CEO's Guide to **Differential Privacy**

GeorgianPartners

Table of Contents

Executive Summary

In a world where the risks and costs associated with privacy are on the rise, differential privacy offers a solution.

1. Differential privacy is a mathematical definition for the privacy loss that results to individuals when their private information is used to create an AI product.

2. Specifically, differential privacy measures how effective a particular privacy technique — such as inserting random noise into a dataset — is at protecting the privacy of individuals within that dataset.

3. This mathematical measurement of privacy can then be used to build trust with customers, in turn increasing customer willingness to participate in data aggregation.

4. The aggregation of data is critical for many machine learning projects because model performance improves with increases in the data available for training.

5. Data aggregation also solves the cold-start problem, i.e., not being able to immediately provide insights to a new customer because you have to wait for historical data.

6. Adopting differential privacy early will generate cumulative benefits over time and provide a differentiating advantage over competitors.

Introduction

Privacy risks are on the rise.

In recent years the growth of data collected on individuals has been explosive. Data now flows from every business transaction, from our mobile and social media activity, and increasingly from sensors embedded into our daily lives. While this data can be a treasure trove of possibilities, enabling everything from personalized shopping to personalized healthcare, it also brings with it significant privacy risks both for individuals and for the organizations collecting and using that data.

One solution for reducing that privacy risk is differential privacy, a mathematical definition for the privacy loss that results to individuals when their private information is used to create an AI product. A relatively new field, differential privacy has been the domain of academics and researchers since its inception a decade ago. Today, it's finding its way into products from Apple, Google and Uber, along with those of innovative earlier-stage companies.

Companies that implement differentially private AI solutions can mathematically enforce and demonstrate privacy guarantees. Those privacy guarantees can then be used to build greater trust with customers. The result is an increase in willingness among customers to share their data and have it included in an aggregated dataset.

Aggregated datasets are important because when used to train machine learning models they can result in improved model performance thanks to the larger amount of data available. Importantly, that aggregated data can also be used to solve the cold-start problem, i.e., the challenge of providing insights to a new customer when you don't yet have sufficient historical data for that customer.

Big technology companies such as Apple, Google and Uber have already recognized the underlying value of differential privacy and are implementing it within their various products and services. Smaller companies are starting to follow their lead as they seek to get access to richer, aggregated datasets on which to train machine learning models. If you're the chief executive of a startup that collects and uses customer data — whether directly from consumers or from other businesses — differential privacy is a subject that you need to understand.

Differential privacy is a mathematical definition for the privacy loss that results to individuals when their private information is used to create an AI product.

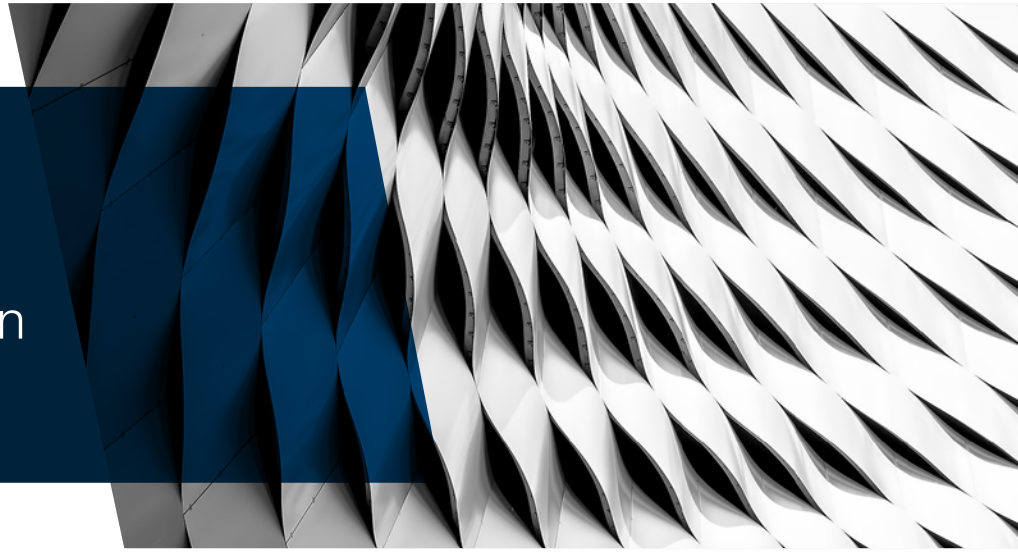

The De-Anonymization Problem

Differential privacy is a direct answer to the problem of de-anonymization, which is what happens when enterprising individuals — whether benevolent researchers or nefarious hackers — attempt to reassemble user identities by matching up disparate datasets. Unfortunately, de-anonymizing data is easier than you might think.

One well-known example of de-anonymization in the public domain involved Netflix's effort to crowdsource a better recommendation system back in 2007. To that end, Netflix released 10 million movie ratings sourced from half a million customers, with each person's name and personal details removed.

Researchers at the University of Texas at Austin subsequently set out to prove that such information could be de-anonymized. The researchers compared the released data with several dozen reviews that people had left on the Internet Movie Database (IMDb) and successfully reidentified two customers. "Releasing the data and just removing the names does nothing for privacy," noted professor Vitaly Shmatikov at the time.[1] "If you know their name and a few records, then you can identify that person in the other [private] database."

Another well-known case took place in 1997, when the Massachusetts Group Insurance Commission released anonymized data on state employees showing all of their hospital visits. Names, addresses and Social Security numbers were removed. Like Netflix, the commission hoped to benefit from crowdsourcing. The goal in this case was to uncover better insights into how the hospital system worked.

Latanya Sweeney, a graduate student at the Massachusetts Institute of Technology, quickly showed how easy it was to decode the data:

"At the time GIC released the data, William Weld, then Governor of Massachusetts, assured the public that GIC had protected patient privacy by deleting identifiers. In response, then-graduate student Sweeney started hunting for the Governor's hospital records in the GIC data. She knew that Governor Weld resided in Cambridge, Massachusetts, a city of 54,000 residents and seven ZIP codes. For twenty dollars, she purchased the complete voter rolls from the city of Cambridge, a database containing, among other things, the name, address, ZIP code, birth date, and sex of every voter. By combining this data with the GIC records, Sweeney found Governor Weld with ease. Only six people in Cambridge shared his birth date, only three of them men, and of them, only he lived in his ZIP code. In a theatrical flourish, Dr. Sweeney sent the Governor's health records (which included diagnoses and prescriptions) to his office.[2]"

1 "Researchers reverse Netflix anonymization," SecurityFocus, December 4, 2007.

2 "'Anonymized' data really isn't — and here's why not," ArsTechnica, September 8, 2009.

Sweeney, who went on to serve as chief technology officer for the U.S. Federal Trade Commission and now works as the director of the Data Privacy Lab in the Institute of Quantitative Social Science at Harvard University, also showed in 2000 that 87 percent of all Americans could be uniquely identified by cross-referencing only three bits of information: their ZIP codes, birth dates and genders.

More recently, a 2015 research study led by Yves-Alexandre de Montjoye at MIT was able to reidentify 90 percent of shoppers from a group of 1.1 million by matching anonymized credit card transactions to publicly available information such as Twitter and Instagram posts. “The old model of anonymity doesn’t seem to be the right model when we are talking about large-scale metadata,” he told *The New York Times.*[3]

In each of the above examples, the results were arrived at by researchers, but they could just as easily have been perpetrated by malicious hackers. That may indeed be happening and we just don’t know about it, since criminals tend to not self-report the techniques they use. The fact that the above examples weren’t discovered by black-hat hackers owes more to luck than anything else, experts say. As with general breaches, catastrophic de-anonymization incidents becoming commonplace is all but inevitable.

3. “With a Few Bits of Data, Researchers Identify ‘Anonymous’ People,” *The New York Times*, January 29, 2015.

How Differential Privacy Works

A differentially private solution may use a technique such as injecting noise into a dataset (or into the output of a machine learning model) in order to protect individual privacy.

Differential privacy achieves this by calibrating the noise level to the sensitivity of the algorithm to noise. The end result is a differentially private dataset (or model) that cannot be reverse-engineered by an attacker.

Ian Goldberg, professor and research chair at the Cheriton School of Computer Science at the University of Waterloo, explains the process as being similar to randomized response, an anonymization technique that has been utilized in surveys for decades.

Goldberg uses the example of asking someone if they are an undocumented immigrant. Answering "yes" can be stigmatizing. To protect the survey respondents' privacy, a surveyor could use a randomized response technique. In that approach, the person answering the question flips a coin that the surveyor cannot see. If the result is "heads" then the person answers truthfully. If it's "tails" the person flips a second coin not visible to the surveyor. If the second coin toss is "heads," the person answers "yes" and if it's "tails," the answer is "no" regardless of the true response. This way, the individual answers do not reveal any personal information, while the survey still yields useful statistical conclusions.

Inserting randomized results into data doesn't completely prevent attempts to cross-reference that data with other sources, but as Goldberg points out, it does make it virtually impossible to identify specific individuals with full certainty.

Inserting randomized results into data makes it virtually impossible to identify specific individuals with full certainty.

Differential Privacy in Practice

To date the Georgian Impact team has worked with three portfolio companies to apply differential privacy to their solutions.

The first of these R&D projects was with Bluecore, the leading decisioning platform for commerce. Bluecore uses data to help companies identify their best customers and keep them for life. It does this by enabling companies to see what their customers are doing on their website in real time, such as what they're clicking on and searching for. With its Propensity to Convert functionality, Bluecore can also tell companies how likely a particular website visitor is to buy something based on that person's activity. Meanwhile, its Affinity Spaces functionality serves as a powerful recommendation engine.

For Bluecore, it's these add-on services — Propensity to Convert and Product Affinity — that are among its most popular and sought-after offerings. However, as with all B2B SaaS companies that provide data-driven insights, Bluecore has long faced what is known as the cold-start problem for new customers. Specifically, these offerings work best once sufficient data has been collected from a new customer to build an accurate predictive model for the respective customer. Collecting that data can take anywhere from several weeks to six months.

Bluecore knew that the only way to get around this cold-start issue was to build a predictive model that used all of its customers' data collectively. That way, new customers could instantly benefit from the model, even if they didn't yet have enough data of their own. However, given the high levels of sensitivity around data privacy, cross-customer data aggregation never seemed like a viable option. That's because most customers are concerned that using their data to build models for other customers may leak trade secrets or transactional information to their competitors.

Bluecore knew that the only way to get around this cold-start issue was to build a predictive model that used all of its customers' data collectively.

The project has shown that new and existing customers of Bluecore would both benefit from the use of differentially private data aggregation. New customers of Bluecore would be able to immediately start predicting from which website visitors are more likely to buy. Existing customers could benefit from predicted gains in sales of 10 percent — a significant uplift.

“We’ve shown a significant improvement for the cold-start problem, which is essentially when a partner first signs on with Bluecore and we don’t have that much data on that partner,” says Zahi Karam, Bluecore’s Director of Data Science. “By using a differential privacy-based approach, we can deliver insights for a new customer from day one rather than waiting a month or two to get them started.”

Another Georgian Partners portfolio company, Toronto-based Integrate.ai, is also working with the Georgian Impact team to implement differential privacy. The company’s AI-powered platform helps consumer-facing enterprises make customer interactions more natural and valuable.

“Companies don’t feel comfortable having their data directly intermingled with other companies’ data,” says Kathryn Hume, Integrate.ai’s Vice President of Product and Strategy. Taking a differentially private approach gives us the ability to join two interesting insights across different verticals together, which results in a whole that’s greater than the sum of its parts.”

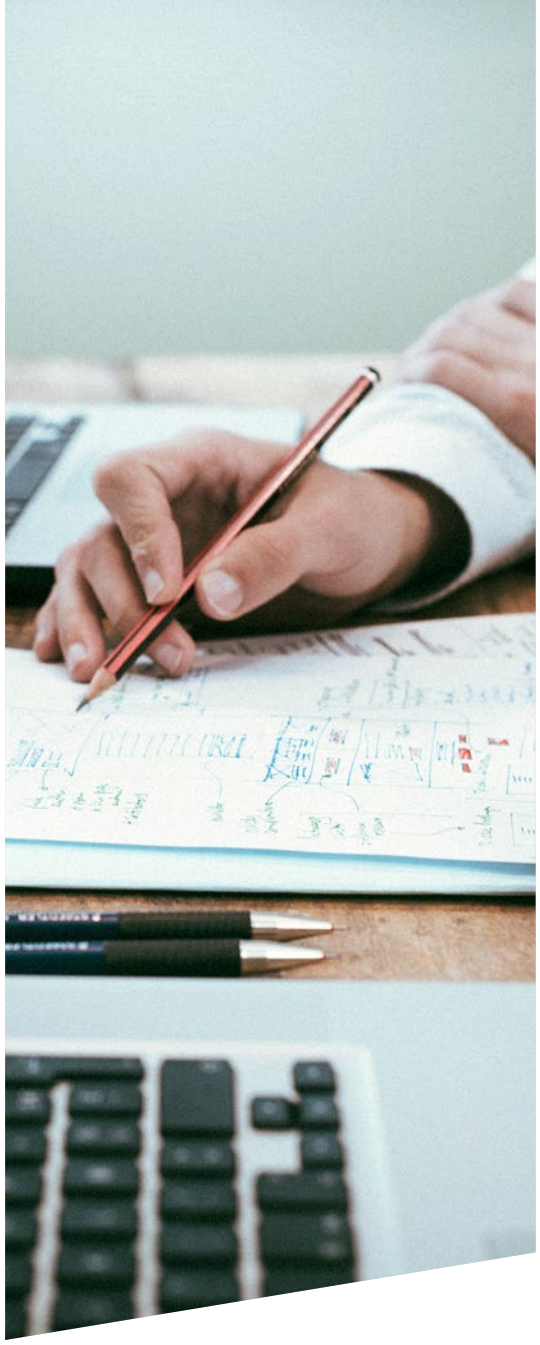

Getting Started

Before implementing differential privacy, first take a step back and consider the kind of information you're gathering.

Specifically, discuss with your team about whether or not you really need to collect the data you're collecting or are planning to collect. Even just a few years ago, the common approach was to amass as much customer information as possible with the goal of eventually finding a way to monetize it. Today, however, that approach needs revisiting given the continued and growing risk of data breaches and the emergence of new legislation such as the General Data Protection Regulation (GDPR) in Europe. If your company doesn't need to collect personal information, consider simply avoiding doing so.

"Data breaches are extremely common today and the only surefire way to avoid one is to not have the data," says the University of Waterloo's Goldberg. "It's now an explicit and foreseeable risk for a company to hold customer data."

Next, for the relevant data that your business needs, have your team look at what tools are available to support your adoption of differential privacy. Depending on your project focus and requirements, options may include:

- Georgian Partners Epsilon v1.0, our differential privacy product for Logistic Regression and Support Vector Machines (https://georgianpartners.com/products).
- Uber's open source tool that adds differential privacy for SQL queries (https://github.com/uber/sql-differential-privacy).
- Google's differentially private reporting tool RAPPOR (https://github.com/google/rappor).
- Immuta's differential privacy-enabled data science platform (https://www.immuta.com/).

Finally, differential privacy is still an area of emerging research. As such, continue to invest in talent, and hire strong data scientists and people with strong foundations in mathematics. Implementing differential privacy successfully in your company will be less about the lines of code needed and more about understanding and applying complex mathematics.

The Bottom Line

Companies should consider differentially private approaches to avoid data leakage from their deployed machine learning models. While differential privacy isn't a field to get into lightly, when properly applied, differential privacy can provide very useful results while safeguarding customer information.

By providing a way to mathematically guarantee the privacy of an individual's data, differential privacy can help an organization build trust with its customers. This trust is required to aggregate even more customer data in the future, resulting in better AI solutions and the ability to offer AI product capabilities to new customers from the start.

The effects of differential privacy are also likely to be cumulative over time. Early adopters will have a leg up on latecomers that replicate existing models and applications. Every business dealing with data needs to be investing in understanding and adopting differential privacy.

GeorgianPartners

info@georgianpartners.com // georgianpartners.com

About Georgian Partners

Georgian Partners is a thesis-driven growth equity firm that invests in SaaS-based business software companies. We look for companies that use foundational technology trends such as applied artificial intelligence, conversational business and security first to dominate their markets.

Founded by successful entrepreneurs and technology executives, at Georgian Partners we leverage our deep software expertise to directly impact the success of our portfolio companies. That expertise spans areas as diverse as machine learning, analytics, deep learning, cryptography, linguistics, natural language processing, differential privacy, software engineering, information security and cloud computing.