SoK: Usability Studies in Differential Privacy

Onyinye Dibia*
Onyinye.Dibia@uvm.edu
University of Vermont

Prianka Bhattacharjee*
Prianka.Bhattacharjee@uvm.edu
University of Vermont

Brad Stenger*
jstenger@uvm.edu
University of Vermont

Steven Baldasty
Steven.Baldasty@uvm.edu
University of Vermont

Mako Bates
Mako.Bates@uvm.edu
University of Vermont

Ivoline C. Ngong
kngongiv@uvm.edu
University of Vermont

Yuanyuan Feng
Yuanyuan.Feng@uvm.edu
University of Vermont

Joseph P. Near
jnear@uvm.edu
University of Vermont

Abstract

Differential Privacy (DP) has emerged as a pivotal approach for safeguarding individual privacy in data analysis, yet its practical adoption is often hindered by challenges in the implementation and communication of DP. This paper presents a comprehensive systematization of existing research studies around the usability of DP, synthesizing insights from studies on both the practical use of DP tools and strategies for conveying DP parameters that determine privacy protection levels, such as epsilon(ε). By reviewing and analyzing these studies, we identify core usability challenges, best practices, and critical gaps in current DP tools that affect adoption across diverse user groups, including developers, data analysts, and non-technical stakeholders. Our analysis highlights actionable insights and pathways for future research that emphasizes user-centered design and clear communication, fostering the development of more accessible DP tools that meet practical needs and support broader adoption.

1 Introduction

In an era where data is increasingly viewed as a vital asset, the need for robust privacy protections has never been more critical. Differential privacy (DP) has emerged as a cornerstone of modern data privacy, offering robust privacy guarantees that allow for meaningful data analysis while protecting the privacy of individuals [12, 13]. Government agencies, large corporations, and research institutions have integrated DP into various data practices [21, 32], but its real-world adoption among small- to med-sized organizations and its broader acceptance among the public remain limited.

The inherent complexity of DP leads to barriers for real-world implementation and user understanding. First, the effective implementation of DP requires adequate understanding of its underlying mechanisms, such as the concept of privacy budget (ε) that determines the level of privacy protection [32]. The implementation often requires specialized DP tools, which can be difficult to use

even for data practitioners with technical backgrounds [35]. Second, it is hard to accurately communicate DP's privacy protection levels, which can prevent the general public and other stakeholders from fully understanding the privacy implications of a given DP implementation. Failure to convey DP's technical details may lead to misconceptions and decreased acceptance [50, 51].

These barriers underscore the challenges around **DP's usability**, which we define in this systemization of knowledge (SoK) paper as **the ease and effectiveness that different users engage with DP in ways that suit them.** Specifically, this is a two-pronged definition of usability considering the needs of two distinct user groups: **technical users** (e.g., data practitioners, software developers) who often need to use DP tools to implement DP in their systems, where we adopt the classic usability definition for software (e.g., learnability, efficiency, error prevention) [36] to consider DP tools' usability; and **end users** (e.g., data subjects or the general public) of DP implementations, where we consider usability as the effectiveness of communicating the privacy implications of DP, since this is the primary way end users engage with DP.

To navigate the challenges around DP's usability, this paper aims to synthesize and critically evaluate existing research on both the usability of DP tools [1, 16, 19, 23, 29, 35, 40] and strategies for communicating DP concepts effectively to diverse audiences [3, 9, 15, 21, 22, 30, 32, 41, 48–51]. By systematically reviewing the literature, we explore two central aspects of DP usability: the creation and selection of DP tools, and the communication of DP-specific parameters. Usability studies, in this context, refer to empirical research studies that examine not only the ease of use of DP tools but also the methods for effectively communicating about DP guarantees. Our work aims to provide researchers with an overview of usability studies in DP and to highlight important directions for future work.

In summary, we make the following contributions:

- **Synthesizing existing usability studies**: We review and summarize findings from 27 studies on DP tool usability and communication strategies, analyzing their methodologies, target audiences, and study outcomes.
- **Identification of Best Practices**: We identify best practices for conducting usability studies in DP, with the objective of improving the design and evaluation of DP tools, as well as the communication strategies used to convey DP concepts.

---

*These authors contributed equally to this work.

- **Highlighting Research Gaps**: We highlight open research challenges in DP usability, providing a roadmap for future research. These challenges include the need for more standardized communication methods, the development of more user-friendly tools for technical users without DP expertise, and the integration of usability considerations into implementations.

Overview of this paper. The technological development of DP is on a trajectory where usability increasingly influences the decisions made by DP adopters. Our analysis framework divides DP usability into two main strata, "DP tools" and "DP communication", each with its own set of comparison points and discussion themes.

Section 2 summarizes the basics of DP and recent relevant research. Section 3 describes our process for finding and choosing papers to include in this SoK work. Section 4 presents our analysis of the methodologies employed these usability studies related to both DP tools and DP communication. Section 5 delves into the findings of usability studies centered on DP tools and DP communications, separately. The first part of the section summarizes how DP software tools designed are assessed; the second part of this section reviews usability studies around DP communication. Section 6 identifies specific takeaways and challenges that emerged from our review, along with our recommendations and proposed solutions; Section 7 discusses important open questions for future research.

2 Background & Related Work

Differential privacy (DP). Differential Privacy (DP) [12, 13] is a formal privacy definition designed to enable the statistical analysis of data while safeguarding individual privacy. At its core, DP introduces randomness, typically in the form of noise, into the analysis process to obscure the impact of any single data point, thereby preventing the identification of individuals within a dataset. Formally, two datasets D and D' are considered neighboring datasets if they differ by only one individual's data. A mechanism M satisfies (ε, δ)-DP if, for all neighboring datasets D and D' and for all possible outcomes S, $\Pr[M(D) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D') \in S] + \delta$. Here, epsilon ($\varepsilon$) represents the privacy parameter or privacy budget, where a smaller ε value indicates stronger privacy guarantees, while a larger ε value indicates weaker privacy, and delta (δ) is a small positive value that accounts for unlikely outcomes. (δ) is set to zero for *"pure"* DP. Noise, often drawn from distributions such as the Laplace or Gaussian, is used to achieve differential privacy [12, 13].

SoK work in DP. Mainstream DP research has examined various aspects of the design of DP algorithms, covering the theory of DP [13], various DP definitions [11], and interpretations of the guarantee [46]. Previous DP SoK work focused on specific algorithms for histograms [34], database queries [33], mobility trajectories [27], and graphs [28]. More recently, the DP research community has begun to acknowledge the importance of making DP more usable to different stakeholders [8]. Cummings and Sarathy [10] further identified research gaps around usable DP, such as the need to tailor DP communication for different audiences and the lack of clear guidance for interpreting privacy-loss parameters (ε, δ).

However, there is no SoK work on existing efforts to make DP more usable, which is essential for broadening the real-world adoption of DP. Therefore, we set out to perform a systematic review of existing research on the usability of DP, which broadly refers to studies at the intersection of DP and human-computer interaction that involve usability evaluations or user studies. In our SoK, we refer to this body of work as "DP usability studies." Our goal is to summarize results in this area for researchers in the DP and human-computer interaction research communities, and to highlight important open questions for future research.

DP & the law. A significant body of work in the law community has examined the connection between formal privacy definitions like DP and legal requirements. Solove [43] presents a taxonomy of privacy informed by laws and regulations; the taxonomy does not consider formal privacy definitions. Solove's taxonomy is based on a deep analysis of legal opinions, and is one starting point for subsequent work connecting formal privacy definitions like DP to legal and regulatory requirements. Nissim and Wood [38] analyze the gaps between *normative* definitions of privacy (including laws and regulations) and *technical* definitions of privacy (including differential privacy). They identify significant gaps between the two, and recommend both improving privacy regulations and considering additional layers of technical definition (e.g. contextual integrity [37]) to bridge these gaps. Sokolovska and Kocarev [42] develop a legal and a technical framework for privacy, then analyze the differences. Like Nissim and Wood, they find that additional layers of formalism may be needed to align with legal requirements.

GDPR compliance is particularly well-studied. Holzel [18] examines the specific question of how DP relates to GDPR's anonymization requirement, and identifies several important gaps between the GDPR's requirements and the standard central-model approach for deploying DP. Cohen and Nissim [7] examine the specific legal concept of "singling out" in the GDPR. They provide a formal definition of this concept, and prove that DP ensures protection against singling out.

Additional work has focused on specific deployments as case studies for analyzing legal issues associated with DP. Chin and Klienfelter [5] analyze the connection between DP and the law via the case study of Facebook's advertising interface. They identify specific properties of public-facing deployments that make DP an appropriate choice for legal compliance, based on Solove's taxonomy and their own analysis of legal requirements. Cohen et al. [6] analyze the specific legal requirements associated with the Decennial Census, and find that the DP-based approach adopted by the US Census Bureau generally satisfies them. These requirements include both privacy and certain specific utility requirements (e.g. for balancing electoral districts).

3 Review Procedure

We identified 27 research papers from a literature search across four digital libraries. Here, we detailed our review procedure.

Paper Selection. To summarize the research development around the usability of Differential Privacy (DP), we conducted a structured literature search across four digital libraries: Google Scholar, ACM Digital Library, IEEE Xplore, and Semantic Scholar. The search was performed on July 25, 2024, and targeted papers published after 2010. We used the following search string: `"differential privacy" AND ("usable" OR "usability") AND "participants"`. This combination was chosen to capture empirical studies that explore human-centered aspects of DP—specifically, studies that involved

participants and addressed usability either in tool interaction or in communication of privacy guarantees. Some relevant studies may use alternate terminology such as "comprehension," "perceptions," or "accessibility." Our manual review of alternative search strings including these terms did not reveal additional relevant papers missed by our original search terms.

Our inclusion criteria required that papers (1) conduct a study with human subjects (2) related to the ease and effectiveness that different users engage with DP. We compiled the search results into a codebook (as a spreadsheet) containing 757 potentially relevant papers for manual review. Since Google Scholar does not expose an API, we manually scraped the first 200 results. Because Semantic Scholar ranks results by relevance, we reviewed the first 200 entries from that source as well. IEEE and ACM returned fewer than 200 papers in response to the search query.

Codebook development. We developed a structured codebook through a collaborative, iterative process involving all seven authors over the course of several months. We began by dividing the initial list of 757 potentially relevant papers among the authors. Each author reviewed abstracts and applied predefined inclusion criteria: papers had to either conduct a user study or propose a software or methodological tool aimed at improving usability in DP. We defined usability as encompassing one or more of the following: end-user understanding of DP concepts, curator ability to set appropriate privacy levels, and analyst ability to leverage data. After the initial individual round of review, the authors met to finalize inclusion decisions, resolve disagreements, eliminate duplicates, and arrive at the final list of 27 papers for in-depth analysis.

We completed the codebook for the remaining 27 papers by deductive coding [2] informed by prior literature and existing frameworks on evaluating usability. We iteratively refined the codes through team discussions. The final codebook included categories for both methodology and study findings. Each paper was double-coded by at least two authors. In cases of discrepancy, the authors discussed their reasoning during weekly team meetings and reached consensus through group discussion. These meetings also served to refine the interpretation of codes and resolve ambiguities. For example, we aligned on how to distinguish between "objective" and "subjective" understanding and standardized how we annotated outcomes like satisfaction and willingness to share.

Limitations. Our review of relevant papers relied on selecting the most appropriate keywords, and could have inadvertently excluded some papers. Our use of specific databases (and the omission of certain libraries like DBLP) could also have contributed to the omission of relevant research. Furthermore, the scope of this review may be limited by its focus on English-language publications, potentially overlooking valuable contributions in other languages. Finally, usability in DP is an active area of research, and our review may miss papers released after our review.

4 Methodology of DP Usability Studies

We analyzed and categorized the methodologies used in the usability studies reviewed in our SoK. Key area covered include studies' recruitment, study design and instruments, and evaluation criteria

that measure various aspects of DP tool's usability or user's DP understanding. Tables 1 and 2 summarize our analysis results of the 27 papers on DP software tools and DP communication, respectively.

4.1 Recruitment and Sample

Recruitment methods. The "Recruitment" column in Tables 1 and 2 outlines the specific methods that these DP usability studies used to recruit study participants. Studies evaluating DP software tools predominantly recruited participants through professional networks and direct outreach. Because these studies primarily targeted **technical users** of DP and did not require large samples, these tailored recruitment methods are suitable to reach potential participants with technical or data science skills required for study participation. For example, Ngong et al. [35] leveraged email lists, Slack channels, and professional networks. PSI studies [24, 29] used social science listservs and student job boards at local universities.

In comparison, studies focusing on communicating DP employed a variety of recruitment methods. Some study also leveraged professional networks [16, 17, 19] and outreach efforts [29] to reach **technical users**. However, most of these studies [15, 32, 41] resorted to online crowdwork platforms (e.g., Amazon Mechanical Turks, Prolific, Qualtrics) because they sought potential **end users** of DP. These platforms enabled researchers to quickly recruit large samples from the general public with the ability to control certain demographic variables.

In summary, most DP usability studies chose appropriate recruitment methods to reach the their targeted populations (see 4.1) and the sample sizes suitable for their study designs (see 4.1). However, current recruitment methods in DP usability studies lack width and diversity. Specifically, researchers largely relied on their own professional networks to recruit **technical users** and on crowdwork platforms to recruit **end users**. Despite the convenience, inherent sampling biases from small professional networks and participants on crowd-scouring platforms are unavoidable. In our opinion, diversifying recruitment strategies in future DP usability studies will increase the overall validity and reliability of research in this area.

Geo-location. The "Geo-location" column revealed that all DP usability studies in Tables 1 and 2 were conducted by researchers in the US and the EU with samples from developed Western countries. We believe reaching participants across countries and cultures will improve the external validity of future DP usability studies, particularly for studies targeting end users.

Sample Populations. The "Sample Populations" column in Tables 1 and 2 shows that DP usability studies in this SoK recruited end users, technical users, and technical users with DP expertise. Generally, the research questions of a specific study determined the target sample populations. Studies on DP communication primarily recruited end users from crowdwork channels like Amazon Mechanical Turk, Prolific, Qualtrics etc. because many seek to examine the effectiveness of DP explanation and user understanding where end users with diverse background are appropriate. In addition, some DP communication studies [48, 50, 51] also involves technical users because the authors wanted to get perspective from users from various technical background for better understanding and comparison with end users. Also, some communication studies

Citation	Recruitment	Geo-location	Sample Populations	Sample Size	Study Methods & Instruments	Educational Materials	Evaluation Methods					
							Objective Understanding	Subjective Understanding	Task Success Rate	Time on Task	Error Rate	Satisfaction
[19]	Net	EU	Tech	8	FOC	✓	◐	●	○	○	○	●
[35]	Net	US	TechDP	24	SURV, INT, TASK, EB	✓	●	◐	●	●	●	●
[16, 29]	Out	US	Tech	20	TASK	✓	◐	◐	●	●	●	●
[23]	Net	EU	TechDP	22	SURV, INT, TASK, EW	✓	●	●	○	◐	○	○
[1]	Out	US	Tech	8	SURV		○	○	○	○	○	◐
[31]	Out, Net	US	Tech	14	TASK, INT	✓	◐	◐	●	◐	◐	○
[40]	Net	US	Tech	19	SURV, INT, TASK		◐	◐	●	○	◐	●
Recruitment key: Net: Network Out: Outreach											
Geography key: US: United States EU: Europe											
Sample Populations key: Tech: Technical users TechDP: Technical users with DP expertise End: End users											
Study design & instruments key: SURV: Survey FOC: Focus Group INT: Interview TASK: Task-based Usability Test EB: Between-subjects EW: Within-subjects											

Table 1: Methodologies used in studies of DP software tools.

only used technical users because they require DP knowledge to interpret certain study materials [16, 17, 29].

In contrast, studies examining DP tools often require participants with certain technical backgrounds, such as [16] to be able to use DP tools. Specifically, we observe that some studies [23, 35] further differentiate technical users by their prior DP expertise This is necessary because those technical users integrate software tools in their product pipeline to add DP to their existing system.

Sample size. The sample sizes in these DP usability studies vary greatly but generally align with their chosen study methods (detailed in 4.2). Studies that involved in-depth qualitative methods (e.g. focus groups, interviews) or time-consuming task-based evaluation had small sample sizes of no more than 30 participants. These include all studies evaluating DP software tools (Tables 1) and a handful of studies that focused on communicating DP [16, 17, 19–21, 29, 30, 47, 48]. Studies that employed the web-based survey method [3, 9, 14, 15, 22, 32, 41, 47, 49–51] had relatively large sample sizes, from a few hundred to more than a thousand participants. Most of these studies focused on communicating DP to end users and leveraged crowdwork platforms for recruitment.

Studies targeting technical users [1, 19, 31, 40, 51], including individuals with data science background, programming skills or prior DP knowledge, tended to have small samples. This is partially due to the difficulty to recruit participants with technical expertise. Overall, mostly of the DP usability studied reviewed in this SoK had adequate sample sizes compared to published work in usable security and privacy research area utilizing similar study methods.

4.2 Study Methods & Instruments
DP usability studies reviewed in this SoK employed different research methods and various instruments to gather quantitative

and qualitative data. Below, we describe how these methods and instruments were used in DP usability studies.

Survey (SURV). Surveys are structured questionnaires to collect quantitative and qualitative data relative large samples. This method is suitable to gauge user perceptions, comprehension, and decision-making, or perform web-based experiments. For instance, [9, 15] utilized surveys to evaluate DP communication effectiveness and user expectations. Survey questionnaire can also be a supplemental instrument for other methods. For example, Ngong et al.[35] used surveys in task-based usability test to gauge pre-task and post-task data from participants.

Interview (INT). Interviews enable in-depth exploration of participants' experiences, reasoning, and feedback, making them valuable for qualitative insights. For example Nanayakkara et al. [32] and Xiong et al. [50] both used this approach to compare more versus less visually descriptive DP representations.

Focus Group (FOC). Focus groups gather qualitative insights into collective perceptions and attitudes from a group of participants. For example, [19, 47] used focus groups to explore privacy risks and trade-offs, informing better communication strategies.

Task-based Usability Test (TASK). Task completion exercises are a research instrument to assess a tool's usability by observing users complete key tasks using the tool. Collected data include success rates, efficiency, or qualitative observations of user behavior. For example, [19, 23, 35] used this method to evaluate user performance and decision-making with differential privacy tools.

Experiment (EW and EB). Experiments are a research method to compare the results of at least two groups of users in different controlled conditions. The study design of experiments can be within-subject or between-subject. Within-subject design (denoted by **EW**) involves each participant in multiple conditions, allowing for comparative analysis of different tool configurations with relative small samples. For example, Bullek et al. [3] employed

Citation	Recruitment	Geo-location	Sample Populations	Sample Size	Study Methods & Instruments	Educational Materials	Evaluation Methods				
							Objective Understanding	Subjective Understanding	Satisfaction	Willingness to Share	Perception
[15]	AMT	US	End	343	SURV, EB		●	●	○	○	◐
[9]	AMT	US	End	1216	SURV, EB		○	◐	○	◐	◐
[20]	Pro	US	End	22	INT	✓	◐	◐	○	○	○
[19]	Net	EU	Tech	8	TASK, FOC	✓	○	○	○	●	○
[21]	Pro	EU	End	30	TASK	✓	○	●	○	●	●
[51]	AMT	US	TechDP, End	300	SURV, EB	✓	◐	●	○	◐	●
[17]	Net	EU	Tech	24	SURV	✓	○	○	●	◐	○
[16]	Net	US	Tech	20	SURV, TASK	✓	◐	○	○	◐	○
[22]	AMT	EU	End	990	SURV, EB	✓	◐	◐	○	◐	○
[48]	Pro	US, EU	TechDP, End	28	SURV, EB	✓	◐	◐	◐	◐	◐
[49]	AMT	US	End	465	SURV, EB		◐	◐	○	◐	○
[49]	AMT	US	End	581	SURV, EB	✓	◐	◐	○	◐	○
[49]	AMT	US	End	468	SURV, EB	✓	◐	◐	○	◐	○
[49]	AMT	US	End	278	SURV, EB	✓	◐	◐	○	◐	○
[49]	AMT	US	End	540	SURV, EB	✓	◐	◐	○	◐	○
[3]	AMT	US	End	228	SURV, EW, EB	✓	○	◐	◐	◐	
[41]	Q	US	End	300	SURV, EB	✓	●	●	●	◐	○
[29]	Out	US	Tech	28	TASK		◐	◐	◐	○	◐
[50]	AMT	US	TechDP, End	300, 295	SURV, EW, EB	✓	◐	◐	○	◐	○
[32]	Pro	US	End	963	SURV, EB	✓	◐	◐	●	◐	○
[47]	Net	EU	End	18	FOC		○	○	○	◐	○
[47]	Con	EU	End	243	SURV, EB		○	○	◐	○	●
[14]	AMT	US	End	308	SURV, EB	✓	●	○	○	●	○
[30]	Out	US	End	16	TASK, EW, INT	✓	●	●	○	○	○

Recruitment key:	**AMT**: Amazon Mechanical Turk **Pro**: Prolific **Q**: Qualtrics **Net**: Network **Out**: Outreach **Con**: Contractor
Geography key:	**US**: United States **EU**: Europe
Sample Populations key:	**Tech**: Technical users **TechDP**: Technical users with DP expertise **End**: End users
Study design & instruments key:	**SURV**: Survey **FOC**: Focus Group **INT**: Interview **TASK**: Task-based Usability Test **EB**: Between-subjects **EW**: Within-subjects

Table 2: Methodologies used in studies communicating DP

this approach to compare user preferences across multiple settings. Between-subject design (denoted by **EB**) ensures each participant is only exposed to one condition, which is beneficial for isolating the effect of the specific condition. This design is well-suited for studies utilizing crowdwork platforms, where larger, more diverse samples can be recruited, to assign each participant into a unique condition. Most DP usability studies utilizing experiments adopted a between-subjects design (e.g. [3, 9, 14, 15, 22, 32, 35, 41, 47–51]).

Educational materials. Some DP usability studies utilized a special type of research instrument to educate participants with basic DP knowledge before study procedures, as shown in the "Educational Materials" column of Tables 1 and 2. For example, Gaboardi et al. [16] and Murtagh et al. [29] incorporate educational resources like handouts, tutorials, or short videos to establish a "common frame of reference" among participants. DP educational materials help establish a consistent baseline of knowledge among participants, when such knowledge is needed for the study procedure.

Many studies in this SoK employed various methods in pursuit of variations of the research question: Can different groups of users understand and properly apply differential privacy given a range of realistic scenarios? Table 3 lists best practices for Experiment Design & Instrument Type. Next-generation DP usability studies will use the same basic set of methods and instruments, but will also benefit from greater adoption of DP, which in turn should provoke focused research questions that help solve critical DP issues.

4.3 Evaluation Metrics

The DP usability studies reviewed in this SoK used various evaluation metrics to assess the usability of DP tools and the effectiveness of DP communication.

4.3.1 Shared metrics. Both studies on DP tools and DP communication measured the following aspects regarding DP's usability.

Objective understanding. This involves quantitative or qualitative measurements of participants' knowledge about what DP terms mean, their ability to reason about DP, or their ability to solve DP tasks. For example, Ngong et al. [35] assessed objective understanding by evaluating participants' abilities to accurately interpret DP parameters and apply DP concepts in practical scenarios. This is typically measured through pre- and post-study quizzes, task-based assessments, and scenario-based questions, which help accurately gauge participants' understanding with standard questions.

Subjective understanding. This refers to how participants conceptualize DP, including their mental models, personal interpretations, concerns, and ethical intuitions about DP [9]. Researchers assess it using qualitative methods such as interviews, focus groups, and self-reported surveys [15, 32]. Nanayakkara et al. [32] echo methods applied previously in Franzen et al. [15], where they measured self-assessments of confidence for DP understanding by Likert scales and analyzed statistically. Nanayakkara et al. [32] took the extra step to have subjects describe DP privacy protection in their own words. While subjective understanding provides insights into how individuals internalize and process DP concepts, personal biases and varying levels of technical literacy may influence responses, requiring careful interpretation.

Satisfaction. This measures how satisfied participants are that differential privacy provides adequate protection for personal information or that a fair balance between privacy and utility has been struck. This is measured through post-study surveys and interviews to gather participants' satisfaction ratings and qualitative feedback. Satisfaction surveys are effective for gauging overall user sentiment, though they can be subjective. Cummings et al. [9] equated satisfaction with subjects' met (or unmet) expectation for privacy protection based on task outcome. Ngong et al. [35] used common recommendation comparison metrics (Net Promoter Score, System Usability Score) to compare satisfaction levels for different DP tools.

4.3.2 Metrics for DP tools. Many studies on DP tools used classic software usability metrics to evaluate DP tools, as detailed below.

Task success rate. This measures how often participants complete tasks using DP tools. This metric is crucial for understanding the practical usability of the tools. It is typically measured through direct observation and logging of task completion rates during usability testing sessions. Task success rate is a straightforward and reliable metric for assessing practical usability. Quite a few papers assessed study participants' understanding of DP through task completion [16, 19, 21, 29–31, 35, 40]. A subset of these studies used talk-aloud methods to capture participants' thought process and decision-making [30, 31, 35, 40].

Time on task. This quantifies how long participants take to complete tasks and reflects "efficiency" for tool-based problem solving according to Nielsen's "5 Components of Usability" [36]. Efficiency encapsulates the speed and ease of the workflow leading to task completion. Systems that are inefficient in comparison can help identify usability issues or learning curves associated with the tools. Studies by Ngong et al. [35] and Murtagh et al. [29] included time on task as a key metric to evaluate the efficiency of the tools being tested. This is measured by timing participants during task completion exercises. This metric effectively highlights efficiency and potential areas where the user experience can be improved.

Error rate. Error rate measures the frequency of errors made by participants while using the tools. Pilot studies help to scope possible errors in advance of formal investigation, especially since error rates depend on having clear definitions for what does and does not constitute a task- or tool-related error. Several studies [1, 23, 35] used error rate to gauge how often users make mistakes and the types of errors encountered.

4.3.3 Metrics for DP communication. Some DP communication studies also measured other subjective constructs of end users.

Willingness to share. This evaluates how willing participants are to contribute their own private data in scenarios where differential privacy (DP) techniques are applied. This is often measured through surveys and interviews, where participants are asked about their comfort levels and willingness to share personal data under various conditions. Several studies have examined participants' willingness to share data after being informed about DP implementations, finding that clear explanations and illustrations significantly influenced their willingness to share data [3, 48–51].

Perception. This evaluates participants' trust, confidence, and attitudes toward DP implementations and the organizations deploying them. It examines how participants judge the appropriateness, fairness, and effectiveness of real-world DP deployments. It assesses how people feel about DP applications and their confidence in the privacy guarantees they provide. Hands on studies with interactive DP tools like DP Creator [40], ViP [30], PSI [16, 29] and DPP tool [19] compelled users to determine for themselves appropriate DP parameters like privacy budget and epsilon in order to complete a simulated task. Each study followed up the exercises with interviews or focus groups to fully understand users' perception of DP parameter appropriateness.

5 Results
5.1 DP Software Tools

This section corresponds with Table 3, which summarizes the findings from various studies that investigated different software tools designed for Differential Privacy (DP). The table highlights key aspects such as the types of tools developed, the target users' expertise and role, the support for privacy budgeting, utility analysis, level of automation, flexibility, correctness checking, and deployment model. These elements are crucial in evaluating how well these tools facilitate the implementation and understanding of DP across diverse user groups.

5.1.1 Tool Design. **API tools.** API tools assist practitioners in implementing DP solutions by writing code in a mainstream programming language. These tools are designed to be used to prepare

Citation	Type of tool	Target expertise	Target Role	Parameter-setting	Utility analysis	Automation	Flexibility	Correctness checking	Deployment model
[19]	V	**Tech, TechDP**	Curators	●	◐	○	●	○	LDP
[35]	A	**Tech, TechDP**	Curators, analysts	●	○	○	●	○	CDP
[17]	H	**TechDP**	Analysts	●	○	○	●	●	LDP, CDP
[16, 29]	V	**Tech, TechDP**	Owners, curators, analysts	●	●	○	○	○	CDP
[25]	H	**Tech**	Curators, analysts	◐	○	◐	◐	○	LDP, CDP
[31]	V	**Tech, TechDP**	Curators, analysts	●	●	◐	◐	◐	CDP
[40]	V	**Tech, TechDP**	Curators, analysts	●	◐	●	○	●	CDP
Type of tool key: V: Visual, A: API-based, H: Hypothetical									
Target expertise key: **Tech**: Technical users **TechDP**: Technical users with DP expertise **End**: End users									

Table 3: Key aspects examined in studies of DP software tools

DP data releases or implement new DP systems, without needing to implement DP mechanisms from scratch.

Visual tools. Visual tools for DP are designed to provide a user-friendly interface for preparing a DP data release without writing any code. These tools often include visual aids and interactive elements that guide users through the process of setting privacy parameters and analyzing data, making them particularly valuable for end users and technical users without DP expertise [16, 19, 40].

Hypothetical tools. Some tools have been proposed, but not implemented. Some of these hypothetical tools align with the current implementation of API or visual tools [17]. Others are entirely new [25, 29], and results of users studies on mockups of these tools can provide valuable guidance for tool design.

Target role: who uses the tool? DP systems and DP data releases involve collaboration between people in different roles, and different tools are designed to help different roles. Tools may be designed to help the *data analyst*, who designs and implements the pipeline that takes sensitive data as input and produces DP results as output; the *data subject* (sometimes called *data owner*), who submits the sensitive data for analysis in the first place; the *data curator* (or *data steward*), who accept data submissions and maintain the database of sensitive data. Most existing tools are designed to help the data analyst to implement software pipelines that successfully enforce DP with the desired privacy parameters and level of utility [16, 17, 25, 29, 35]. A subset of tools are also designed to help the data curator [29, 40].

Deployment model. Current tools assume the central model of DP, in which the sensitive data is collected on a central server by a trusted data curator. Tool support for other deployment models, including the local and shuffle models, is currently lacking.

5.1.2 Target Expertise. Many tools claim to improve usability for technical users without DP expertise. Studies of these tools group participants according to DP expertise.

Tech: This group includes technical users *without* DP expertise. These users benefit the most from visual tools that provide a user-friendly interface, require minimal DP knowledge and guide users towards the right usage of DP. Tools like PSI [16] and DP Creator [40] are particularly valuable for this group as they simplify the process of setting privacy parameters through editable forms and visual aids, reducing the need for DP expertise.

TechDP: This group includes technical users *with* DP expertise. These users are the typical target audience for API-based tools that offer extensive customization and flexibility. Such tools are designed to integrate with various programming languages and data processing frameworks.

5.1.3 Support for Setting Parameters. Every API tool [35] and most visual tools [16, 40] we examined provided a method for setting the privacy budget (ε or corresponding parameter(s)).

Visual tools like PSI [16] and DP Creator [40] provide user-friendly interfaces for setting privacy budgets. These tools offer default options and allow users to allocate the total privacy budget across different queries through editable forms. PSI, for example, lets users reserve a portion of the budget for specific queries, with the rest divided among data attributes, while DP Creator allows for updating privacy parameters for different dataset entities.

API-based tools offer more flexibility for technical users with DP expertise users to set and adjust the privacy budget programmatically. This flexibility, while powerful, increases complexity and requires a good understanding of both DP concepts and the API itself to avoid misconfigurations [35]. In all cases, users need to understand the privacy parameters, either through their own knowledge of DP or by consulting an expert—a challenging task, as indicated by studies on tools [35, 40] and the communication studies described in Section 5.2.

Some hypothetical tools aim to further improve support for budgeting by automating the process and providing clearer guidance to users. For example, the hypothetical tools discussed by Murtagh et al. [29] do not recommend specific privacy parameters but instead focus on providing mechanisms to estimate the risk of data sharing.

Another important aspect of support for budgeting is the setting of clipping parameters, which define the range of allowable data values. Both visual [16, 40] and API tools [35] allow users to set these bounds, and they often include warnings to highlight the importance of appropriate clipping to maintain both privacy and utility. Inappropriate clipping parameters can lead to poor data utility, especially if they affect a significant portion of the data.

Support for utility analysis. Utility analysis should tell us if the DP release provides sufficient accuracy to enable the desired downstream analyses. Some tools provide accuracy metrics to help the target user understand the utility of the release and navigate the privacy/utility tradeoff [16, 40]. Support for this kind of analysis is important, since analysts report concern that the DP release will not be sufficiently accurate for the desired uses [17].

Oishee Bhattacharjee, Brad Stenger, Steven Baldasty, Mako Bates, Ivoline C. Ngong, Yuanyuan Feng, and Joseph P. Near

PSI [16] calculates mean absolute error for every selected column on the output as utility metric, and allows adjusting the confidence level. DP Creator [40] provides a similar accuracy report, but is less transparent about how error is calculated.

Level of automation and flexibility. Level of automation for DP software tools indicate the tradeoff between how many parameter values have to be decided by user and how many the system will determine automatically. Both visual and API tools require the analyst to provide values for most parameters [16, 35, 40] (e.g. lower and upper bounds for clipping) and to select the precise mechanism to be used. In both contexts, analysts struggle with setting these parameters [35, 40], suggesting that increased automation would be helpful for users.

Level of flexibility refers how much control the user has in determining how DP will be achieved. Existing visual tools are less flexible: they provide wizard-like interfaces that support a small, fixed set of analyses [16, 40]. API tools differ in their levels of flexibility [35]: some provide access to low-level mechanisms and enable users to build new mechanisms on top (e.g. OpenDP and DiffPrivLib), while others aim to provide a higher-level API that requires less expertise (e.g. Tumult Analytics). Existing tools (both visual and API-based) tend to require data to be stored in a particular format or underlying data store, limiting integration with existing data processing infrastructure.

Level of correctness checking. In the context of Differential Privacy (DP) tools, correctness checking refers to the tool's ability to ensure that the deployed system is free of privacy bugs. Interactive tools like DP Creator [40] prevent bugs by limiting available functionality to correct operations. In flexible API-based tools, ensuring correctness is more difficult. Ngong et al. [35] highlight that even experienced practitioners can inadvertently implement DP mechanisms incorrectly when using API-based DP tools. Despite their confidence, these users often made mistakes that resulted in insecure or flawed DP implementations. This finding underscores the necessity of built-in correctness-checking tools to ensure the accuracy and security of DP implementations.

5.2 Communicating DP

5.2.1 Communicating Using. Effective communication in differential privacy utilizes descriptions, visual aids - visualizations, and pictures/diagrams. Each method caters to different aspects of comprehension and informational needs.

Descriptions (Text). Textual descriptions serve as the foundation for explaining DP concepts. They are most effective when they simplify complex ideas and relate them to everyday situations. Narratives that contextualize DP settings in familiar scenarios can significantly improve comprehension, as evidenced by the textual strategies recommended by Cummings et al. [9]. Their study highlights the importance of using simple, relatable language to explain differential privacy, making it accessible to end users. Franzen et al. [15] emphasize using risk communication formats to articulate the implications of DP parameters, such as epsilon (ε), clearly and directly. This approach helps users understand the trade-offs between privacy and utility, enhancing their overall comprehension. Xiong et al. [51] also support the use of clear textual explanations to describe DP models, noting that simplifying complex mathematical

Citation	Communicating about		Communicating using		
	Epsilon or other parameters	Deployment model (CDP vs LDP)	Descriptions (text)	Visualizations	Pictures & diagrams
[15]	◑	○	●	○	○
[9]	◑	○	●	○	○
[20]	○	○	◑	◑	◑
[21]	◑	●	○	●	○
[51]	◑	●	◑	○	●
[17]	○	◑	◑	○	○
[44]	○	○	◑	○	○
[22]	○	●	●	○	○
[48]	●	○	◑	○	○
[49]	●	●	●	○	○
[3]	○	◑	◑	●	○
[41]	●	●	●	○	○
[32]	●	○	●	●	●
[47]	○	○	◑	○	●
[14]	●	○	●	●	◑
[30]	●	○	○	●	○

Table 4: Key aspects examined in studies communicating DP

principles significantly aids non-technical audiences in understanding. Also, the work of Garrido et al. [17] and Kühtreiber et al. [22] highlights the importance of detailed textual descriptions in improving users' comprehension and data-sharing decisions. These studies suggest that clear, concise descriptions can effectively communicate the nuances of DP, thereby increasing users' willingness to share their data.

Despite these advantages, textual descriptions alone may not always suffice to convey the complexities and nuances of differential privacy. Additionally, the current descriptions of DP are often insufficient to help users make informed decisions, lacking consistency and standardization [9]. This emphasizes the need for new, standardized descriptions to improve user understanding and trust in DP implementations.

5.2.2 Visualizations. Across studies, visualization has emerged as a natural strategy for communicating with end users. Our review suggests that visualizations—especially icon arrays, metaphoric illustrations, and interactive sliders—are more effective than text alone in communicating DP concepts.

Visualizations of privacy. Franzen et al. [14] evaluated five interface designs to support data donors in making privacy decisions in a between-subjects online experiment with 378 end users. The study measured *concern-decision consistency*—the alignment between participants' stated privacy concerns and their sharing choices—as the primary outcome. Results showed that interfaces incorporating visual risk representations led to more consistent and informed decision-making than text-only designs.

Nanayakkara et al. [32] developed and compared three explanation formats for communicating the privacy implications of the ϵ

parameter in DP: a textual odds description (ODDS-TEXT), a visual odds representation using frequency-framed icon arrays (ODDS-VIS), and concrete output examples (SAMPLE REPORTS). The study was conducted as a randomized vignette-based online experiment with 963 end users. The evaluation measured objective risk comprehension, subjective privacy understanding, and willingness to share data. Results showed that ODDS-VIS improved both objective and subjective understanding compared to the other formats.

Wen et al. [48] conducted a design-space exploration to evaluate how different explanation formats affect user understanding of Local Differential Privacy (LDP). They tested three formats: (1) textual descriptions, (2) data tables illustrating algorithmic behavior, and (3) explanatory illustrations, such as lottery-ball drawings designed to visualize probabilistic noise injection. The study involved a large-scale randomized online survey with 228 end users. Findings showed that participants in the illustration condition achieved significantly higher comprehension scores compared to those in the text or table conditions. The metaphor-driven visuals—particularly the lottery analogy—helped users form more accurate mental models of local noise mechanisms and privacy protection.

Bullek et al. [3] investigated whether visualizing the obfuscation process in the Randomized Response technique (RR)—a mechanism consistent with Local Differential Privacy—could enhance users' trust and understanding of privacy guarantees. They developed a visual explanation illustrating how responses are randomly flipped to ensure privacy. The study involved 228 participants. Outcome measures included self-reported trust, comfort with the mechanism, comprehension of how RR works, and willingness to share data. Results showed that participants exposed to the visual explanation reported higher trust and greater comfort in using RR. However, some participants made riskier data-sharing choices despite grasping the mechanism.

Visualizations of uncertainty for navigating the privacy-utility tradeoff. Nanayakkara et al. [30] introduced the ViP (Visualizing Privacy) tool to support data curators in configuring ϵ for DP data releases. The tool featured interactive sliders and real-time visual feedback that dynamically illustrated how varying ϵ values affected both accuracy and disclosure risk, including changes in confidence interval width and the probability of extreme (noisy) outputs. The tool was evaluated through a within-subjects study involving 16 clinical research practitioners with statistical training but limited familiarity with DP. Participants completed a series of ϵ-setting tasks using both ViP and a spreadsheet-based control interface that lacked visualization. The study assessed participants' accuracy in interpreting probabilistic outcomes, and their ability to select ϵ values aligned with specified privacy or accuracy goals. Results showed that ViP significantly improved performance on both measures. Participants using the tool demonstrated better understanding of the privacy-utility trade-off and made better parameter choices.

Calero Valdez and Ziefle [47] used a conjoint analysis (N=521) to study user preferences around data sharing in health recommender systems. Scenarios varied in privacy level (including DP), data type (mental vs. physical health), and purpose (scientific vs. commercial). While not focused on visualization per se, the study relied on interface sliders and privacy parameter displays. Participants favored stronger privacy guarantees and demonstrated contextual

sensitivity, showing that visual presentation of privacy settings influences preferences even without explicit DP education.

Diagrams for describing deployment models. Karegar et al. [21] developed a set of pictorial metaphors to explain the privacy mechanisms of Central and Local DP to technical users without DP expertise in the context of eHealth data sharing. These visuals used imagery such as blurred portraits and obscured shadows to convey concepts like data perturbation, anonymity, and the privacy-utility trade-off. The evaluation was conducted through 30 semi-structured online interviews. Each participant was shown one or more metaphors tailored to specific DP models and asked to describe their understanding of how the mechanism protected their data. Results indicated that the metaphors effectively conveyed general ideas such as the presence of noise and the loss of identifiability. However, several participants developed overgeneralized or inaccurate mental models—such as assuming complete anonymity—particularly when visual cues were not anchored to real-world scenarios.

Xiong et al. [51] designed explanative illustrations to compare the three major DP models: Central DP, Local DP, and Shuffler DP. Each illustration included a textual explanation and a data flow diagram that visually conveyed the model's trust assumptions, noise injection points, and privacy-utility trade-offs. The diagrams were designed using dual coding theory to reinforce learning through verbal and pictorial cues. A utility heatmap was also included to visualize the impact of DP mechanisms on data accuracy at the aggregate level. The evaluation involved an online experiment with 300 end users recruited from Amazon Mechanical Turk. Participants were assessed on comprehension of each DP model's mechanism, perceived privacy and utility, and willingness to share data in public-good and commercial-interest scenarios. Findings showed that participants exposed to the visual explanations achieved significantly higher comprehension of model structure and implications, particularly for Shuffler DP. The visualizations also improved perceived privacy protection and influenced data-sharing preferences. The study concluded that explanative illustrations can improve understanding of complex DP models and should be paired with textual scaffolding to support informed user decisions.

5.2.3 Communicating About. Effectively communicating complex Differential Privacy (DP) concepts is crucial for fostering understanding and trust among users. This section discusses the best approaches to articulate key DP parameters and the differences between common deployment models such as Centralized Differential Privacy (CDP) and Local Differential Privacy (LDP).

Clarifying parameters like epsilon is essential for explaining DP's privacy guarantees. Table 4 summarizes various methods—such as text descriptions and visualizations—used in the studies we reviewed to communicate these concepts effectively. Understanding the differences between CDP and LDP is vital for informed decision-making. Table 4 also highlights the use of text and visuals to enhance users' comprehension of these models.

Epsilon and other parameters. "Epsilon" (ε) is the primary parameter in differential privacy (DP) that quantifies the level of privacy guarantee, encapsulating the trade-off between privacy and

data utility. Our literature review reveals that effectively communicating privacy parameters hinges on simplification and contextualization. Studies employ a variety of methods, with textual explanations or descriptions being the most prevalent [9, 15, 49]. These studies such as those by Cummings et al. [9] and Franzen et al. [15] demonstrate that presenting complex mathematical concepts in simpler terms significantly enhances comprehension among non-technical audiences. Franzen et al. [15] advocate for the use of risk communication formats from the medical field to articulate ε's implications clearly and directly.

Previous studies have been inconclusive about the benefit of a particular approach for textual description of differential privacy. For example, odds-based explanation methods from the study conducted by Nanayakkara et al. [32] were found to be effective in improving objective risk comprehension, subjective privacy understanding, and self-efficacy compared to examples-based methods. However, these approaches still leave room for improvement [9].

Visual tools, while less commonly used, have shown considerable effectiveness in aiding understanding when used in conjunction with textual descriptions [32, 51]. For instance, Xiong et al. [51] designed explanative illustrations that clarify how ε functions within DP models, leading to improved user comprehension. Moreover, Nanayakkara et al. [32] explore effective methods of explaining epsilon ε, particularly through visualization techniques that help users visualize the privacy-utility trade-offs. This body of work highlights the necessity for transparency in communicating epsilon ε, suggesting that a well-informed public is more likely to make knowledgeable data-sharing decisions [9].

Deployment and threat models. Central Differential Privacy (CDP) and Local Differential Privacy (LDP) are two principal frameworks for implementing differential privacy. The two correspond to very different threat models from a security perspective, so effectively communicating these models is crucial.

Xiong et al. emphasize the importance of clear textual descriptions and visual aids to explain the role of the trusted curator and the benefits of minimal noise addition in CDP. Their research shows that visualizing the data flow from collection to anonymization and publication helps users understand how their data is protected at various stages [50]. Additionally, Xiong et al. found that descriptions focusing on the implications of DP rather than the technical definitions improved comprehension and willingness to share data [49]. Karegar et al. suggest using analogies and metaphors, such as adding noise to a signal before transmission, to help users grasp the concept of LDP. Interactive tools that allow users to experiment with data perturbation can demonstrate the effectiveness and implications of LDP, helping users visualize the privacy-utility trade-offs and understand the benefits of the model [21].

Research by Xiong et al [50, 51] and Karegar et al [21] suggests that visual tools are particularly effective in explaining these models. Their work shows that when users are provided with clear illustrations that demonstrate the data flow and noise addition in CDP versus LDP, they better understand the implications of each model on their privacy and the utility of their data [51]. Additionally, Wen et al. used diagrams and pictures to illustrate the data perturbation process of Local DP to protect user privacy, further supporting the effectiveness of visual aids in communicating DP models [48].

Existing studies do not go beyond LDP and CDP to more complicated threat models. Recent developments include the shuffle model [4], and the use of trusted hardware [26] or cryptography [39] to implement DP. The DP community has not settled on specific terminology for communicating these threat models, and comparing threat models remains challenging even for experts.

6 Takeaways, Challenges, and Recommendations

This section summarizes the major takeaways of our review, with a particular focus on challenges and limitations identified by the studies we reviewed and our recommendations for addressing those challenges. Section 6.1 summarizes DP-specific takeaways associated with study methodology; Section 6.2 summarizes takeaways from studies on communication about DP; and Section 6.3 summarizes takeaways from studies on DP software tools.

6.1 Methodology

The DP usability studies in this SoK generally followed the best practices of human-computer interaction research methods. We believe the methodological challenges for future work are recruiting diverse DP users into research effectively, and increasing the validity and reliability in measuring DP usability. Below, we detailed these challenges our recommendations to overcome them.

Diversifying participant recruitment strategies. For DP usability studies to achieve greater generalizability and to investigate deeper DP-related research questions, there is an increasing need to recruit both representative and specialized participants. Assessing the effectiveness of DP communication often requires large samples from general populations, so many studies resorted to crowdwork platforms to quickly reach scalability and recruitment goals [3, 9, 14, 15, 20–22, 32, 48–51]. However, research has revealed systematic bias in crowdsourced data from these platforms [24]. Even it is possible to control certain demographic variables to achieve a level of representativeness on platforms like Prolific, fully relying on crowdwork platforms for recruitment inherits the sample bias of the crowdworker population. Therefore, we recommend that future DP communication studies include more representative samples from the community in addition to crowdsourced samples to ensure the generalizability of research results.

On the other hand, in-depth DP usability studies, including those evaluating DP tools and DP communication among data practitioners, require participants to have certain technical and DP backgrounds. Existing studies had success recruiting from DP-specific listservs, job boards and professional networks [16, 23, 29–31, 35, 40], but samples tended to lack demographic or geographic diversity. Additionally, work-intensive studies that involve task completion exercises are resource-consuming, researchers had to proceed with smaller samples [17, 35]. To address the specific challenge to recruit specialized participants, we hope to see community efforts to establish a DP usability interest group consisting of representative groups of technical users with and without DP expertise who are interested in participating in future DP usability studies. The community can help in this process by supporting and publicizing resources for reaching DP stakeholders, such as the OpenDP mailing list and Slack channel.

Researchers in human-computer interaction acknowledged systematic bias of publishing user studies focusing on Western, Educated, Industrialized, Rich, and Democratic (WEIRD) participants [45]. The sample in current DP usability studies suffer from this bias. Growing the pool of potential study participants is the surest way to make experiments' subject more representative. DP research collaborations can extend their reach by broadening the geographic location, discipline and culture among project partners.

Accommodating different technical users. DP usability studies with specialized samples face the unique challenges of assessing and managing participants' prior technical and DP expertise. The modern data professional has proficiency in either R or Python (specifically including the Pandas and Numpy packages), and those skills are prerequisite for working with DP tools. Pre-enrollment screening shows a baseline technical proficiency [1, 35]. Screening can also indicate a level for participants' DP knowledge [35]. can help assess participants' prior DP knowledge, which are necessary for more valid testing of DP tools. Convenience samples, without formal screening, are adequate for pilot testing but when they are applied to rigorous investigation of DP usability they are a missed opportunity to define the educational path toward DP literacy and competence. Our review shows that education materials (e.g., short paragraphs, handouts, tutorials, videos) used in studies [16, 29, 35] help level the playing field of DP understanding among participants. We see an opportunity to establish standardized DP educational materials that articulate the path from technical proficiency for data analysis in R or Python to effective use of DP tools.

Improving validity through mixed-methods approach. Our SoK shows the advantages of the mixed-methods approach to examine the usability of DP. For example, Ngong et al. [35]'s usability study combined both quantitative usability metrics with qualitative data from surveys and interviews to generate deeper insights. Similarly, studies incorporating both controlled experiments and post-study interviews articulated the difficulties users faced in interpreting DP parameters [9, 30, 48–51]. The mixed-methods approach allows for sophisticated study designs, but increase the difficulty to compare results from different studies.

Improving reliability by standardizing evaluation. Section 4.3 summarizes the varying constructs and metrics used in these studies to evaluate usability. Theses measures of users' understanding and satisfaction, of users' task performances and of users' self-described decision-making are tried and true human-computer interaction methodology. Even though the evaluation methods used in DP usability studies largely align with the human-computer interaction evaluation practices, they vary in terms of specific measurements and execution, making cross-study comparison difficult. We recommend that researchers in this area should develop standardized metrics and instruments that are specific to DP. For example, creating and validating a set of standardized multiple choice questions that cover a comprehensive range of DP knowledge (e.g. DP parameters, privacy budget, deployment model) will allow future studies to use a consistent metrics to measure participants' DP understanding. This will increase measurements' reliability and enable easier cross-study result comparison.

6.2 Communication

A growing number of studies investigate communication strategies for DP, mostly with end users. We summarize the takeaways and remaining challenges in this area, and provide some recommendations for how to communicate DP guarantees to both technical users and end users.

Everyone struggles to understand DP parameters. Epsilon (ε) and other DP parameters define privacy guarantees, yet their meaning remains opaque and non-intuitive since they require users to reason about probabilistic guarantees. Unlike encryption, where higher bit lengths clearly indicate stronger security, DP parameters lack direct real-world analogies, making their implications difficult to grasp, [9, 15, 32]. Even participants with technical backgrounds struggle to interpret how changes in ε impact privacy risk and data utility [15]. Mathematical complexity further complicates communication. ε and δ rely on probability theory and worst-case guarantees, concepts unfamiliar to most users [51]. Oversimplified explanations often distort or mislead privacy guarantees, while formal definitions remain inaccessible to end users [15]. The privacy-utility trade-off is nonlinear and context-dependent, making it difficult to communicate effectively. Users often assume that lowering ε always improves privacy, overlooking its dependence on dataset size, sensitivity, and query structure [9].

No single communication strategy works universally. Text-based explanations, visual aids, and risk-based descriptions each offer partial improvements, but usability of DP studies indicate that communicating DP parameters require tailored communication strategies or adapting explanations of DP parameters for diverse audiences [51]. For instance, technical users with DP expertise might prefer mathematical descriptions, but policymakers and data practitioners often require context-driven explanations. Furthermore, DP lacks standardized benchmarks for defining "safe" ε or δ values, making it challenging for practitioners and policymakers to assess privacy guarantees [15]. Improving communication requires using accurate layman's terms, contextual examples and relatable scenarios, and risk communication formats to explain DP parameters clearly [32]. Combining textual descriptions with visual aids enhances comprehension for non-technical users [51]. Standardized, intuitive frameworks for explaining DP parameters and trade-offs are necessary to support better decision-making across stakeholders.

End users struggle to understand deployment models. Communicating deployment models (e.g. central DP vs local DP) is challenging due to misinterpretation of trust assumptions. Users often assume that LDP always provides stronger privacy or misunderstand the role of a data curator in CDP [32]. The privacy-utility trade-off remains difficult to convey, as users struggle to grasp why LDP introduces more noise while CDP preserves higher data utility [15, 51]. Without clear benchmarks, these trade-offs remain unclear. Inconsistent messaging across tools and studies further complicates understanding, as many fail to explicitly clarify which model they use and its implications, leading to confusion among practitioners, policymakers, and end-users [9].

To improve communication, explanations of DP deployment models should be simplified with clear, accessible descriptions of the privacy-utility trade-offs and security implications of each model [9]. Visual aids, such as comparative illustrations and interactive

tools, can enhance understanding of data flow and noise addition processes [21, 30, 51]. Contextual examples and real-world scenarios should be used to clarify operational nuances [32]. Metaphors and analogies can help make these abstract concepts and DP models more relatable, further aiding comprehension [21]. Standardized frameworks incorporating these approaches are necessary to support informed decision-making across diverse audiences.

End users struggle with text descriptions. To make communicating straightforward for the audience, text-based methods are popular as educational materials for differential privacy. However, simplifying complex DP concepts while maintaining accuracy remains difficult, as technical details often overwhelm users [15]. Balancing clarity and depth is essential, yet finding the right level of detail for diverse audiences is challenging [50]. Additionally, making textual explanations relatable and applicable across different user contexts requires tailored approaches, as users interpret privacy risks differently based on their background and expertise [9, 32, 51]. To improve comprehension, textual descriptions should use clear, layman-friendly language and integrate relatable narratives to illustrate key concepts [15]. Combining text with visual aids enhances understanding and reduces cognitive load [50]. Contextual examples that situate DP settings in familiar scenarios may help users grasp abstract concepts more intuitively [9, 32].

Visual descriptions are more effective for end users. Non-textual methods, including visualizations, diagrams, and pictures, aim to simplify DP concepts, but several challenges persist. Ensuring accuracy without misinterpretation is difficult, as abstract DP principles do not always translate clearly into visual formats [32, 50]. Users often misinterpret visualizations without sufficient context, and engagement remains a challenge when designing intuitive and informative visuals [21]. Visual metaphors must balance simplicity and accuracy to avoid misleading interpretations of privacy guarantees [51]. Effective non-textual descriptions should integrate structured visual elements with clear textual explanations. Privacy-utility trade-off graphs, interactive models, and animations can help users explore different DP settings dynamically [30, 32]. Analogies and step-by-step diagrams can clarify DP mechanisms like data perturbation and aggregation [21, 51]. Iterative user testing should guide the refinement of these tools to ensure accessibility and accuracy across diverse audiences [50].

6.3   Software Tools

DP software tools are maturing quickly and have significantly improved in the past several years. Here, we summarize some of the main challenges still remaining for DP software tools, along with our recommendations.

Tools require DP expertise. Current tools—even those aimed at technical users without DP expertise—currently require significant DP expertise. In particular, API tools require technical expertise for effective use [35], and visual tools still demand an understanding of DP concepts [40]. We recommend that tool designers be conservative in their descriptions of expertise requirements, since deployments by technical users leveraging these tools may be incorrect. For tools that target technical users both with and without DP expertise, we recommend separate interfaces for different levels

of DP expertise, employing progressive disclosure to reveal complexity only when necessary, and high-quality documentation with examples for all user levels.

Tools require manual parameter-setting. Current tools require manual setting of DP parameters, which means deep understanding of DP is required to use them [29, 35, 40]. This can be incredibly challenging, even for DP experts: for example, even experts do not currently agree on how to set or interpret the privacy budget. Moreover, manual configuration of privacy parameters is complex and error-prone. We recommend adding automation for parameter-setting where possible, to reduce the burden on users. However, tool designers should be careful to avoid defaults that may violate DP, even if they improve automation.

Tools lack flexibility on deployment model and integration. Existing tools are designed for specific deployment models; they provide limited flexibility for altering the deployment model and are often unclear about the threat model they implement (for example, what level of trust in the data curator is required). We recommend that tool designers publish a clear description of each stakeholder's role and responsibilities, expected expertise, and expected level of trust. Moreover, current tools support only the central model of DP; we hope to see future tools support other deployment models.

Current tools provide limited integration with existing data processing pipelines and infrastructure. We recommend moving toward flexibility in integrating DP tools with non-DP data infrastructure, including modular design that allows experts to develop new extensions to integrate with existing data infrastructure.

Tools do not communicate utility. Current tools provide minimal support for understanding the utility of the DP results. Visual tools like DP Creator are beginning to investigate simple metrics for communicating expected utility to analysts, but API-based tools do not provide this support. We believe that tools should support negotiating the privacy-utility tradeoff among various stakeholders, by allowing the exploration of the impact on utility of various privacy parameter choices.

Some tools lack correctness checking. When using DP tools, technical users with and without DP expertise make mistakes that violate DP [35]. Automatic correctness checking can catch these mistakes, but reduces user satisfaction [35]. Many existing tools are research prototypes or aimed at users with DP expertise, and do not ensure correctness. However, since even DP experts routinely make mistakes when implementing DP deployments, we recommend that tools *always* perform correctness checks, and raise errors when DP might be violated. Automated parameter-setting can help avoid DP mistakes, but automated approaches should cautiously avoid defaults that could silently violate DP. To improve user satisfaction, tools should follow existing APIs when possible, and perform correctness checking transparently.

7   Open Questions

Our review identified many exciting avenues for future research that have not yet been addressed by prior work. In this section, we summarize open questions in all areas of DP usability.

Can end users understand DP? Understanding DP requires understanding complex mathematical concepts and threat models. Prior work has shown that communicating these aspects of a

guarantee using textual descriptions is difficult, but that simplified explanations tend to be more effective than detailed ones. Additional research is needed to determine if simple textual descriptions can sufficiently communicate important features of a DP guarantee.

Can visualizations help users understand DP? Visual tools seem to help individuals understand DP guarantees, including both the mathematical guarantee itself [32] and the surrounding deployment details [50] (e.g. LDP vs CDP). Research is needed to determine whether visual tools are useful for technical users, and what level of objective understanding can these tools conveyed.

How should we communicate with other stakeholders? The vast majority of research on communicating about DP has focused on end-users. Additional research should focus on stakeholders other than end users—including developers (who need assistance setting parameters), downstream data users (who need to understand utility) and policymakers (who need to understand the strength of a guarantee). Recruiting from these groups is challenging, but this research will be an important step toward increased DP adoption and transparency in DP deployments.

How should we communicate parameters beyond ϵ? The vast majority of studies with end-users have investigated methods for communicating about DP itself or the value of ϵ, and have excluded other important parameters like the unit of privacy. These parameters can be just as important as ϵ in determining the strength of the privacy guarantee, but little research have tackled them.

How should we communicate threat models? Prior work demonstrates that threat models in DP deployments are difficult for end-users to understand [50], and many DP deployments do not fully document their threat models. We recommend that tools specify the threat model(s) they implement, and that DP deployments document the threat model in a standardized way—but how to develop and communicate the threat model of a DP deployment effectively to various stakeholders remains an open question.

How should we communicate about utility? We currently lack clear, standardized metrics to evaluate the privacy-utility trade-off, especially in domain-specific ways. We believe future studies should examine the effectiveness of utility and accuracy metrics for navigating the privacy-utility tradeoff, and also explore the use of visual tools for this purpose.

How should we measure understanding? Most of the studies in our review attempted to measure whether participants understood some concept (e.g. the definition of DP or the strength of the guarantee), but different studies used very different approaches to measuring understanding. Some evaluated understanding directly, by asking participants to reason about possible outcomes in order to demonstrate understanding; others evaluated understanding indirectly, by asking participants about their willingness to share data in a given situation. How best to measure participants' understanding of DP concepts, and whether standardized measures of understanding could help enable comparisons between studies, remains an important open question.

How do legal requirements influence DP deployments? Legal or regulatory requirements, such as Europe's General Data Protection Regulation (GDPR) and California's Consumer Privacy Act (CCPA), are important motivating factors for the deployment of DP in industry. As described in Section 2, a significant amount

of work in the law community has sought to build connections between legal requirements and DP's formal guarantees. However, no studies have directly examined how legal requirements shape the use of DP and related techniques in practice. Garrido et al. [17] found that "*[privacy practitioners] measure privacy based on the fulfillment of data protection regulation*," suggesting that regulation is the *primary* motivation for DP deployments in industry—but the study did not examine this question in detail, and we are not aware of any other studies that examine this issue. We recommend future research targeting practitioners involved in legal questions of privacy, to investigate the connection between the law and decisions to implement DP, and to provide guidance on how to consider legal requirements in the decision-making process for a DP deployment.

Can users without DP expertise use DP tools effectively? Current research suggests that technical users without DP expertise struggle to use API-based tools [35], but less research has investigated if visual tools can help technical users [40]. New wizard-based visual tools that combine guidance for parameter-setting with automation may enable technical users without DP expertise to correctly implement DP.

Can automation ease parameter-setting? Technical users with or without DP expertise struggle to set DP-specific parameters (e.g. clipping parameters) and choose the most appropriate mechanisms [35, 40]. Some DP tools provide partial automation for setting some parameters, but no studies have investigated the impact of these features on usability.

Can verification tools be made accessible for all users? Software tools implement correctness checking in different ways, and existing work has shown that the specifics make a big difference in usability [35]. Future work should explore methods for DP verification tools that are usable by users regardless of their DP expertise.

8 Conclusion

In this paper, we systematically reviewed existing research on the usability of DP tools and DP communication. Our analysis highlights persistent difficulties in simplifying DP concepts, conveying the privacy-utility trade-off, inconsistent messaging and lack of standardized explanations and that current DP tools often require significant DP expertise making them less accessible to some technical users. To address these challenges, we recommend user-centered tool design, clearer communication strategies, and standardized messaging and educational resources to improve accessibility and adoption. In addition, establishing a usability of DP interest group can also help in recruiting representative samples for usability of DP studies.

Future research should refine communication methods for DP parameters, investigate the effectiveness of visual aids for different user groups, and explore the intersection of DP practices with legal requirements.

Acknowledgments

References

[1] Mayda Alrige. *Understanding natality data users' perceptions and preferences towards a query based noise addition approach based on differential privacy.* University of Maryland, Baltimore County, 2013.

[2] Andrea J Bingham and Patricia Witkowsky. Deductive and inductive approaches to qualitative data analysis. *Analyzing and interpreting qualitative data: After the interview*, 1:133–146, 2021.

[3] Brooke Bullek, Stephanie Garboski, Darakhshan J Mir, and Evan M Peck. Towards understanding differential privacy: When do people trust randomized response technique? In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3833–3837, 2017.

[4] Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I 38*, pages 375–403. Springer, 2019.

[5] Andrew Chin and Anne Klinefelter. Differential privacy as a response to the reidentification threat: The facebook advertiser case study. *NCL Rev.*, 90:1417, 2011.

[6] 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.

[7] Aloni Cohen and Kobbi Nissim. Towards formalizing the gdpr's notion of singling out. *Proceedings of the National Academy of Sciences*, 117(15):8344–8352, 2020.

[8] Rachel Cummings, Damien Desfontaines, David Evans, Roxana Geambasu, Yangsibo Huang, Matthew Jagielski, Peter Kairouz, Gautam Kamath, Sewoong Oh, Olga Ohrimenko, et al. Advancing differential privacy: Where we are now and future directions for real-world deployment. *arXiv preprint arXiv:2304.06929*, 2023.

[9] Rachel Cummings, Gabriel Kaptchuk, and Elissa M Redmiles. "I need a better description": An investigation into user expectations for differential privacy. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 3037–3052, 2021.

[10] 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.

[11] Damien Desfontaines and Balázs Pejó. Sok: Differential privacies. *Proceedings on Privacy Enhancing Technologies*, 2020(2):288–313, 2020.

[12] 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.

[13] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[14] Daniel Franzen, Claudia Müller-Birn, and Odette Wegwarth. Communicating the privacy-utility trade-off: Supporting informed data donation with privacy decision interfaces for differential privacy. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1):1–56, 2024.

[15] 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.

[16] Marco Gaboardi, James Honaker, Gary King, Jack Murtagh, Kobbi Nissim, Jonathan Ullman, and Salil Vadhan. Psi (Ψ): a private data sharing interface, 2018.

[17] Gonzalo Munilla Garrido, Xiaoyuan Liu, Florian Matthes, and Dawn Song. Lessons learned: Surveying the practicality of differential privacy in the industry. *arXiv preprint arXiv:2211.03898*, 2022.

[18] Julian Holzel. Differential privacy and the gdpr. *Eur. Data Prot. L. Rev.*, 5:184, 2019.

[19] Mark F. St. John, Grit Denker, Peeter Laud, Karsten Martiny, Alisa Pankova, and Dusko Pavlovic. Decision support for sharing data using differential privacy. In *2021 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pages 26–35, 2021.

[20] Bailey Kacsmar, Vasisht Duddu, Kyle Tilbury, Blase Ur, and Florian Kerschbaum. Comprehension from chaos: Towards informed consent for private computation. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, CCS '23, page 210–224, New York, NY, USA, 2023. Association for Computing Machinery.

[21] Farzaneh Karegar, Ala Sarah Alaqra, and Simone Fischer-Hübner. Exploring {User-Suitable} metaphors for differentially private data analyses. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 175–193, 2022.

[22] Patrick Kühtreiber, Viktoriya Pak, and Delphine Reinhardt. Replication: the effect of differential privacy communication on german users' comprehension and data sharing attitudes. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 117–134, 2022.

[23] Marcus Land. Supporting diverse stakeholders to make informed decisions about the use of differential privacy with a web-based e-learning application.

[24] Ho Cheung Brian Lee, Sulin Ba, Xinxin Li, and Jan Stallaert. Salience bias in crowdsourcing contests. *Information Systems Research*, 29(2):401–418, 2018.

[25] Sascha Löbner, Sebastian Pape, and Vanessa Bracamonte. User acceptance criteria for privacy preserving machine learning techniques. In *Proceedings of the 18th International Conference on Availability, Reliability and Security*, pages 1–8, 2023.

[26] Pieter Maene, Johannes Götzfried, Ruan De Clercq, Tilo Müller, Felix Freiling, and Ingrid Verbauwhede. Hardware-based trusted computing architectures for isolation and attestation. *IEEE Transactions on Computers*, 67(3):361–374, 2017.

[27] Àlex Miranda-Pascual, Patricia Guerra-Balboa, Javier Parra-Arnau, Jordi Forné, and Thorsten Strufe. Sok: differentially private publication of trajectory data. *Proceedings on Privacy Enhancing Technologies*, 2023.

[28] Tamara T Mueller, Dmitrii Usynin, Johannes C Paetzold, Daniel Rueckert, and Georgios Kaissis. Sok: Differential privacy on graph-structured data. *arXiv preprint arXiv:2203.09205*, 2022.

[29] Jack Murtagh, Kathryn Taylor, George Kellaris, and Salil Vadhan. Usable differential privacy: A case study with psi. *arXiv preprint arXiv:1809.04103*, 2018.

[30] Priyanka Nanayakkara, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers. Visualizing privacy-utility trade-offs in differentially private data releases. *arXiv preprint arXiv:2201.05964*, 2022.

[31] Priyanka Nanayakkara, Hyeok Kim, Yifan Wu, Ali Sarvghad, Narges Mahyar, Gerome Miklau, and Jessica Hullman. Measure-observe-remeasure: An interactive paradigm for differentially-private exploratory analysis. *arXiv preprint arXiv:2406.01964*, 2024.

[32] 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.

[33] Joseph P Near, Xi He, et al. Differential privacy for databases. *Foundations and Trends® in Databases*, 11(2):109–225, 2021.

[34] Boel Nelson and Jenni Reuben. Sok: Chasing accuracy and privacy, and catching both in differentially private histogram publication. *arXiv preprint arXiv:1910.14028*, 2019.

[35] Ivoline C. Ngong, Brad Stenger, Joseph P. Near, and Yuanyuan Feng. Evaluating the usability of differential privacy tools with data practitioners, 2024.

[36] Jakob Nielsen and Jonathan Levy. Measuring usability: preference vs. performance. *Communications of the ACM*, 37(4):66–75, 1994.

[37] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.

[38] Kobbi Nissim and Alexandra Wood. Is privacy privacy? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128):20170358, 2018.

[39] Martin Pettai and Peeter Laud. Combining differential privacy and secure multiparty computation. In *Proceedings of the 31st annual computer security applications conference*, pages 421–430, 2015.

[40] Jayshree Sarathy, Sophia Song, Audrey Haque, Tania Schlatter, and Salil 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.

[41] Mary Anne Smart, Dhruv Sood, and Kristen Vaccaro. Understanding risks of privacy theater with differential privacy. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–24, 2022.

[42] Ana Sokolovska and Ljupco Kocarev. Integrating technical and legal concepts of privacy. *Ieee Access*, 6:26543–26557, 2018.

[43] Daniel J Solove. A taxonomy of privacy. *U. Pa. l. Rev.*, 154:477, 2005.

[44] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. Privacy-aware eye tracking using differential privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pages 1–9, 2019.

[45] Christian Sturm, Alice Oh, Sebastian Linxen, Jose Abdelnour Nocera, Susan Dray, and Katharina Reinecke. How weird is hci? extending hci principles to other countries and cultures. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2425–2428, 2015.

[46] Michael Carl Tschantz, Shayak Sen, and Anupam Datta. Sok: Differential privacy as a causal property. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 354–371. IEEE, 2020.

[47] André Calero Valdez and Martina Ziefle. The users' perspective on the privacy-utility trade-offs in health recommender systems. *International Journal of Human-Computer Studies*, 121:108–121, 2019.

[48] Zikai Alex Wen, Jingyu Jia, Hongyang Yan, Yaxing Yao, Zheli Liu, and Changyu Dong. The influence of explanation designs on user understanding differential privacy and making data-sharing decision. *Information Sciences*, 642:118799, 2023.

[49] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. Towards effective differential privacy communication for users' data sharing decision and comprehension. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 392–410. IEEE, 2020.

[50] Aiping Xiong, Chuhao Wu, Tianhao Wang, Robert W Proctor, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Using illustrations to communicate differential privacy trust models: an investigation of users' comprehension, perception, and data sharing decision. *arXiv preprint arXiv:2202.10014*, 2022.

[51] Aiping Xiong, Chuhao Wu, Tianhao Wang, Robert W Proctor, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Exploring use of explanative illustrations to communicate differential privacy models. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 67, pages 226–232. SAGE Publications Sage CA: Los Angeles, CA, 2023.