smaller counts require more noise to ensure privacy guarantees, resulting in greater distortions.

Full URLs Dataset (Facebook)

In April 2018, Facebook announced the release of the interactions of their users with public URLs. The first version of the dataset was released nearly two years later, after an important investment and a long delay.[d]

This dataset includes interactions between Facebook users and the URLs they publicly share, whether original posts or reshares. These interactions include user actions such as views, likes, shares, and comments

[a] https://www.census.gov/library/publications/2023/decennial/c2020br-03.html.
[b] https://www2.census.gov/programs-surveys/decennial/2020/program-management/data-product-planning/2010-demonstration-data-products/04-Demonstration_Data_Products_Suite/2023-04-03/2023-04-03_Privacy-Loss_Budget_Allocations.pdf.
[c] https://www2.census.gov/programs-surveys/decennial/2020/program-management/data-product-planning/2010-demonstration-data-products/04-Demonstration_Data_Products_Suite/2023-04-03/2023-04-03_Factsheet.pdf.
[d] https://socialscience.one/blog/public-statement-european-advisory-committee-social-science-one.

on those URLs. Other relevant statistics on users reporting posts containing spam and false news URLs are also provided. Privacy protection is enforced by adding DP-calibrated noise to the user counts for a given action. In addition, only URLs shared by at least 100 users are included in the release. This filters out URLs with low-frequency interactions, which are the ones that most likely reveal potentially sensitive user attributes.

Privacy in the released URL action counts was enforced via zCDP with a total privacy budget at the user level of $\epsilon = 1.617$ with $\delta = 10^{-5}$. According to Messing et al.,[6] δ was set based on the number of unique URL actions per user (≥ 100, with no maximum value specified). This value can be considered robust as it is highly unlikely that a single user would have interacted with a number of URLs larger than 10^5. Compared to other releases of very disaggregated count data (such as the aforementioned 2020 Decennial Census), Facebook's data are highly aggregated; tables contain only user action counts, which have a small global sensitivity, and they are preprocessed to remove URLs with spare interactions. This is why this release can use a small ϵ value while keeping the added perturbation at bay.

This dataset has been updated 10 times. Each update released statistics on the users' actions gathered since the previous release. This is adequate because releasing DP-protected updates on the *same* data would rapidly dilute privacy guarantees. However, data across different updates were considered independent and therefore not subjected to sequential composition. In fact, such independence is questionable since the successive behaviors of users in their URL interactions tend to be correlated in the midterm (that is, several months). In the worst case, the effective ϵ would accumulate over time due to sequential composition, resulting in a greater privacy loss. This accumulated ϵ is also referred to as a *unit of privacy*.

Regarding data utility, Allen et al.[7] examined potential bias in the Facebook dataset and concluded that censoring data (for example, discarding URLs with fewer than 100 interactions) may introduce bias even in large-scale datasets, potentially overestimating the shares of clicked news and fake news. Hence, the calculated proportions of fake news might be significantly higher than they would be if data from all URLs were considered.

DC

This type of noninteractive application performs continuous DC from data subjects (for example, the company's end users).

Privacy at Scale (Apple)

Apple used DP for user DC on both iOS and macOS in 2017.[e] The process begins on the client side by removing device identifiers from the data. Then, LDP (in its pure form) is applied on the remaining client data. LDP-protected data are transmitted through an encrypted channel to a central server, where they are finally aggregated. Both the ingestion and aggregation stages occur within a secure access environment. The aggregator creates DP histograms for different use cases, which can then be safely shared with the interested teams within the company.

Apple used this mechanism in a variety of applications, including lookup hints, emoji suggestions, health types, QuickType, Safari domains identified as causing high-energy use or crashes, and Safari-autoplay intent detection. Although the same DP architecture is used in all applications, the base ϵ depends on the application, ranging from 2.0 to 8.0. However, since data can be collected up to twice a day for some applications and data are retained up to a maximum of three months, the accumulated values ϵ are much larger (see Table 1). For instance, even for a small base privacy budget (that is, $\epsilon = 2$), if users' data are collected twice a day and retained for three months, the effective ϵ would reach 360.

Learning Iconic Scenes (Apple)

In 2023, Apple expanded its portfolio of DP applications to include photos taken in frequently visited locations. For this newer approach, Apple combined LDP with *secure aggregation* to enhance privacy. Specifically, (ϵ, δ)-DP was used instead of pure DP for local noise addition. When a photo is taken, a random location annotation is selected from a set of possible annotation pairs (for example, <Central Park, Recreation>) and encoded into a one-hot vector. DP is then applied by flipping each bit with a certain probability. The result is divided into two shares. Each share, by itself, is just meaningless noise, but when combined, they reconstruct the original noised vector. When a cohort of 150,000 devices has shared data, the decryption and aggregation process begins.

The learned histograms in this architecture satisfy (ϵ, δ)-DP with $\epsilon = 1$ and $\delta = 1.5 \cdot 10^{-7}$. To limit the effect of sequential composition, only one category from one picture per device is selected at random. This restriction on continuous DC allows this application to use smaller ϵ than the previous ones.

It is reported that 4.5 million <location, category > pairs were discovered. If each device provides one pair (as stated previously), we may conclude that

[e] https://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf.

Table 1. The characterization of base and accumulated privacy budgets of the analyzed DP applications.

Organization	Application	Year	DP Model	DP Type	ϵ	Accumulated ϵ
			DRs			
USCB	Census 2020	2023	zCDP, $\delta = 10^{-10}$	Central	52.83	N/A
Facebook	Protected URLs	2023	zCDP, $\delta = 10^{-5}$	Central	1.617	16.17^{+}
			DC			
Apple	Lookup hints	2017	ϵ-DP	Local	4	720
Apple	Emoji suggestions	2017	ϵ-DP	Local	4	360
Apple	Health types	2017	ϵ-DP	Local	2	180
Apple	QuickType	2017	ϵ-DP	Local	8	1,440
Apple	Safari domain	2017	ϵ-DP	Local	4	720
Apple	Safari-autoplay detection	2017	ϵ-DP	Local	8	1,440
Apple	Iconic Scenes	2023	(ϵ, δ)-DP, $\delta = 1.5e^{-7}$	Local	1	1
Google	RAPPOR (Chrome)	2014	ϵ-DP	Local	0.5343	16.03^{*}
Google	RAPPOR (Windows)	2014	ϵ-DP	Local	1.0743	32.23^{*}
Google	Mobility Reports	2020	ϵ-DP	Central	2.64	79.2^{*}
Google	Search Insights	2021	(ϵ, δ)-DP, $\delta = 10^{-5}$	Central	2.19	65.7^{*}
Microsoft	Telemetry data	2017	ϵ-DP	Local	1.672	207.328
Microsoft	Assistive AI	2020	(ϵ, δ)-DP, $\delta = 10^{-7}$	Central	4	Unknown
			REP			
Google	DP-SQLP (Google Shopping)	2024	(ϵ, δ)-DP, zCDP, $\delta = 10^{-9}$	Central	1	Unknown
Google	DP-SQLP (Google Trends)	2024	(ϵ, δ)-DP, zCDP, $\delta = 10^{-10}$	Central	2	Unknown
LinkedIn	Labor Market Insight – Hiring	2020	(ϵ, δ)-DP, $\delta = 10^{-10}$	Central	1.2	14.4
LinkedIn	Labor Market Insight – Avg. jobs	2020	(ϵ, δ)-DP, $\delta = 10^{-10}$	Central	1.2	14.4
LinkedIn	Labor Market Insight – Skills	2020	(ϵ, δ)-DP, $\delta = 10^{-10}$	Central	0.1	0.3
			IS			
LinkedIn	Audience Engagement API	2020	(ϵ, δ)-DP, $\delta = 10^{-10}$	Central	0.15	34.9

*"Year" corresponds to the (latest) release of the application that has been analyzed; *, the monthly budget (maximum DC period not provided); +, the accumulated budget for all 10 dataset updates (worst case); Avg.: average; REP: data reporting.*

4.5 million is the size of the population. Even though $\delta = 1.5{\cdot}10^{-7} < 1/4.5$ million, the actual value is still quite close to the upper bound and therefore could result in leakages as the population size increases.

RAPPOR (Google)

Google's Randomized Aggregatable Privacy-Preserving Ordinal Response (RAPPOR) is an application for anonymous DC from end users.[8] It uses the classical randomized response (RR) mechanism, by which true answers from end users are sent with probability p and false answers are provided with probability $1 - p$.

A *permanent* RR creates a noisy answer that is memoized by the client and reused permanently instead of the actual answer. Subsequently, an *instantaneous* RR reports on the noisy answer over time to the server, gradually revealing it completely.

The fact that the randomizations revealed over time to the server have not been computed on the original value but on a permanently randomized value protects against privacy degradation. RAPPOR was tested in various environments, including collecting usage statistics from Chrome users and evaluating Windows operating system (OS) process names.

Regarding Chrome, RAPPOR was used to collect daily data from approximately 14 million Chrome users, with a privacy budget $\epsilon = 0.5343$ for daily statistics. For the Windows Process Name use case, up to 186,000 reports from 10,000 computers were collected daily, using a privacy budget $\epsilon = 1.0743$. Although the daily values appear to ensure strong privacy, the privacy unit accumulates over time. For example, monthly ϵ values are double-digit in both use cases (see Table 1). Due to the constraints imposed by pure RAPPOR/DP, Google replaced it with a system that uses mix nets as a "privacy amplification," thereby allowing the amount of noise added to the raw statistics to be significantly lowered.

COVID-19 Community Mobility Reports (Google)

In 2020, Google introduced the COVID-19 Community Mobility Reports,[9] designed to provide information on changes in mobility patterns to help combat the COVID-19 pandemic. Google collected data from users' Location History Up to seven different categories (for example, recreation, pharmacies, and groceries) were tracked at various geographical levels. Each user could contribute up to four <category, location> pairs daily at the country level. Each place visited by a user was protected with $\epsilon = 0.44$, leading to a daily maximum $\epsilon = 1.76$ for the four <category, location> pairs. However, when aggregating the privacy budgets for residential and workplace locations, each with $\epsilon = 0.44$ and one daily contribution, the privacy budget increases to a daily $\epsilon = 1.76 + 0.44 + 0.44 = 2.64$ and a monthly ϵ as high as 79.2. Once again, the accumulation of single-digit ϵ values over time results in weak privacy guarantees.

COVID-19 Vaccination Search Insights (Google)

Google also provided anonymized trends compiled from searches related to COVID-19 vaccination.[10] They offered a dashboard and a complete dataset where counts are grouped by the week each query was performed, associated with the triple <week, region, and category>. Data were aggregated at three levels: state, county, and postal code.

The daily search activity data of users were masked with DP parameters $\epsilon = 2.19$ and $\delta = 10^{-5}$, which results in a monthly budget $\epsilon = 65.7$. Considering that user data are grouped according to population size, $\delta = 10^{-5}$ is barely sufficient for small counties and is inadequate for larger population sizes (that is, >100,000).

Telemetry Data (Microsoft)

In Ding et al.,[11] LDP and memoization are considered to collect DP telemetry data from Windows Insiders during the Windows 10 Fall Creators Update, similar to the technique presented in RAPPOR.[8] However, this application presents a notable privacy limitation for telemetry numerical counters; small but frequent updates can potentially reveal sensitive user patterns over time.

To improve privacy in this scenario, the values are discretized. This generalization ensures that individual observations reveal less precise information, making it harder for the data collector to infer sensitive details from repeated responses. As a downside, discretization causes a loss of precision and may hamper correlation analyses. Note that this precision loss adds to the one caused by DP noise perturbation.

The objective of the application was to collect data on application time usage from 3 million users to analyze behavior patterns. A single DC round is performed with a privacy budget $\epsilon = 1.672$, with new rounds executed every 6 h for 31 days. Although this base budget is smaller than in other DP implementations, it accumulates up to $\epsilon = 6.688$ daily and can reach a maximum of $\epsilon = 207.328$ monthly. In scenarios where higher collection frequencies are needed, the accumulated budget can be even higher.

Assistive AI (Microsoft)

Assistive AI aims to improve responses to e-mails or text messages by generating predefined response options.[f] According to the source, the system utilizes a deep neural network trained on hundreds of millions of message-reply (MR) pairs.

Even if the information on the implementation of DP in this application is scarce, it is mentioned that, from millions of MR pairs, short popular responses are initially selected "as-is" for analysis. DP is then applied to further limit the set of possible responses that can be revealed. The privacy parameters chosen are $\epsilon = 4.0$ and $\delta = 10^{-7}$. Unfortunately, the lack of details (on how often per-user data are collected, how popular responses are filtered, and the potential correlations between data collected from per-user e-mails and chats) prevents an accurate evaluation of this application.

REP

The applications in this category release noninteractive statistical reports, such as histograms.

[f] https://www.microsoft.com/en-us/research/articles/assistive-ai-makes-replying-easier-2/.

DP Stream Processing at Scale (Google)

In Zhang et al.,[12] Google recently presented DP SQL Pipelines (DP-SQLP), a scalable stream aggregation system to release DP histograms from raw data with unknown domains and continual observations. The system provides user-level privacy so that all actions (that is, user events) are protected simultaneously.

DP-SQLP enforces (ϵ, δ)-DP on the number of unique users who can contribute. Also, if there are enough user contributions for the requested histogram, zCDP is enforced on the user contributions.

The system has been deployed on Google Shopping and Google Trends. For Google Shopping, DP-SQLP was used to anonymize the page-view counts that are used to prioritize web page crawling. However, the DP implementation caused an overall accuracy loss of 39% for page-view data, which is significant.

The user data employed to generate the DP histograms are limited to a daily contribution protected with $(\epsilon = 1, \delta = 10^{-9})$-DP. However, the maximum number of contributions for a single user across days or the time period considered is not disclosed.

Google Trends, on the other hand, analyzes the interests of search queries for particular topics. In this case, user contributions are limited to one event per query and are protected with $(\epsilon = 2, \delta = 10^{-10})$-DP. Again, no information on the maximum number of contributions considered (per day or in total) is provided. Although the baseline ϵ is reasonably low, the lack of information on the total number of user data points considered to build the histograms prevents assessing the budget accumulated after sequential composition.

Labor Market Insight (LinkedIn)

During the COVID-19 pandemic, LinkedIn published data tracking labor market trends to help governments and industries make informed decisions.[13] We consider this application noninteractive data reporting because aggregate data are represented in static sets of visualization tools and do not offer interactive features.

The data released focused on the following three main aspects:

1. Who is hiring?
2. What jobs are available?
3. What skills are needed?

The "Who is hiring?" service includes each hiring event in four different monthly reports: country level, region level, country-industry level, and region-industry level. Therefore, as described in Rogers et al.,[13] given that each report fulfills $(\epsilon = 1.2, \delta = 10^{-10})$-DP for each user, the combination of the four reports results in a cumulative $(\epsilon = 4.8, \delta = 10^{-10})$-DP per month. Also, considering that a monthly histogram is computed from the last three months, the privacy budget increases to $\epsilon = 14.4$ because each user's data will be used in three consecutive monthly reports. For "What jobs are available?," the same DP parameterization applies. Even with relatively low values of ϵ per report, the overall monthly privacy budget can become significant. Furthermore, if consecutive reports must be generated sequentially (for example, quarterly and annual reports), the applied cumulative privacy loss may be impractical for implementation.

Finally, the skills required for these jobs are published with a robust $(\epsilon = 0.1, \delta = 10^{-10})$-DP thanks to the very low global sensitivity of the skill values w.r.t. individuals' hiring events. Similarly to the U.S. 2020 Decennial Census, we can consider $\delta = 10^{-10}$ an acceptably small value if the focus is on the U.S. population, although the population size is not explicitly defined.

IS

This is the scenario DP was designed for, in which DP is enforced on the answers to interactive aggregate queries performed against a remote database.

Audience Engagement API (LinkedIn)

The Audience Engagement API serves as a platform for marketers to gain consolidated insights into members' interactions with content.[14] The application programming interface (API) offers DP-protected interactive histograms and top-k results tailored to the user's requirements.

The number of queries analysts can perform is limited according to the overall privacy loss accumulated. Each query provides a privacy guarantee of $(\epsilon = 0.15, \delta = 10^{-10})$-DP, while the monthly privacy budget is set to $\epsilon = 34.9$. Although the monthly privacy budget allowed is considerably high, the IS allows calculating the cumulated ϵ and refusing to answer further queries as soon as the cumulated budget is deemed too high.

Discussion

Table 1 summarizes and compares the applications analyzed, highlighting the privacy parameters they employ. For continuous DC/DR and interactive applications, we report both the base privacy budget ϵ and the accumulated ϵ (unit of privacy), the latter being calculated by sequential composition according to the collection/querying constraints imposed by each application. It should be noted that while the base ϵ was clearly advertised in the information sources of applications, the accumulated ϵ was not communicated in most cases and had to be calculated manually by carefully analyzing

the specificities of each application (to the extent that such information was available). When relaxations involving δ are used, this parameter also increases with composition. However, the influence of δ, which represents the "probability" of a breach of privacy under ϵ, is less relevant when the values ϵ themselves are already too loose for the DP guarantee to hold. For these reasons, we do not detail the accumulated values of δ.

DP is a neat privacy definition that offers a strict and ex ante privacy guarantee but also demands a neat implementation to keep that guarantee. Otherwise, as mentioned previously, the DP guarantee is mostly nominal.[1] Unfortunately, the latter seems to be the case for the vast majority of discussed applications; out of the 21 applications, only LinkedIn's Skills for Labor Market Insight, Apple's Iconic Scenes, Google's Stream Processing, and the Facebook dataset *apparently* employed reasonably safe ϵ. These applications achieved this either by dealing with data so aggregated that even releasing them unprotected would hardly be privacy disclosive (Facebook and LinkedIn) or by severely constraining DC (Apple and Google). The latter has the undesirable side effect of affecting longitudinal data analysis. Moreover, Facebook did not account for the privacy loss incurred by publishing successive data updates from nondisjoint sets of users, and Apple employed the (ϵ, δ)-DP relaxation with an unreasonably large δ, which entails a sizable probability of not satisfying the ϵ-DP guarantee.

For the other applications using two-, three- or four-digit accumulated ϵ, the protection obtained is comparable to using plain noise addition, sometimes very mild. This is a statistical disclosure control technique that, due to its lack of ex ante privacy guarantees, requires an ex post risk evaluation, which is something very rarely done for the analyzed applications. In this respect, most applications avoid the ex post risk assessment by invoking the alleged DP guarantee, although such a guarantee only exists in name.

Notice that the previous discussion focuses on the accumulated ϵ, which is the one that measures the effective privacy achieved. Not only was this value not properly communicated in most applications, but their documentation often seems to minimize or simply ignore the effect of continuous DC and incremental releases of nondisjoint data.

Pure DP was used mainly in older implementations. More recent applications replace it with relaxations. Although relaxations allow DP guarantees to be breached with a certain probability, their implementations indicate that large privacy budgets are still needed to achieve reasonable utility.

Several applications—such as Facebook's, Microsoft's, or Apple's Iconic Scenes—combined (weak) DP with additional protection techniques—filtering of the rarest and most disclosive data, data discretization, or secure aggregation, respectively. In these cases, most of the protection came from the additional techniques applied rather than from DP. However, the credit for the protection achieved was attributed to DP.

All applications, except for LinkedIn's Audience Engagement API, focus on non-ISs. This shows that these settings are far more desirable (due to their flexibility and lack of constraints regarding secondary uses) than interactive ones. Given that DP was not designed for non-ISs, it is not surprising that engineers struggled to apply it to them.

Lessons Learned

Next, we gather lessons learned from the previous analysis and the experiences shared by the application designers in the respective information sources. For some items, we point out existing alternative solutions that might be more suitable than DP to solve specific problems.

- Practical applications have transitioned from using pure DP (applications before 2020) to a variety of relaxations in an attempt to reap the much-needed data utility and/or lower the privacy budget. But relaxations should not be accounted as "DP" because the privacy guarantees they offer may be significantly different/weaker than those expected from pure DP. In fact, there is a tendency, both in academia and in industry, to label any DP-like implementation as "DP protection" (therefore implying DP guarantees) even if it uses relaxations or very weak DP parameters for which the DP guarantee is ineffectual. This should be avoided as it misleads one about the actual privacy protection enforced.
- The implementation of DP via noise addition is always *perturbative*, meaning that the protected outcomes do not truthfully correspond to the original data due to the *random* noise added, no matter how small that noise is. In some critical domains where truthfulness is vital, this may undermine conclusions, as shown in Kenny et al.[5] In the statistical disclosure control literature, one can find a variety of nonperturbative masking techniques, such as data suppression, sampling, or generalization, which keep protected data truthful. Specifically, even though partial suppression and sampling suppress some outputs, those that are reported coincide with the original values; generalization, on the other hand, may report less detailed categories, but the generalized categories remain truthful. Some of these techniques can be enforced under the scope of privacy models such as *k*-anonymity, which shows that ex ante privacy guarantees are compatible with truthful outcomes.

- Protecting aggregate queries (as in LinkedIn's interactive application) is the ideal scenario for DP as it is precisely the one it was designed for. Aggregate queries provide inherent protection, which DP can transform into a formal guarantee with relatively little noise addition. Still, it should be noted that 1) the query types to be supported must be specified in advance because the noise required by DP depends on each query's global sensitivity, and 2) the number of times a query can be performed is limited by the privacy budget.
- Micro-DRs are challenging for DP due to large global sensitivities associated with detailed data publishing requiring significant noise and resulting in unacceptable utility loss. An alternative model actually designed for data publication like k-anonymity is better suited for this scenario: 1) k-anonymity provides a privacy guarantee (probability of reidentification at most $1/k$) that can be intuitively understood regardless of the value k, 2) it allows better privacy/utility tradeoffs than DP in return for making assumptions about the attackers' knowledge, and 3) it uses nonrandom masking for disclosive records, resulting in less distortion and better utility. Notice that the weaknesses typically attributed to k-anonymity, that is, dependency on the choice of the attributes to be masked (quasi-identifiers), and attributed disclosure due to the potential lack of diversity of the confidential attributes, can be respectively addressed by 1) masking *all* attributes (as DP does, with the consequent loss of utility) and 2) resorting to k-anonymity extensions offering attribute disclosure protection (such as l-diversity and t-closeness). Nevertheless, k-anonymity lacks DP's composability, which makes it suitable only for static (that is, nonincremental) DRs.
- In continuous DC, the accumulated ϵ should reflect both the base ϵ per collection and the frequency of overlapping collections. Since DP limits repeated sampling from the same subjects, it proves unsuitable for longitudinal analyses. Federated learning may be a more suitable solution if the ultimate goal of DC is to train machine learning models as it has the advantage that only updates of models trained locally by the data owners are transmitted. This makes it a more privacy-preserving alternative to centralized raw DC, especially when combined with other protection techniques such as secure aggregation.
- For DRs, the privacy budget should consider previous and planned incremental releases, which decrease the privacy of individuals. Similarly to continuous DC, DP is not well suited for this scenario, but neither is any other method in the privacy literature. These are open privacy problems that are likely to be hard to solve.

Today, the effectiveness of DP is mostly limited to scenarios similar to those for which it was originally designed, that is, protecting aggregate responses. However, according to the deployments analyzed, these scenarios represent only a small fraction of what industry and society demand.

Having observed the issues and long development periods incurred during the application of DP to real-world scenarios, we can conclude that DP needs careful evaluation and design for each use case. In particular, DP requires per-record and per-attribute domain analyses to calibrate the perturbation needed to attain the desired privacy guarantee. Moreover, quoting the inventors of DP,[1] there is still "no clear consensus on how to choose ϵ, nor agreement on how to approach this and other key implementation decisions."

In this respect, we recall that the DP guarantee establishes a mathematical *upper* bound on how much the inclusion of an individual's data can influence the results of any analysis, that is, it limits the individual's privacy loss relative to ϵ. This is especially true for sequential composition, which accounts for a worst-case scenario in which *all* data are assumed nonindependent. In practice, this means that sufficiently private results *might* be achieved for larger (accumulated) ϵ, but this is uncertain because privacy is no longer *guaranteed* by DP. In these cases, the protection comes from the perturbation introduced by the noise addition mechanism rather than the DP guarantee. Therefore, the residual disclosure risk should be empirically evaluated ex post on an application basis, as has been a common practice in the statistical disclosure control community for decades.[3] This has not yet been done frequently enough outside of that community. ■

Acknowledgment

Partial support for this work has been received from the Government of Catalonia (ICREA Acadèmia Prizes to J. Domingo-Ferrer and to D. Sánchez and Grant 2021SGR-00115), MCIN/AEI/10.13039/501100011033 and "ERDF A way of making Europe" under Grant PID2021-123637NB-I00 "CURLING," the EU's NextGenerationEU/PRTR via INCIBE (project "HERMES" and INCIBE-URV cybersecurity chair), and CIMI-Université de Toulouse.

References

1. C. Dwork, N. Kohli, and D. Mulligan, "Differential privacy in practice: Expose your epsilons!," *J. Privacy Confidentiality*, vol. 9, no. 2, pp. 1–22, 2019, doi: 10.29012/jpc.689.
2. G. Garrido, X. Liu, F. Matthes, and D. Song, "Lessons learned: Surveying the practicality of differential privacy in the industry," *Proc. Privacy Enhancing Technol.*, vol. 2023, no. 2, pp. 151–170, 2023, doi: 10.56553/popets-2023-0045.

3. A. Hundepool et al., *Statistical Disclosure Control*. Chichester, U.K.: Wiley, 2012.
4. A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber, "Privacy: Theory meets practice on the map," in *Proc. IEEE 24th Int. Conf. Data Eng.*, 2008, pp. 277–286, doi: 10.1109/ICDE.2008.4497436.
5. C. T. Kenny, C. McCartan, S. Kuriwaki, T. Simko, and K. Imai, "Evaluating bias and noise induced by the U.S. Census Bureau's privacy protection methods," *Sci. Adv.*, vol. 10, no. 18, 2024, Art. no. eadl2524, doi: 10.1126/sciadv.adl2524.
6. S. Messing et al., *Facebook Privacy-Protected Full URLs Data Set*. (2020). Harvard Dataverse. Online]. Available: https://doi.org/10.7910/DVN/TDOAPG
7. J. Allen, M. Mobius, D. Rothschild, and D. Watts, "Research note: Examining potential bias in large-scale censored data," *Harvard Kennedy School Misinformation Rev.*, vol. 1, Jul. 26, 2021. [Online]. Available: https://misinforeview.hks.harvard.edu/article/research-note-examining-potential-bias-in-large-scale-censored-data/
8. U. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA: ACM, 2014, pp. 1054–1067, doi: 10.1145/2660267.2660348.
9. A. Aktay et al., "Google COVID-19 community mobility reports: Anonymization process description (version 1.1)," 2020, *arXiv:2004.04145*.
10. S. Bavadekar et al., "Google COVID-19 vaccination search insights: Anonymization process description," 2021, *arXiv:2107.01179*.
11. B. Ding, J. Kulkarni, and S. Yekhanin, "Collecting telemetry data privately," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst. (NIPS)*, Red Hook, NY, USA: Curran, 2017, pp. 3574–3583.
12. B. Zhang et al., "Differentially private stream processing at scale," 2024, *arXiv:2303.18086*.
13. R. Rogers et al., "A members first approach to enabling LinkedIn's labor market insights at scale," 2020, *arXiv:2010.13981*.
14. R. Rogers et al., "LinkedIn's audience engagements API: A privacy preserving data analytics system at scale," *J. Privacy Confidentiality*, vol. 11, no. 3, pp. 1–27, 2021, doi: 10.29012/jpc.782.

Luis Del Vasto-Terrientes is a postdoctoral researcher at the Universitat Rovira i Virgili, CYBERCAT-Center for Cybersecurity Research of Catalonia, 43007 Tarragona, Catalonia. His research interests include decision making and data privacy. Del Vasto-Terrientes received a Ph.D. in computer science from Universitat Rovira i Virgili. Contact him at luismiguel.delvasto@fundacio.urv.cat.

David Sánchez is a full professor and an ICREA-Acadèmia research professor at the Universitat Rovira i Virgili, CYBERCAT-Center for Cybersecurity Research of Catalonia, 43007 Tarragona, Catalonia. His research interests include data semantics, machine learning, and data privacy. Sánchez received a Ph.D. in computer science from the Technical University of Catalonia. He is a Senior Member of IEEE. Contact him at david.sanchez@urv.cat.

Josep Domingo-Ferrer is a distinguished full professor of computer science and an ICREA-Acadèmia research professor at Universitat Rovira i Virgili, CYBERCAT-Center for Cybersecurity Research of Catalonia, 43007 Tarragona, Catalonia, where he also leads CYBERCAT. He is currently an invited professor at LAAS-CNRS, Université de Toulouse, 31400 Toulouse, France. His research interests include data privacy, data security, and ethics in IT. Domingo-Ferrer received an M.A. in philosophy at Université Paris Nanterre. He is a Fellow of IEEE. Contact him at josep.domingo@urv.cat.