Cross-Domain Comparison of Differential Privacy Tools: Google DP vs. DP-OPT

Yurui Zheng (Y.Zheng-32@student.tudelft.nl) || Responsible Professor: Dr. Zeki Erkin || Supervisor: Dr. Roland Kromes

1 Background & Motivation

Differential Privacy (DP)
DP is a mathematical framework that ensures the output of a computation does not change significantly when any single individual's data is modified. Formally, a randomized algorithm M is said to be **(ε, δ)-differentially private** if, for all adjacent datasets D and D′ (differing by a single individual), and for all possible outputs S, the following equation holds:

$$\Pr[M(D) \in S] \leq exp(\varepsilon) \Pr[M(D') \in S] + \delta$$

However, DP is just a theoretical concept that needs to be applied in practice. This study compares two such implementations:

Google Differential Privacy Library
- Used for statistical analysis on tabular data
- Adds calibrated noise to queries (SUM, COUNT, AVG, etc) using the Laplace or Gaussian mechanisms
- Tracks privacy loss using standard **(ε, δ) accounting**
- Limits the number of contributions a single user can make to the data

Differentially Private Offsite Prompt Tuning (DP-OPT)
- Used to adapt large language models (LLMs) without exposing sensitive data through prompt tuning
- Generates private prompts locally
- Tracks privacy using **Rényi Differential Privacy (RDP)**
- Designed for scenarios where model weights are not changed or the model is closed-source

Nowadays, we have seen an increase in the number of applications incorporating both **statistical analysis and machine learning models** in their tasks; therefore, the need for a **cross-domain comparison** has become even more relevant for this purpose. This study aims to bridge that gap by evaluating the **utility, performance, and privacy accounting mechanisms** trade-offs across these tools.

Research Questions

How do DP-OPT and Google's DP Library compare when accounting for different factors in different contexts?:

- How do DP-OPT and Google's DP Library compare in their privacy-budget accounting mechanisms?
- What are the performance trade-offs (runtime, memory) of each tool on representative ML and analytics tasks?
- How does the output utility of DP-OPT compare to that of Google's DP Library across different use cases?

2 Methodology

Literature-Based Review

Start with foundation papers on Google DP and DP-OPT
- Use snowballing to identify additional studies
- Focus on peer-reviewed benchmarks and technical reports

No Empirical Experiments

Reason:
- Tool Complexity (esp. DP-OPT)
- Time/Resource Constraints
- Availability of high-quality existing evaluations

Evaluation Framework

Metric	Google DP	DP-OPT
Utility	RMSPE on statistical queries	Accuracy drop on classification tasks
Performance Overhead	Runtime/memory cost of queries	Prompt generation time and memory usage
Privacy Accounting	(ε, δ) composition	Rényi DP tracking

Synthesis & Comparison

Summarize trade-offs, strengths, and use-case suitability to provide practical guidance for tool selection

3 Results

Metric	Google DP	DP-OPT
Utility	Low error on statistical queries	High accuracy on ML tasks ($\Delta < 2\%$)
Runtime	~10–30% overhead	Fast prompt gen. (~1h), low compute cost
Memory Use	≤ 3% increase	Relatively low (no gradients/backprop)
Privacy Model	(ε, δ) → transparent & auditable	RDP → tighter but less interpretable
Best Use Case	Structured analytics with tabular data	LLM adaptation under strict privacy

4 Discussion

Context is important:
- Google DP is best suited for **structured data analysis** with low computational cost and high interpretability
- DP-OPT is more effective for **machine learning tasks**

Utility trade-offs depends on the complexity class of the task:
- Google DP improves with larger datasets and well-tuned parameters
- DP-OPT maintains high accuracy even under strict privacy budgets by using more powerful ML models

Performance changes with the system's design:
- Google DP adds minimal overhead in both memory and runtime
- DP-OPT requires more computational resources, but remains memory efficient since it avoids gradient-based training

Privacy accounting differs in interpretability:
- Google DP uses standard (ε, δ) accounting, which is transparent and auditable
- DP-OPT uses RDP, which offer tighter privacy guarantees but require more expertise to interpret

No one-size-fits-all solution:
- The right DP tool depends on **where in the pipeline** privacy must be enforced
- Early (data aggregation) → Google DP
- Late (model tuning) → DP-OPT
- In some cases, both might be the solution

5 Conclusion

- Google DP performs well for structured data analysis with low overhead and strong interpretability
- DP-OPT performs well for LLM adaptation tasks with good utility under strict privacy constraints
- **Tool choice depends on the task and pipeline stage,** choosing one over the other presents trade-offs in different areas
- **Limitations:**
  - Findings are based on existing literature and bechmarks
  - Results may not generalize to all data types or system configurations
- **Future work:**
  - Conduct empirical experiments to valide findings across different types of data and settings
  - Evaluate the practical usability of both tools
  - Test scenarios where Google DP and DP-OPT are implemented in the same pipeline