WHITE PAPER

Why Differential Privacy Should Be Top of Mind for Data Science and Governance Teams

MAR 2019

Sophie Stalla-Bourdillon
Senior Privacy Counsel & Legal Engineer, Immuta

IMMUTA

WWW.IMMUTA.COM

Note: The information contained in this presentation is not intended to be and should not be construed to be legal advice. Organizations should not rely on the information herein, and they should obtain legal advice from their own legal counsel or other professional legal services provider.

I recently taught a masterclass on de-identification organised by the Future of Privacy Forum. The audience comprised industry representatives, policy and law-makers mainly from the European Union. My goal was to cover the topic of differential privacy in an easy to digest manner, and to highlight the potential of this relatively new (at only 12 years old) privacy enhancing technology (PET) to decision-makers. It was fascinating to be able to challenge common assumptions about a variety of data processing activities. Microsoft Distinguished Scientist Cynthia Dwork wrote the first seminal contribution on the topic in 2006.

The view that data analytics methods always require access to individual-level data is widely spread, as much as the view that Big Data analytics is antithetical to privacy or data protection. When choosing to embrace PETs, though, a more nuanced – and exciting – world starts to take shape. And differential privacy should be playing an important role in this new world. Unsurprisingly, most of the participants at the event had not heard about differential privacy.

So why is it that data science teams and/or governance personnel should think about differential privacy more often?

It is a strong method to mitigate re-identification risks while deriving insights and utility from data. In fact, it is one of the strongest de-identification PETs, and therefore important to spread the word and make it easily accessible to data scientists.

How to explain the main features of differential privacy to a non-technical expert?

Differential privacy is based on the injection of randomised noise into the data analysis process. This is why compared to other PETs – which are not process based such as masking or generalisation techniques – with differential privacy it is possible to calibrate the noise to the query each time a query is made, and therefore to precisely navigate the trade-off between utility and privacy.

There are two common approaches when it comes to differential privacy:

1 **Global differential privacy** ensures that the individual whose data is being queried is in a position to **deny his/her participation to the data set**, meaning to make sure that she is able to deny that her data was included in the data set used to produce analysis results. The promise of global differential privacy is that the participation in the data set will not significantly increase the likelihood of re-identification.

This type of differential privacy enjoys a number of desirable properties. It is "immune against post-processing," as stated by Cynthia Dwork. What's more, this includes adversaries in possession of external or even future information! This is because differential privacy modifies the analysis process to ensure that the result to the queries depends weakly on each item (or data point) within the database. Further, its clever use of randomization promises that the final result is not only a possible result on a version of the database that does not include this item, it is almost as likely! Thus, observing this particular result does not tell one much about whether or not any data point is present.

Note that under this type of differential privacy you can only ask questions that will generate **aggregates** (e.g. minimum, maximum, average, count and sum). You could ask, for example, the following question: how many people have bought product B after having bought product A? Let's take the example of a customer database produced by Company Z in which John Smith appears as having bought products. The presence of John's record will have only a slight influence on the resulting count. This is because differential privacy ensures that whatever result is obtained by running the account occurs with nearly the same probability over a version of the database that does not include John's data. As a result, John will be able to argue, plausibly, that he has never bought any product from Z.

2 If aggregates are hard to work with for the kinds of analysis you wish to perform— say you wish to have access to individual-level data — you could use **local differential privacy**. Note, however, that aggregates can in principle be used in a variety of use cases (e.g., to analyse data in order to improve a manufacturing process, to analyse data in order to improve products/services, to create customer profiles to ensure maintenance of products/services, to derive insights in order to offer new goods or services, the list goes on...). The creation of profiles on the basis of aggregates is probably the least obvious use case and requires skill and expertise – but is feasible.

Unlike global differential privacy, with local differential privacy, an individual cannot deny her participation in the data set, but she can deny the contents of her record. The output of the process is therefore individual-noised records. This method has a great potential for supervised machine learning and is certainly under used!

The key to fully implementing these techniques is to understand that machine learning models should be built within controlled environments, which rely upon strict access control and allocation of roles between several lines of defence[1].

Within such an environment, it is indeed possible to proceed in different steps. First, build a version of the model without differential privacy. Do not release the model to the public at this stage. You would note its baseline performance and then throw away the model. You would then iteratively build models with more and more noise until you reach a minimum acceptable threshold for performance, or a maximum acceptable threshold for privacy loss. Assuming, then, that the privacy loss is acceptable, you could release the model into production.

[1] Beyond Explainability: A Practical Guide to Managing Risk in Machine Learning Models (go.immuta.com/beyond-explainability-white-paper)

COMPLIANCE PLAYBOOK

GDPR

Highlighting how to preform techniques critical for using your data in a GDPR compliant manner with Immuta's data management platform.

JULY 2018

IMMUTA

www.immuta.com

The Immuta GDPR Compliance Playbook for 2019 includes new best practices required for legal and compliant use of EU data for AI and Machine Learning, with a focus on Data Protection by Design. Learn purpose-based restrictions, how to map GDPR data protection principles to the Immuta platform's global policies, and guidance on implementing specific controls within the Immuta platform, such as masking and differential privacy. To download the playbook, visit: https://go.immuta.com/gdpr-compliant-ai-playbook.

Sophie Stalla-Bourdillon is a leading expert on the EU GDPR, and Senior Privacy Counsel | Legal Engineer at www.immuta.com. She is responsible for examining current data protection and model risk frameworks, helping customers to embed aspects of these frameworks within the Immuta platform, and framing these practices into digestible, easy-to-scale methods so they can better control risk across their data science programs.

IMMUTA

7878 Diamondback Dr, Suite C, College Park, MD 20740 | immuta.com | (800) 655-0982