Position: Challenges and Opportunities for Differential Privacy in the U.S. Federal Government

Amol Khanna
Booz Allen Hamilton
Boston, MA
Khanna_Amol@bah.com

Adam McCormick
Booz Allen Hamilton
Washington, DC
McCormick_Adam@bah.com

Andre Nguyen
Booz Allen Hamilton
Washington, DC
Nguyen_Andre@bah.com

Chris Aguirre
Booz Allen Hamilton
Austin, TX
Aguirre_Chris@bah.com

Edward Raff
Booz Allen Hamilton
Syracuse, NY
Raff_Edward@bah.com

Abstract

In this article, we seek to elucidate challenges and opportunities for differential privacy within the federal government setting, as seen by a team of differential privacy researchers, privacy lawyers, and data scientists working closely with the U.S. government. After introducing differential privacy, we highlight three significant challenges which currently restrict the use of differential privacy in the U.S. government. We then provide two examples where differential privacy can enhance the capabilities of government agencies. The first example highlights how the quantitative nature of differential privacy allows policy security officers to release multiple versions of analyses with different levels of privacy. The second example, which we believe is a novel realization, indicates that differential privacy can be used to improve staffing efficiency in classified applications. We hope that this article can serve as a nontechnical resource which can help frame future action from the differential privacy community, privacy regulators, security officers, and lawmakers.

1 Differential Privacy

Despite executive orders and guidance encouraging the use of revolutionary privacy enhancing technologies to reduce the risk related to the rise of big data, artificial intelligence (AI), and machine learning (ML), institutional barriers impede the widespread deployment of differential privacy throughout the U.S. federal government. Released over the past year, "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," and NIST SP 800-226 "Evaluating Differential Privacy Guarantees" both explore the need for adopting privacy enhancing technology to mitigate privacy risks arising from increased data processing [1, 2].

Differential privacy, a privacy-preserving algorithmic framework, is one such privacy preserving technology that can be applied to a wide variety of common government agency use cases and has already been studied, documented, and deployed in commercial industry. Differential privacy measures the privacy risk to individuals by comparing the output of an algorithm with an individual's data to the output without that data. If the output can be heavily swayed by a single datapoint, the algorithm can betray the privacy of specific individuals. If this cannot happen, the output does not indicate if a datapoint was used as an input [3]. Differential privacy employs quantitative parameters to set the amount that an output can change with each additional datapoint, and by modifying these parameters, practitioners can set the allowable privacy risk of an algorithm. Differential privacy

2nd Workshop on Regulatable ML at NeurIPS 2024.

is the only existing privacy enhancing technology which can make a statistical guarantee of data privacy, indicating that released statistics, synthetic datasets, and even machine learning models cannot be reverse engineered to identify individuals in their source datasets. Unlike other methods, the privacy guarantee is not conditional and cannot be attacked, which makes it an attractive option for guaranteeing privacy in privacy-sensitivity applications [4].

2 Attempts and Challenges of Federal Differential Privacy Implementations

Within the federal government, differential privacy has the potential to reduce potential harms caused by data leaks and enable more public releases of data and statistics, which can lead to better informed political and economic decisions [5]. Indeed, the U.S. Census Bureau has already taken steps in this direction by using differential privacy to mask some of the sensitive statistics it released from the 2020 Census [6]. By using differential privacy, the Bureau guaranteed that specific households cannot be identified from the privatized statistics while still releasing information to lawmakers [7].

Despite the U.S. government touting the benefits of differential privacy and encouraging its adoption for a variety of use cases, the use of the technology has largely been restricted, to wit there are no other prominent examples. We argue that the limited government usage of the technology can be attributed to a lack of awareness among government program managers, challenging deployments in large-scale systems, and most critically, unclear guidance to security officers challenging deployments in large-scale systems. We discuss each of these challenges in turn:

1. Lack of Awareness: differential privacy is a statistical method for guaranteeing privacy, which departs from older privacy-preservation techniques that were deterministic, but ineffective. Older methods, like record anonymization, removing sensitive attributes, and K-anonymity, operated directly on datasets [8]. These methods could typically be used as a dataset preprocessing step prior to using a standard algorithm for statistics or machine learning. However, these methods are still very vulnerable to privacy attacks; indeed, government datasets that employed the aforementioned methods have been attacked in the past [9]. Differential privacy is robust to all attacks, but effectively using it requires a deep understanding of probability and statistics, which many privacy professionals are unfamiliar with. Additionally, differential privacy's statistical nature often requires modifying an entire algorithmic pipeline; instead of simply preprocessing the rows of a dataset, significant modifications must be made to a target algorithm to produce useful results [10]. This statistical nature can also impact decision-makers, as the notion that data is secure but probabilistically protected is a departure from better-understood methods like cryptography. Finally, current implementations of differential privacy use functional programming APIs, which can be less accessible to data scientists than more common object-oriented frameworks [11].
2. Challenging Deployments: differential privacy can be challenging to use for tasks with multiple goals and intermediate computations [12]. In such tasks, there are many potential statistical missteps that can produce sources of privacy leaks. Large technology companies have approached this challenge by focusing on deploying differential privacy for targeted tasks and including differential privacy experts to oversee deployments. However, government agencies are often tasked with releasing social and political information to the public, and these datasets often require multiple sources and released values.
3. Unclear Guidance: most importantly, many overseeing security officers in the government are reluctant to use differential privacy since no government authority has released official approvals for using the technology in privacy-critical tasks. Since bureaucratic approvals exist for older technologies, security officers choose to use these technologies despite their increased susceptibility to privacy attacks such as membership inference and model inversion attacks [13, 14]. We believe that for differential privacy to permeate through government applications, a government authority must approve its use in multiple privacy-critical applications, since this can set a government standard.

Each of these challenges was realized in the Census Bureau's differential privacy deployment. First, approvals for differential privacy use did not exist, and the Census used privacy-enhancing dataset preprocessing techniques in conjunction with differential privacy [15]. Next, a number of released values had to be modified since the Census contains many codependent statistics, and adding

noise to many individual statistics can significant diminsh the utility of the final result. Finally, many statisticians and social scientists critiqued the Bureau's use of differential privacy for having multiple potential privacy leaks and discrepancies among statistics used to inform political and economic decisions [16, 17, 18, 19]. This was exacerbated by the fact that differentially private noise disproportionately affects underrepresented communities in datasets, and social scientists analyzing Census data often seek to identify and help marginalized communities. Ultimately, the Census Bureau has made a massive step forward in the practice of real-world differential privacy, and securing citizens' information, which are hard-fought lessons we wish to see replicated in other agencies.

3 Differentiating Private Releases by Perceived Trustworthiness

As a state-of-the-art privacy enhancing technology, we believe that differential privacy should be incorporated into more government agency applications. In addition, we argue that differential privacy's parametric approach enables data security officers to further improve customizability and efficiency in government applications. Specifically, security officers can inject their knowledge of a dataset's sensitivity and an algorithm's end users through differential privacy's quantitative parameters [20]. This customization means that different versions of the same algorithm can be released with varying levels of privacy based on a perceived level of risk.

For example, if an agency working with personal health information built a machine learning model on patient records, the agency could enable physicians to access a less private and more accurate model while releasing a more private version to academics and insurance companies. This is because the agency may see registered physicians as a trusted group, while researchers and insurance companies are unknown public entities. By releasing the two different versions of the model, the agency enables physicians to access a state-of-the-art model for disease diagnosis and treatment while spurring advancements in research and policy with a more private model. Using other privacy enhancing methods without this fine-grained control may have resulted in a less accurate model for physicians or a health data breach, either of which are unsatisfactory outcomes.

4 Improving Efficiency with Privacy: An Unrealized Frontier

The previous section's example highlights differential privacy's applicability in civilian agencies. We also argue that differential privacy could enable significant efficiency gains in classified settings, and we believe that we are the first to identify this opportunity. Government defense agencies often generate statistics and models on datasets classified, e.g., as Top Secret. Under current guidance, every person with access to any part of these modeling pipelines, such as data annotators, data scientists, machine learning engineers, software engineers, deployment specialists, and even eventual users, must have a Top Secret clearance. This is because all of these individuals can either access raw Top Secret data or derivatives of this data, which could be reverse-engineered into the Top Secret data. However, staffing technical roles at such a high clearance level is challenging due to a shortage of cleared staff, and as such, agencies often choose not to pursue projects in highly classified settings due to a lack of cleared and capable personnel [21, 22, 23]. We believe that if differential privacy was used when creating data derivatives, certain positions in these pipelines could be staffed at lower clearance levels. For example, if a sufficient level of privacy was used when generating statistical reports on Top Secret datasets, software engineers and deployment specialists would be unable reverse engineer these reports to identify specific Top Secret datapoints. Thus, there is a reduced risk associated with these reports, meaning that personnel handling, deploying, and using these reports can have lower clearance levels. Staffing at lower clearance levels is significantly easier and cheaper, so adopting this policy on highly classified projects can significantly reduce time and monetary costs and allow agencies to pursue more initiatives.

5 Conclusion

This work seeks to communicate the challenges our team of differential privacy researchers, privacy lawyers, and data scientists have had when communicating about differential privacy to decision-makers at government agencies. We have noticed that senior data scientists are wary of using the technology due to its significant departure from previous methods and its difficulty in scaling to large applications and unstructured datasets. Security officers often avoid novel privacy preserving

technologies due to a lack of government guidance and specifications on deployments. We note that each of these challenges has been discussed in prior literature on differential privacy, but is particularly relevant to the highly structured and regulated setting of data processing at government agencies [24, 25].

We then proceed to highlight two examples of government tasks that could be improved with differential privacy. The first demonstrates that in civilian agencies, the government can use the quantitative nature of differential privacy to release different versions of the same output to different populations based on a perceived level of privacy risk. This can enable government agencies to serve each stakeholder of their released statistics and models effectively while managing the privacy risk of sensitive data. Our next use case highlights that in classified settings, using differential privacy can allow the government to staff positions handling data derivatives at lower clearance levels, which can significantly improve efficiency. We note that we have not seen this use case discussed in any other literature; this is likely because classification and clearances are only strongly enforced in government settings. As such, we hope that by communicating this use case clearly, government security officers can consider lowering classifications of data derivatives when using differential privacy.

Finally, we conclude by highlighting four areas of future action critical to the success of deploying differential privacy in government applications.

1. The differential privacy community must develop an effective method for communicating privacy guarantees to security officers and decision makers. The community is already aware of this critical need and recent research works on this problem [10, 26].
2. Differential privacy researchers should develop better methods for privatizing large, unstructured datasets. These forms of datasets are underexplored in the differential privacy literature but are very common in government applications.
3. There is a significant gap between theoretical works on differential privacy and practical differential privacy deployments. We identify works authored in-part by government employees focused on empirical methods for differential privacy, and believe that government agencies would be best served by empirical works and implementations of theoretical methods as these provide evidence that differential privacy is ready for deployment [27, 28, 29, 30, 31, 32, 33].
4. Regulators, lawmakers, and security officers must decide on a framework for evaluating, accepting, and deploying differentially private systems on government datasets.

We believe that effectively addressing these three challenges will unlock significant use cases for differential privacy in the government setting.

References

[1] Joseph R Biden. Executive order on the safe, secure, and trustworthy development and use of artificial intelligence. 2023.

[2] Joseph P Near, David Darais, Naomi Lefkovitz, Gary Howarth, et al. Guidelines for evaluating differential privacy guarantees. *National Institute of Standards and Technology, Tech. Rep*, 2023.

[3] Joseph P Near and Chiké Abuah. Programming differential privacy. 2021.

[4] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.

[5] Rashmi Krishnamurthy and Yukika Awazu. Liberating data for public value: The case of data.gov. *International Journal of Information Management*, 36(4):668–672, 2016.

[6] John M Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2867–2867, 2018.

[7] Buxin Su, Weijie J. Su, and Chendi Wang. The 2020 united states decennial census is more private than you (might) think, 2024.

[8] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International journal of uncertainty, fuzziness and knowledge-based systems*, 10(05):557–570, 2002.

[9] Daniel Barth-Jones. The're-identification'of governor william weld's medical information: a critical re-examination of health data identification risks and privacy protections, then and now. *Then and Now (July 2012)*, 2012.

[10] Ivoline C Ngong, Brad Stenger, Joseph P Near, and Yuanyuan Feng. Evaluating the usability of differential privacy tools with data practitioners. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, pages 21–40, 2024.

[11] Shiliang Zhang, Anton Hagermalm, Sanjin Slavnic, Elad Michael Schiller, and Magnus Almgren. Evaluation of open-source tools for differential privacy. *Sensors*, 23(14):6509, 2023.

[12] Ying Zhao and Jinjun Chen. A survey on differential privacy for unstructured data content. *ACM Computing Surveys (CSUR)*, 54(10s):1–28, 2022.

[13] OTCnet Deployment Team. Otcnet pii information, 2020.

[14] Staff in the Office of Technology. No, hashing still doesn't make your data anonymous, 2024.

[15] Simson Garfinkel. Differential privacy and the 2020 us census. 2022.

[16] J Tom Mueller and Alexis R Santos-Lozada. The 2020 us census differential privacy method introduces disproportionate discrepancies for rural and non-white populations. *Population Research and Policy Review*, 41(4):1417–1430, 2022.

[17] Danah Boyd and Jayshree Sarathy. Differential perspectives: Epistemic disconnects surrounding the us census bureau's use of differential privacy. *Harvard Data Science Review (Forthcoming)*, 2022.

[18] Michael B Hawes. Implementing differential privacy: Seven lessons from the 2020 united states census. *Harvard Data Science Review*, 2(2):4, 2020.

[19] Aloni Cohen, Moon Duchin, J Matthews, and Bhushan Suwal. Private numbers in public policy: Census, differential privacy, and redistricting. *Harvard Data Science Review*, (Special Issue 2), 2022.

[20] Priyanka Nanayakkara, Mary Anne Smart, Rachel Cummings, Gabriel Kaptchuk, and Elissa M Redmiles. What are the chances? explaining the epsilon parameter in differential privacy. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1613–1630, 2023.

[21] Ginger Groeber, Paul W Mayberry, Brandon Crosby, Mark Doboga, Samantha E Dinicola, Caitlin Lee, and Ellen E Tunstall. Federal civilian workforc e hiring, recruitment, and related compensation practices for the twenty-first century. 2020.

[22] Benjamin F Berger. *US security clearances: Reducing the security clearance backlog while preserving information security*. PhD thesis, Monterey, CA; Naval Postgraduate School, 2019.

[23] Brenda S Farrell. Personnel security clearances: Additional actions needed to ensure quality, address timeliness, and reduce investigation backlog. 2017.

[24] Daniel Franzen, Saskia Nuñez von Voigt, Peter Sörries, Florian Tschorsch, and Claudia Müller-Birn. " am i private and if so, how many?"–using risk communication formats for making differential privacy understandable. *arXiv preprint arXiv:2204.04061*, 2022.

[25] Simson L Garfinkel, John M Abowd, and Sarah Powazek. Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, pages 133–137, 2018.

[26] Rachel Cummings and Jayshree Sarathy. Centering policy and practice: Research gaps around usable differential privacy. In *2023 5th IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pages 122–135. IEEE, 2023.

[27] Aniruddha Sen, Christine Task, Dhruv Kapur, Gary Howarth, and Karan Bhagat. Diverse community data for benchmarking data privacy algorithms. *Advances in Neural Information Processing Systems*, 36, 2024.

[28] Amol Khanna, Fred Lu, Edward Raff, and Brian Testa. Differentially private logistic regression with sparse solutions. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, pages 1–9, 2023.

[29] Amol Khanna, Edward Raff, and Nathan Inkawhich. Sok: A review of differentially private linear models for high-dimensional data. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 57–77. IEEE, 2024.

[30] Tyler LeBlond, Joseph Munoz, Fred Lu, Maya Fuchs, Elliot Zaresky-Williams, Edward Raff, and Brian Testa. Probing the transition to dataset-level privacy in ml models using an output-specific and data-resolved privacy profile. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, pages 23–33, 2023.

[31] Fred Lu, Joseph Munoz, Maya Fuchs, Tyler LeBlond, Elliott Zaresky-Williams, Edward Raff, Francis Ferraro, and Brian Testa. A general framework for auditing differentially private machine learning. *Advances in Neural Information Processing Systems*, 35:4165–4176, 2022.

[32] Ryan Swope, Amol Khanna, Philip Doldo, Saptarshi Roy, and Edward Raff. Feature Selection from Differentially Private Correlations. In *Proceedings of the 17th ACM Workshop on Artificial Intelligence and Security (AISec'24)*, 2024. URL `https://arxiv.org/abs/2408.10862`.

[33] Amol Khanna, Fred Lu, and Edward Raff. The challenge of differentially private screening rules. *2nd AdvML Frontiers Workshop at 40th International Conference on Machine Learning*, 2023. URL `https://arxiv.org/abs/2303.10303`.