PRIVITAR

Differential privacy

A new level of protection against
a new type of threat

www.privitar.com

Privitar works with the Government Statistical Service (GSS) to highlight new capabilities for managing and mitigating privacy risk

Any release of aggregate statistics (counts, sums and averages, for example) about sensitive datasets carries some element of risk – it's your ability to quantify and manage that risk that's really important.

Determined attackers can use a variety of methods to unravel statistical disclosures and reveal information about individuals. And in a world of increasing data volumes – and increasingly powerful privacy attacks – traditional methods of preventing unintended disclosure about individuals are struggling to keep up.

A new report from the UK Office for National Statistics, who delivered the report on the behalf of the Government Statistical Service (GSS), looks at the changing privacy threat landscape and state-of-the-art methods of managing privacy risk. The GSS asked Privitar experts and Professor Kobbi Nissim, one of the inventors of differential privacy, to co-author a chapter of the report to show how differential privacy can help preserve individuals' privacy in the face of a new kind of threat.

Reconstruction attacks: the new privacy threat

It's the discovery of a new, serious type of attack called a reconstruction attack that's accelerated the need for a new approach to privacy.

In a reconstruction attack, the attacker uses aggregate statistics released about a sensitive dataset (even when protected by traditional privacy methods) to infer with high accuracy the dataset itself.

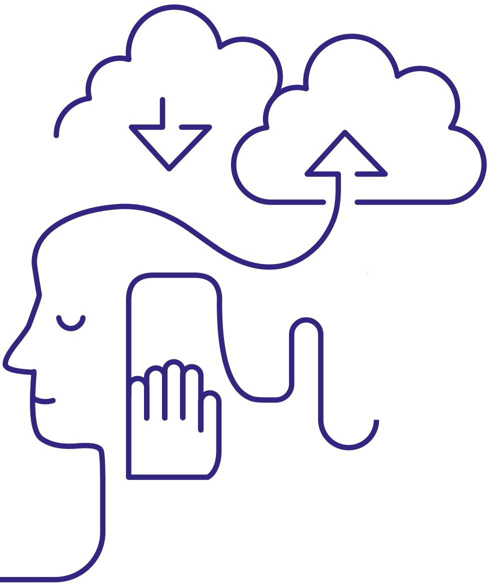

Once considered only a theoretical risk, real-world reconstruction attacks have now been demonstrated in practice by the US Census and independently by a New York Times journalist[1]. The US Census Bureau recently reported “serious vulnerabilities” to reconstruction attacks in its 2000 and 2010 Census data releases, leading it to use differential privacy for the 2020 Census[2].

John Abowd, the US Census Bureau’s chief scientist and associate director of research and methodology, described reconstruction attacks as “the death knell for traditional data publications systems”[3]. And a recent paper co-authored by Abowd concludes:

“The vast quantity of data products published by statistical agencies each year may give a determined attacker more than enough information to reconstruct some or all of a target database and breach the privacy of millions of people. Traditional disclosure avoidance techniques are not designed to protect against this kind of attack.”

Differential privacy: key advantages

> Quantify the level of privacy and utility for each use case

> Gain insights from data which would otherwise be too sensitive to be used

> Defend against even the most sophisticated privacy attacks, including reconstruction attacks

> Future-proof protection, with privacy-preserving methods that make no assumptions about attack strategies

> Quantify privacy risk across multiple statistical releases

> Disclose algorithms and parameters without risk, for complete transparency

> “The database reconstruction theorem is the death knell for traditional data publication systems from confidential sources[4].”
>
> John M. Abowd, Chief Scientist and Associate Director for Research and Methodology, US Census Bureau

[1] https://www.nytimes.com/2018/12/05/upshot/to-reduce-privacy-risks-the-census-plans-to-report-less-accurate-data.html

[2] https://doi.org/10.1145/3219819.3226070

[3] https://digitalcommons.ilr.cornell.edu/cgi/viewcontent.cgi?article=1051&context=ldi

[4] https://www.census.gov/content/dam/Census/newsroom/press-kits/2018/jsm/jsm-presentation-database-reconstruction.pdf

What is differential privacy?

Differential privacy is a new approach to managing privacy risk in today's data-rich world.

Since its invention in 2006, differential privacy has grown beyond its roots in academia to see significant adoption in government agencies and big tech players like Apple and Google.

By distorting data in a precisely calibrated way, differential privacy methods allow organizations to gain insights into groups of individuals while guaranteeing that nothing significant can be learned about any given individual within the group.

Because the degree of distortion relates directly to the trade-off between data utility and individual privacy, differential privacy enables organizations to make informed decisions about publishing statistics from sensitive datasets.

Differential privacy provides a mathematical guarantee to individuals that their privacy risk (how much information specific to them is revealed) is limited. By limiting the information revealed about individuals, differential privacy can defend against attacks such as reconstruction.

Getting started with differential privacy

Although differential privacy has a long background in academia, it's a relatively new approach for large-scale applications.

Like any emerging discipline, there are some practical challenges to overcome to make it work in the real world.

Implementing differentially private algorithms and configuring them to achieve the desired balance of privacy and utility will require qualified staff. And data holders, analysts and subjects need to be educated about the value (and the limitations) of differential privacy.

Our recommendations to the GSS to help it protect privacy when releasing aggregate statistics were:

- Gauge the risk of reconstruction attacks, and other state-of-the-art privacy attacks, in existing statistical releases.
- Identify the right use cases for piloting differential privacy. Lower-sensitivity datasets can be a good place to start while gaining experience with differential privacy.
- Engage with policymakers, legal scholars, differential privacy researchers, and other relevant stakeholders to discuss appropriate levels of privacy protection.
- Strengthen relationships with differential privacy communities in academia and industry, to influence research in directions that are relevant for the organization, and help bridge the gap between theory and practice.

Learn more

Differential privacy can help you provide more robust and transparent protection, address existing vulnerabilities to new threats, and prepare with confidence for the privacy challenges of the future.

For a deep dive into differential privacy – what it is, why it's so important, and how to get it right – download the full GSS report, accessible below.

Download now

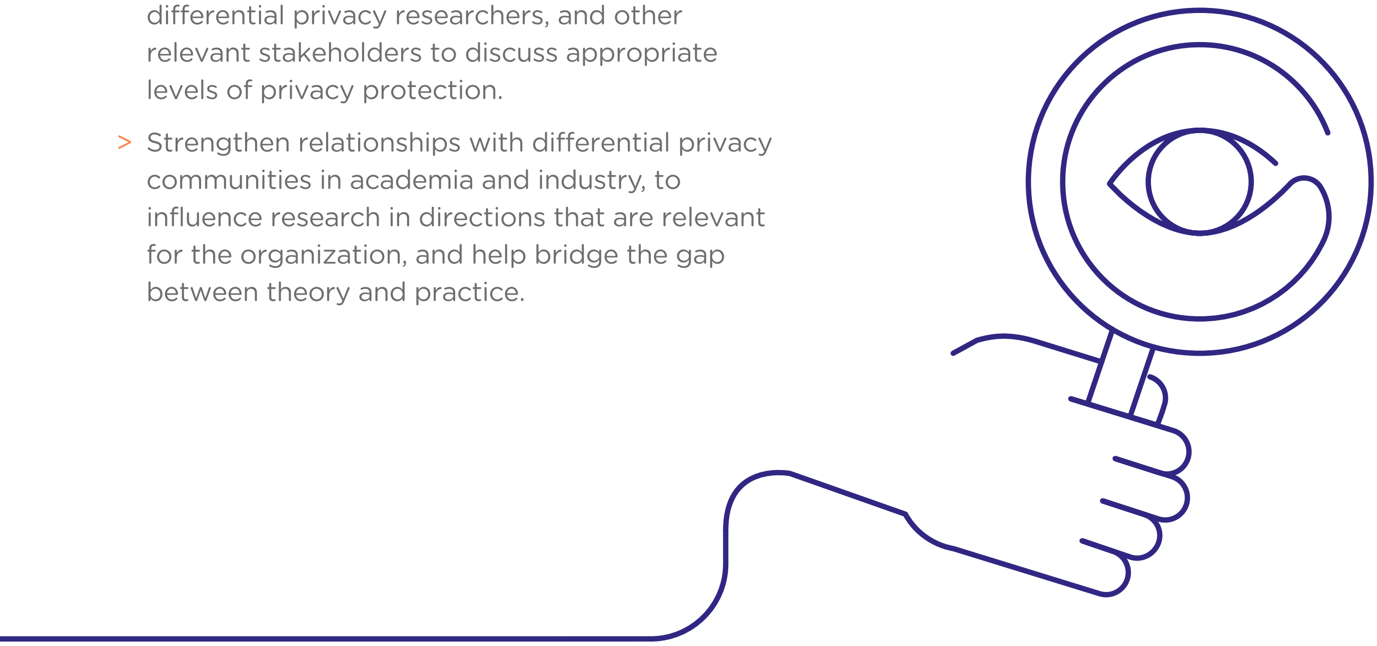

We’re Privitar

We help organizations engineer privacy-preserving data operations, using advanced data privacy techniques that protect sensitive information while retaining data utility. Our software accelerates and automates privacy-safe data provisioning, helping our customers get more business value from their data, generate data-driven insights, and drive innovation.

Contact us:

e: info@privitar.com
t: +44 203 282 7136
w: www.privitar.com

PRIVITAR

@PrivitarGlobal

www.privitar.com