Google DP vs. OpenDP: Empirical Comparison of Differential Privacy Libraries

Author: Stilyan Penchev <spenchev@tudelft.nl>
Supervisors: Dr. Zeki Erkin, Dr. Roland Kromes
Affiliation: EEMCS, Delft University of Technology, The Netherlands

Introduction

Overview: The modern world relies on analyzing large, sensitive datasets, but this creates significant privacy risks. Simply removing names is insufficient, as re-identification through linkage attacks is a known threat [2]. Differential Privacy (DP)[1] has become a key method to guarantee privacy during data analysis.

Differential Privacy is a framework designed to enable useful data analysis while providing strong, mathematical privacy guarantees [1]. Its core principle is that the output of any analysis should not significantly change if a single individual's data is added to or removed from the dataset. This protection is achieved by adding carefully calibrated random noise to the true result of a query. The amount of noise is controlled by a privacy budget, epsilon (ϵ), which creates a fundamental trade-off: stronger privacy (lower ϵ) requires more noise, reducing the accuracy of the result.

Problem & Motivation: To put DP into practice, software libraries like Google's Differential Privacy Library and the OpenDP Library are essential. Although they aim to provide similar privacy guarantees, their internal implementations, performance, and utility can differ significantly. Because these libraries are actively developed, an up-to-date empirical comparison is needed to help practitioners choose the right tool. This study builds on prior work, such as Zhang et al [4]., by providing a focused benchmark on the latest library versions.

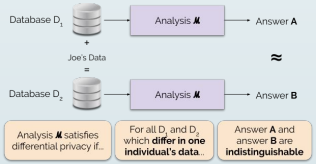

Figure 1. Differential Privacy Diagram (Source: Wikimedia Commons)

Research Question

How do the Google Differential Privacy Library and the OpenDP Library compare in terms of computational performance, scalability, and the utility of their implemented Laplace and Gaussian mechanisms when applied to Count and Sum queries on structured datasets?

Methodology

Experimental Setup:
- Hardware: HP ZBook Power G10, Intel Core i7-13700H, 16GB RAM, Windows 11.
- Software: Python 3.9.18, GoogleDP (PyDP) 1.1.4, OpenDP 0.13.0, pandas, numpy.

Datasets:
- Type: Synthetic datasets with a single numerical column (values uniformly drawn from [-100.0, 100.0]).
- Sizes: Small (10k records), Medium (100k records), Large (1M records).

Queries & Mechanisms:
- Queries Tested: Differentially Private Count and Bounded Sum.
- Mechanism Analyzed: Laplace mechanism. (Note: Gaussian mechanism evaluation was deferred) [3].

Privacy Parameters:
- Epsilon (ϵ): Tested values from 0.1 (stronger privacy) to 3.0 (weaker privacy).

Evaluation Metrics:
- Performance: Average execution time (ms) over 100 runs.
- Scalability: How performance and utility change with dataset size.
- Utility: Mean Absolute Error (MAE) between noisy results and the true value.

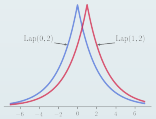

Figure 2. Laplace Mechanism

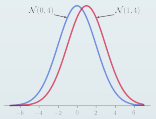

Figure 3. Gaussian Mechanism

References

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography, 2006.
[2] L. Sweeney. k-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 2002.
[3] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 2014.
[4] R. Zhang, B. Niu, et al. Evaluating OpenDP SmartNoise and Google DP with Other Libraries for Differential Privacy. Sensors, 2023.

Results

Performance (vs. Epsilon):
- Execution time was largely insensitive to changes in ϵ for both libraries.
- GoogleDP was consistently faster than OpenDP across all dataset sizes, with the performance gap widening significantly on the Large dataset.

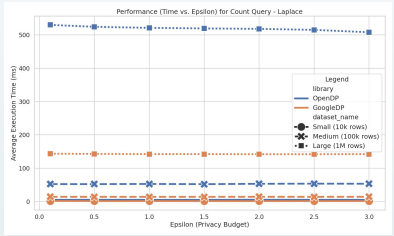

Figure 4. Performance (Time vs. Epsilon) for Count Query

Scalability of Performance:
- Both libraries showed expected increases in execution time with larger datasets, with performance scaling approximately linearly on a log-log scale.
- GoogleDP consistently maintains a significant performance advantage over OpenDP.

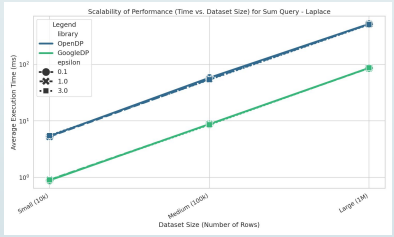

Figure 5. Scalability of Performance for Sum Query

Utility (MAE vs. Epsilon):
- As expected, MAE decreased as ϵ increased for both libraries, showing the privacy-utility trade-off.
- For Count and Sum queries, OpenDP often provided similar or slightly lower MAE (better utility).

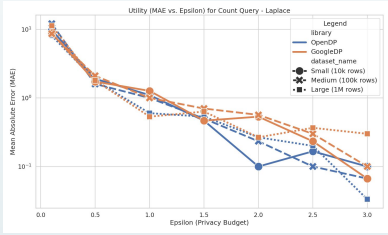

Figure 6. Utility for Count Query

Conclusions and Future Work

Conclusions:
- For Laplace-based queries, **GoogleDP offers a clear advantage in computational speed and scalability**.
- OpenDP provides a competitive and, in some cases, a marginally better utility profile.
- The choice between libraries depends on whether the primary criterion is processing efficiency or achieving optimal utility for specific privacy configurations.

Limitations & Future Work:
- **Limitation:** This study was confined to the Laplace mechanism. A comparative analysis of the Gaussian mechanism was not completed due to implementation status in the tested GoogleDP version.
- **Future Work:** The essential next step is an **empirical comparison of Gaussian mechanism implementations** in both libraries. Further work could also include more query types, real-world datasets, and assessing new library releases.