A Visualization Tool to Help Technical Practitioners of Differential Privacy[*]

Liudas Panavas[†]
Northeastern University

Saeyoung Rho[‡]
Columbia University

Hari Bhimaraju[†]
Columbia University

Wynee Pintado[§]
Barnard College

Rebecca N. Wright[‡]
Barnard College

Rachel Cummings[†]
Columbia University

Abstract

Addressing the complex challenges of implementing differential privacy (DP) requires effective tools that bridge the gap between theoretical concepts and practical application. To this end, we are developing a visualization tool designed to facilitate understanding of DP's privacy/accuracy tradeoffs. We aim to connect quantitative formulas and mathematical accuracy definitions with qualitative, visual insights, making the complexities of privacy/accuracy tradeoffs more tangible. To inform our visualization tool, we are conducting formative interviews with differential privacy practitioners to identify the challenges they encounter in understanding and applying these concepts. Our tool will use interactive visualizations to enhance users' comprehension by translating abstract concepts and theoretical accuracy bounds into digestible visualizations. Ultimately, the tool aims to be a robust, user-friendly resource that enables technical users to independently learn about and explore the complexities of differential privacy.

1 Introduction

Differential privacy (DP) [6] is complex and challenging to implement effectively. The probabilistic nature of its guarantees [6], the vast range of implementation choices and tradeoffs [13], and a shortage of educational materials [12] all contribute to these challenges. Researchers have developed various educational and infrastructural resources to assist in the deployment of differential privacy, but gaps remain in how theoretical concepts are applied in practice. Consequently, bridging this gap in understanding is essential to harnessing the full potential of DP in data privacy.

The difficulty in translating academic theory into practical applications can be attributed to various usability and operational challenges [10, 4, 3]. There is a lack of detailed guidance on hyperparameter tuning in research literature [7] and minimal decision-making support for implementation parameters in open source libraries [12, 13]. Additionally, technical professionals, such as software engineers and data scientists, are expected to possess a high level of background knowledge, yet without prior experience with privacy products, they may struggle to verify DP guarantees and navigate the inherent privacy/accuracy tradeoffs [3].

Visualization has arisen as a proven solution to help users understand their implementation choices and set them correctly. User interfaces have aided individuals in setting implementation parameters in DP deployments (DPCreator [13]), visually comparing DP algorithms (DPComp [8]), and visualizing the effects of the ε parameter on privacy/accuracy tradeoffs (ViP [11]). While these interfaces demonstrate the efficacy of visual solutions, they only allow for the exploration on a limited number of parameters (primarily ε) on pre-set datasets. These interface leave users with only a partial understanding of the nuances of DP implementation by neglecting to illustrate the effects of many of the critical implementation decisions (δ, bounds on data values, accuracy metrics, DP mechanisms, composition across multiple queries, user vs event-level privacy, etc.). Additional visualization tools are needed to help guide practitioners in assessing privacy/accuracy tradeoffs across the full range of implementation decisions [3].

To address this need, we propose the development of a visualization tool that allows practitioners to experiment with privacy/accuracy tradeoffs across a broad spectrum of implementation decisions. The visualization tool would be designed for a technically skilled individual new to differential privacy who either by themselves or on a team is working to implement differential privacy in their organization. To ground the visualization tool in real-world implementation challenges, we will first conduct formative interviews with practitioners actively involved in differential privacy deployments. These interviews will identify the most challenging implementation parameters and determine how visual tools can enhance their understanding. The interactive visual tool developed from these interviews will then complement existing resources like open-source libraries, research publications, and textbooks, helping to bridge the theoretical-practical divide in differential privacy applications.

In this paper, we cover our interview methodology and introduce our visualization prototype. Section 2 covers the interview research questions, protocol, and participant choices.

[*] We acknowledge helpful feedback on our interview scripts and protocols from Jayshree Sarathy, Priyanka Nanayakkara, and Elissa Redmiles.

[†] Email: panavas.l@northeastern.edu

[‡] Emails: {s.rho|mhb2189|rac2239}@columbia.edu

[§] Emails: {wp2279,rwright}@barnard.edu

Section 3 discusses the current visualization prototype and design choices. We have received IRB approval for our interview study, and will complete the interviews during summer 2024. During this time we will refine the visualization tool and be prepared to show the interview findings and the next iteration of the visualization tool during the TPDP conference.

2 Formative Interviews

To explore the practical challenges faced during DP implementations, we are conducting formative interviews with DP practitioners. These discussions will identify educational barriers in DP, uncover how visualization can serve as an aid, and determine the most effective visualizations to show. Informed by their insights, we will refine our initial prototype (Figure 3), engaging in a cycle of iterative development using feedback from potential users.

We have designed a formative (semi-structured) interview study [1] on how interactive visual systems can support differential privacy practitioners. Our interview and analysis are centered around two research questions that will later guide the design of our visualization system.

RQ1 For technical individuals newly introduced to differential privacy, which concepts are most challenging to grasp, and what factors contribute to these challenges?

RQ2: Which visualization strategies would improve new technical users' comprehension and implementation of differential privacy?

Interviews: Our interviews were developed through multiple iterations of pilot studies, refining both the research and interview questions. They target two main areas: educational challenges and visualization strategies. For educational challenges, participants will be asked to share experiences about barriers encountered during recent differential privacy implementations. Following this, the interviewees will conduct a participatory sorting activity using an online whiteboarding tool, where they arrange common DP parameters on a continuum from 'least understood' to 'most understood' (Figure 1). This visual method was chosen because it tangibly highlights variations in understanding, allowing participants to provide more precise and thoughtful explanations [5]. The format will also help identify focus areas for future visualization efforts.

The second part of the interview focuses on an early-stage visualization prototype (presented in Section 3). Having a tangible visualization interface will help anchor the conversations, facilitating easier ideation and engagement with the concepts presented. It will also allow us to receive direct feedback on the usefulness and features of our current prototype. This preview of potential visualizations will help us identify elements that are intuitive or problematic, as well as features that might be missing but are desired by the target audience. It also allows us to understand how users interact with the tool, what assumptions they make about its use, and how they expect it to perform. Additionally, this approach helps validate the underlying design decisions and the overall direction of the project.

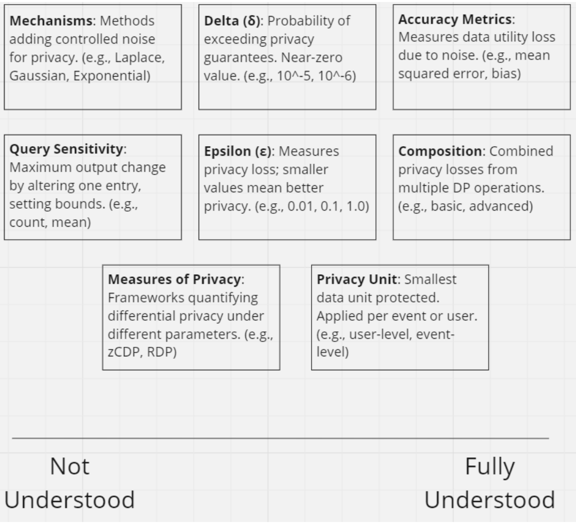

Mechanisms: Methods adding controlled noise for privacy. (e.g., Laplace, Gaussian, Exponential)

Delta (δ): Probability of exceeding privacy guarantees. Near-zero value. (e.g., 10^-5, 10^-6)

Accuracy Metrics: Measures data utility loss due to noise. (e.g., mean squared error, bias)

Query Sensitivity: Maximum output change by altering one entry, setting bounds. (e.g., count, mean)

Epsilon (ε): Measures privacy loss; smaller values mean better privacy. (e.g., 0.01, 0.1, 1.0)

Composition: Combined privacy losses from multiple DP operations. (e.g., basic, advanced)

Measures of Privacy: Frameworks quantifying differential privacy under different parameters. (e.g., zCDP, RDP)

Privacy Unit: Smallest data unit protected. Applied per event or user. (e.g., user-level, event-level)

Figure 1: Screen capture of (a condensed version of) our online whiteboarding tool, where interviewees will sort which implementation parameters cause the most difficulty when deploying differential privacy.

Participants: These interviews will be conducted with technical experts in DP who have experience with the implementation of differential privacy. Our primary target group for interviewees will be technical managers who have implemented differential privacy in their organizations. We focus on these individuals because they 1) have experience managing junior colleagues who may be introduced to DP for the first time, 2) have seen practical challenges across an array of projects, and 3) have a deep technical knowledge of DP and its implementation.

Analysis: Following the interviews we will analyze the data using a thematic analysis [2]. The final themes will be connected back to the research questions and distilled into a set of design requirements for the final version of the prototype.

3 Visualization Prototype

We have developed a web-based tool that features interactive visualizations for users to explore the complexities of privacy/accuracy tradeoffs in differential privacy. This tool demonstrates how various parameters impact the utility of private data. Users can investigate questions like, "Should I use the Laplace or Gaussian mechanism on my data?", "How does changing the data bounds affect my query accuracy", and "Which composition method gives the best per-query ε". Through these visualizations, users can simulate scenarios to compare these mechanisms and delve into critical parameters such as ε, δ, bounds on data ranges, sensitivity of their queries, composition across multiple queries, and compare user-level versus event-level privacy. This tool can serve as a learning environment for technically skilled individuals new to differential privacy by providing a way to explore and understand the nuances of these choices in implementations of differential privacy. Our interface is segmented into clear pages and tabs to help guide the user through the implementation process and showcase different visualizations pages (Figure 2).

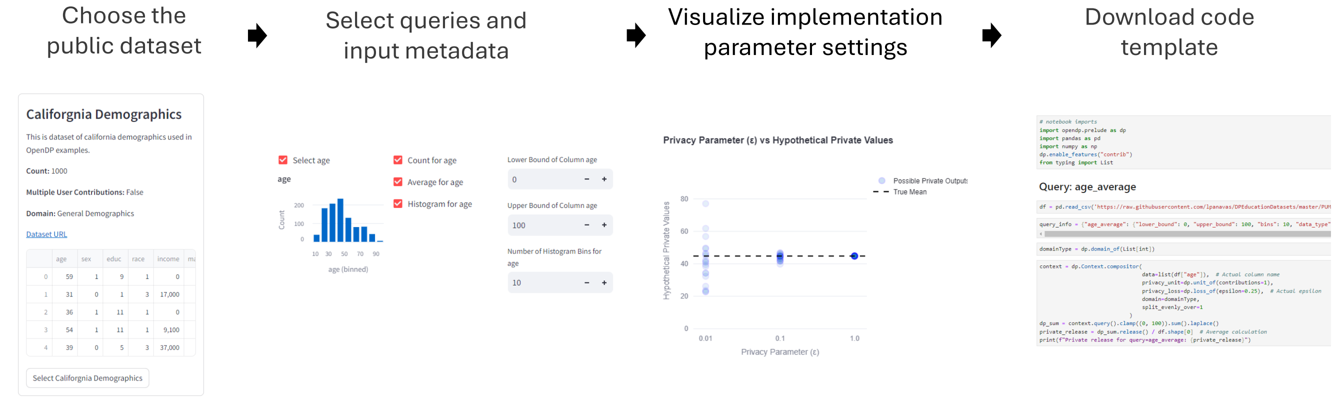
Choose the public dataset
Select queries and input metadata
Visualize implementation parameter settings
Download code template

Califorgnia Demographics

This is dataset of california demographics used in OpenDP examples.

Count: 1000

Multiple User Contributions: False

Domain: General Demographics

Dataset URL

	age	sex	educ	race	income	ma
0	59	1	9	1	0	
1	31	0	1	3	17,000	
2	36	1	11	1	0	
3	54	1	11	1	9,100	
4	39	0	5	3	37,000	

Select Califorgnia Demographics

☑ Select age

☑ Count for age

☑ Average for age

☑ Histogram for age

age
Count
200
100
0
10
30
50
70
90
age (binned)

Lower Bound of Column age: 0

Upper Bound of Column age: 100

Number of Histogram Bins for age: 10

Privacy Parameter (ε) vs Hypothetical Private Values
Possible Private Outputs
True Mean
Hypothetical Private Values
80
60
40
20
0
0.01
0.1
1.0
Privacy Parameter (ε)

```
# notebook imports
import opendp.prelude as dp
import pandas as pd
import numpy as np
dp.enable_features("contrib")
from typing import List
```

Query: age_average

```
df = pd.read_csv('https://raw.githubusercontent.com/lpanavas/DPEducationDatasets/master/PUM
```

```
query_info = {"age_average": {"lower_bound": 0, "upper_bound": 100, "bins": 10, "data_type"
```

```
domainType = dp.domain_of(List[int])
```

```
context = dp.Context.compositor(
                                data=list(df["age"]),  # Actual column name
                                privacy_unit=dp.unit_of(contributions=1),
                                privacy_loss=dp.loss_of(epsilon=0.25),  # Actual epsilon
                                domain=domainType,
                                split_evenly_over=1
                            )
dp_sum = context.query().clamp((0, 100)).sum().laplace()
private_release = dp_sum.release() / df.shape[0]  # Average calculation
print(f"Private release for query=age_average: {private_release}")
```

Figure 2: General workflow of a user of our visualization tool. Users will: 1) select a public dataset, 2) choose the queries they want to investigate, 3) interactively experiment and visualize the private outputs for a variety of implementation settings, and 4) export code to help them release data privately locally.

3.1 Visualization Features

Our current visualization prototype has three panels for exploring parameter choices and privacy/accuracy tradeoffs, illustrated in Figure 3. The first panel displays variations in private outputs based on different parameters (Figure 3A,C). The second panel visualizes privacy-accuracy tradeoff curves (Figure 3B). The third panel enables composition comparisons for privacy parameters (Figure 3D). Each panel is further detailed below and a current working prototype can be found at: https://dpeducation.streamlit.app/.

Implementation parameter variations: The first visualization panel (Figure 3A,C) allows practitioners to qualitatively compare the private outputs across a range of parameters. Textbooks and formulas can give users mathematical definitions of accuracy, but showing how different implementation settings affect private outputs through visualizations can help give users a qualitative sense of the accuracy under various parameter settings. We allow users to experiment with the queries of Count, Average, and Histogram (categorical and continuous) and visualize changes for different epsilons, mechanisms (Laplace and Gaussian), bounds on data ranges that affect sensitivity for the Average query, accuracy metrics (Absolute and Relative additive error), and bin discretization for Histogram queries. The visualizations for Count and Average show a sample of privatized outputs generated by the selected mechanism and parameters (Figure 3C) while the histograms visualization shows only one sampled private output to reduce visual clutter (Figure 3A).

In the second visualization panel, we show a larger overview of privacy/accuracy tradeoff curves (Figure 3B), as visualized through two charts. The chart on the left shows the generated private outputs overlaid with theoretical error bounds based on the ε, the selected mechanism, and β (high confidence bound on the error metric). The line chart (right), on the other hand, shows the upper bound of the error for the selected confidence. This means that the error will likely not exceed the number on the y-axis for the ε-value found on the x-axis. The red dots on the lines indicate the user selected ε. This line chart gives users an idea of the exponential increase in error as ε decreases. It can also help users determine the minimum appropriate ε based on the elbow in the curves and their own organization's privacy needs.

Composition comparisons: The composition comparison panel helps illustrate the per query ε and δ based on the user-specified number of queries, global ε, and global δ (Figure 3D). The user can adjust any of these parameters to see the effects and gain a better understanding of which composition methods would work best for their data and analysis needed. This panel in particular is a focus area for future development.

3.2 Design Requirements

Previous research has demonstrated the effectiveness of user interfaces and visualizations in assisting practitioners of differential privacy set and understanding implementation parameters. Our prototype takes inspiration from these interfaces, DPComp [8], DPCreator [13], and ViP [11], as well as informal interviews with DP practitioners. Accordingly, we set out several design requirements for our interface:

1. **No coding:** Like DPComp, our aim is to facilitate easy browsing and performance comparison of privacy implementation choices in a user-friendly manner. Knowing that the current open source differential privacy libraries can require high levels of expertise [12], we ensure that our interface has simple guided selections for all relevant parameters akin to the DPCreator interface.

2. **Provide real datasets to explore:** Real data has complexity and patterns not always seen in simulated data. Certain combinations of parameters may work well on one dataset but fail on another. To illustrate this we provide users with a variety of different public datasets showcasing different parameters and distributions. We also provide users an interface to generate their own synthetic dataset or upload a dataset.

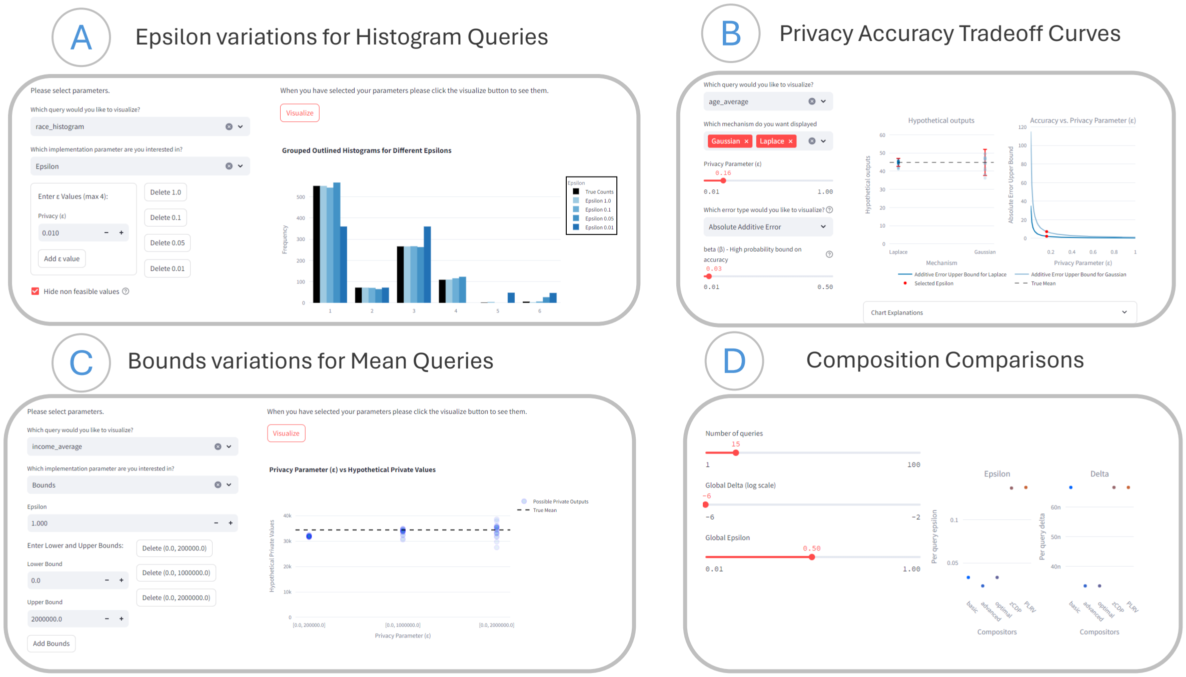
A
Epsilon variations for Histogram Queries
B
Privacy Accuracy Tradeoff Curves
C
Bounds variations for Mean Queries
D
Composition Comparisons

Please select parameters.

Which query would you like to visualize?
race_histogram

Which implementation parameter are you interested in?
Epsilon

Enter ε Values (max 4):
Privacy (ε)
0.010
Add ε value
Delete 1.0
Delete 0.1
Delete 0.05
Delete 0.01

☑ Hide non feasible values

When you have selected your parameters please click the visualize button to see them.

Visualize

Grouped Outlined Histograms for Different Epsilons
Frequency
0
100
200
300
400
500
1
2
3
4
5
6
Epsilon
■ True Counts
■ Epsilon 1.0
■ Epsilon 0.1
■ Epsilon 0.05
■ Epsilon 0.01

Which query would you like to visualize?
age_average

Which mechanism do you want displayed
Gaussian × Laplace ×

Privacy Parameter (ε)
0.16
0.01 1.00

Which error type would you like to visualize?
Absolute Additive Error

beta (β) - High probability bound on accuracy
0.03
0.01 0.50

Please select parameters.

Which query would you like to visualize?
income_average

Which implementation parameter are you interested in?
Bounds

Epsilon
1.000

Enter Lower and Upper Bounds:
Lower Bound
0.0
Upper Bound
2000000.0
Add Bounds
Delete (0.0, 200000.0)
Delete (0.0, 1000000.0)
Delete (0.0, 2000000.0)

When you have selected your parameters please click the visualize button to see them.

Visualize

Privacy Parameter (ε) vs Hypothetical Private Values
Hypothetical Private Values
0
10k
20k
30k
40k
[0.0, 200000.0]
[0.0, 1000000.0]
[0.0, 2000000.0]
Privacy Parameter (ε)
● Possible Private Outputs
True Mean

Number of queries
15
1 100

Global Delta (log scale)
-6
-6 -2

Global Epsilon
0.50
0.01 1.00

Figure 3: Our prototype has a variety of interactive visualizations to explore the privacy/accuracy tradeoffs under different implementation choices. Panels A and C allow users to choose metadata and implementation parameters, and visualize the hypothetical private outputs under those choices. Panel B illustrates a privacy-accuracy curve helping users see a larger overview of the effects of ε and mechanism choice on their private query. Panel D gives users an understanding of how different composition methods will affect their per query privacy parameter (ε, δ) values.

3. **Easy manipulation and exploration of implementation parameters:** In DPComp and ViP, users can easily manipulate and select different parameters to test. Simple sliders and dropdown menus enable quick exploration through engaging interactive elements. Our system does the same by providing easy user input for all the necessary metadata and parameter manipulation.

4. **Clear visual comparisons for privacy/accuracy tradeoffs:** ViP and DPComp give immediate interactive feedback on the private outputs depending on the parameter changes. Our system provides various visualizations to show what the private output might look like and give a visual comparison between different implementation choices.

5. **Explanations of choices and visualizations:** In their study of DPCreator, Sarathy et al. found that users wanted more explanations of the choices being made [13]. DPComp assists users by offering clear textual explanations that contextualize the visuals displayed. ViP gives helpful information in different parts of the interface. Following these examples, an interface should provide clear visuals and give additional information and guidance on what the users are seeing and how they should interpret it.

6. **Data Export:** To help tie the visualizations back to the user's own workflow, we have a method to download Jupyter notebooks with a selected query (Figure 2(4)). Users of open source libraries have previous noted the difficulty of properly implementing basic queries [12], in part due to lack of code examples. The ability to download a working starting example can reduce friction for new users to libraries such as OpenDP.

We choose to implement our system as a web based interface to maximize use and accessibility. We use the OpenDP library [14] as our DP backend and Streamlit [15] along with Plotly [9] for the front end user interface. We chose to use the OpenDP library because it is rigorously tested, with mathematical proofs and regular audits ensuring its reliability. Additionally, a robust community surrounds OpenDP, providing support and continuous improvements. The library also offers a wide range of useful functions, making it a comprehensive resource for implementing differential privacy.

4 Conclusion

Differential privacy can be challenging for new technical users to understand, and when used improperly, it can fail to provide the desired privacy or accuracy guarantees. Our long-term goal is to provide a visualization tool that is useful to differential privacy practitioners as they carry out their technical implementation of differential privacy. We have developed an initial prototype of the tool, and to inform the continued development of our tool, we are conducting a qualitative user study with DP practitioners. We will be prepared to discuss our interview findings and the resulting visualization tool at the workshop.

References

[1] O. A. Adeoye-Olatunde and N. L. Olenik. Research and scholarly methods: Semi-structured interviews. *Journal of the american college of clinical pharmacy*, 4(10):1358–1367, 2021.

[2] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

[3] R. Cummings, D. Desfontaines, D. Evans, R. Geambasu, Y. Huang, M. Jagielski, P. Kairouz, G. Kamath, S. Oh, O. Ohrimenko, N. Papernot, R. Rogers, M. Shen, S. Song, W. Su, A. Terzis, A. Thakurta, S. Vassilvitskii, Y.-X. Wang, L. Xiong, S. Yekhanin, D. Yu, H. Zhang, and W. Zhang. Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment. *Harvard Data Science Review*, 6(1), 2024.

[4] D. Desfontaines. *Lowering the cost of anonymization*. PhD thesis, ETH Zurich, 2020.

[5] L. Dewitz. Engaging participants in online interviews: Lessons learned from implementing a participatory visual approach in two explorative health information behavior studies. *Proceedings of the Association for Information Science and Technology*, 60(1):98–110, 2023.

[6] C. Dwork. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer, 2006.

[7] M. Gaboardi, J. Honaker, G. King, J. Murtagh, K. Nissim, J. Ullman, and S. Vadhan. Psi ({\Psi}): a private data sharing interface. *arXiv preprint arXiv:1609.04340*, 2016.

[8] M. Hay, A. Machanavajjhala, G. Miklau, Y. Chen, D. Zhang, and G. Bissias. Exploring privacy-accuracy tradeoffs using dpcomp. In *Proceedings of the 2016 International Conference on Management of Data*, pages 2101–2104, 2016.

[9] P. T. Inc. Collaborative data science, 2023. Accessed: 2023-05-02.

[10] A. Machanavajjhala, X. He, and M. Hay. Differential privacy in the wild: A tutorial on current practices & open challenges. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1727–1730, 2017.

[11] P. Nanayakkara, J. Bater, X. He, J. Hullman, and J. Rogers. Visualizing privacy-utility trade-offs in differentially private data releases. *arXiv preprint arXiv:2201.05964*, 2022.

[12] I. C. Ngong, B. Stenger, J. P. Near, and Y. Feng. Evaluating the usability of differential privacy tools with data practitioners. *arXiv preprint arXiv:2309.13506*, 2023.

[13] J. Sarathy, S. Song, A. Haque, T. Schlatter, and S. Vadhan. Don't look at the data! how differential privacy reconfigures the practices of data science. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2023.

[14] M. Shoemate, A. Vyrros, C. McCallum, R. Prasad, P. Durbin, S. Casacuberta Puig, E. Cowan, V. Xu, Z. Ratliff, N. Berrios, A. Whitworth, M. Eliot, C. Lebeda, O. Renard, and C. McKay Bowen. OpenDP Library.

[15] Streamlit. Streamlit – the fastest way to build and share data apps, 2023. Accessed: 2023-05-02.