Developing Textual Descriptions of PETs for Ad Tracking and Analytics

Lu Xian Song Mi Lee-Kan Jane Im Florian Schaub
University of Michigan

1 Introduction

Various organizations and companies are increasingly incorporating Privacy Enhancing Technologies (PETs) in their services and products. Platforms are moving away from third-party cookies to more privacy-friendly methods [5, 10]. Apple and Google use local differential privacy *(LDP)* to analyze browser data in Safari [9, 30], Chrome [11], and across other services [2, 6, 15], enhancing user privacy by modifying user data on individual devices before transmission to central servers. Federated learning *(FL)* trains machine learning models on decentralized devices and updates the global model without data centralization [18, 23, 24, 28]. FL has expanded to other applications, including IBM's data analytics services [19]. FL is often paired with LDP *(FL+LDP)* to enhance privacy in model training, such as in Google's Smart Text Selection and Apple's Siri personalization [14, 17]. Additionally, Google has introduced Topics *(GT)* [13, 29] as part of its Privacy Sandbox initiative, which analyzes user browsing data on-device to identify users' interests. This allows advertisers to target ads based on a subset of a user's interests without individual tracking across websites [29].

Given the potential impact of PETs on user privacy regarding online ads and analytics, it is crucial to study how users understand and respond to these technologies. Prior research primarily conducted in health settings shows that users' understanding of PET descriptions significantly affects their willingness to share data [4, 8, 16, 26, 27]. Studies have focused predominantly on effective communication methods for differential privacy (DP) and LDP within the health domain [12, 20, 25, 31, 32]. While PETs can be presented through various formats like images and interactive tools, textual descriptions are most commonly used by PET vendors, implementers (e.g., see Table 2 in Appendix 5.2), and in privacy policy documents. A promising approach for textually describing PETs, developed by Xiong et al. [31], combines an explanation of how a PET works on a process level with a statement on its privacy implications, such as reducing the risk from server-side data breaches. They developed textual descriptions of DP and LDP for health apps and found that mentioning potential data breaches enhances users' objective comprehension and willingness to share health data. However, due to the context-dependent nature of privacy and information sensitivity [1], it is unclear if these findings in the health domain apply to other domains, such as ad tracking and analytics. Furthermore, most research focused on user-centric descriptions of DP/LDP, while descriptions of other PETs, like FL, FL+LDP, and GT, remain underexplored despite their growing industry use. To address this gap, we investigate the following research questions:

RQ1. How does user comprehension differ between our refined LDP description and the minimally modified description from Xiong et al. in ad tracking and analytics?

RQ2. How does including an implications statement, found effective in Xiong et al.'s health context [31], affect user comprehension of the LDP description in ad tracking and analytics?

RQ3. How does including an implications statement affect user comprehension of the descriptions of FL, FL+LDP, and GT in ad tracking and analytics?

RQ4. What aspects of the text describing PETs in ad tracking and analytics do users comprehend accurately and inaccurately?

2 Methodology

PET description development. We developed a total of 10 textual descriptions for different PETs in the context of ad tracking. The descriptions of LDP were adapted from Xiong et al. [31] from the health sector to ad tracking. For FL,

USENIX Symposium on Usable Privacy and Security (SOUPS) 2024.
August 11–13, 2024, Philadelphia, PA, United States.

FL+LDP, and GT, which lack established descriptions, we initially created drafts based on industry documents and academic research. Building on prior research that recommends adding an implications statement to improve comprehension [21, 31], we produced two versions of descriptions for each PET: one describing the PET's approach to data processing and another enhanced by a privacy implications statement ("impl"). The drafts were refined through think-aloud interviews (n=5), two expert reviews, three rounds of pilot testing with a Prolific panel (n=6; 11; 5), and feedback from an industry privacy engineer. Table 1 in Appendix 5.1 shows a full list of finalized descriptions evaluated in the survey.

Survey experiment. We conducted an online survey experiment using Qualtrics. This research obtained an exemption from the university's Institutional Review Board. A total of 357 participants were recruited through Prolific. Participants were predominantly white, with women constituting a slight majority at 51%. Our sample, ranging widely in age with a median of 35, had marginally higher education levels compared to census data, making it broadly representative of the general population in the US. After providing consent, they were randomly assigned to one of 10 experimental conditions. Each condition featured a PET description: a control group, LDP-Xiong, and descriptions both with and without "impl" for LDP, FL, FL+LDP, and GT. The full survey instrument is provided in Appendix 5.3.

Analysis and coding. To measure the effect of the implications statement of PETs descriptions (**RQ1-3**), we conducted non-parametric Mann–Whitney tests to compare pairs of experimental conditions for each PET across three quantitative metrics: confidence in platform use, objective comprehension score, and subjective confidence in comprehension. Qualitative responses were analyzed to understand users' perception of the PETs descriptions we provided (**RQ4**), with all responses singled-coded by the first author using inductive coding [22].

3 Key Findings

We present novel, pairwise comparisons of the effects of an implications statement versus descriptions without it for several increasingly popular PETs (see Figure 1 in Appendix 5.4). We find that **the process- and implications-focused approach to describing PETs is effective not only in health settings but also in other contexts, such as ad tracking and analytics**. Prior work has found positive effects of implication statements for DP/LDP in the health context [21, 31]; yet, **we could not replicate a significant impact for LDP and other PETs in ad tracking and analytics.** Incorporating an implications statement into the descriptions did not measurably increase user understanding in terms of objective comprehension score, subjective comprehension, and confidence in platform use. This result may suggest that implications statements that we adapted from the health context are less effective in ad tracking and analytics. More research is needed to craft implications statements specifically suited for behavioral data and further examine their phrasings and salience, such as the location within the description, to improve usefulness.

Our analysis of users' mental models reveals varied understanding concerning PETs: For FL and FL+LDP, some participants accurately comprehended the role of machine learning models in discerning user interests and the concept of model sharing. Yet, many struggled with the concept of on-device model training and the process of merging models to derive general behavioral patterns. Additionally, the term "noise" in FL+LDP descriptions, similar to challenges with LDP, was difficult to understand for participants. For GT, while some participants clearly understood the on-device processing of data, others overlooked this feature and instead expressed concerns about the potential data tracking and collection practices. These observations underscore the complexities in effectively communicating the functionalities and role of machine learning and "models" in PETs descriptions.

4 Implications for Describing PETs to Users

We show what we can learn to improve descriptions of the PETs we studied:

Simplify technical terminology. We find that the use of technical terms like "noise" in LDP descriptions, "random modification" in FL, and "machine learning models" in both FL and GT can confuse users. Despite avoiding statistical jargon, the term "noise" still baffles due to its everyday usage. Furthermore, the mention of the use of machine learning to infer interests in PET descriptions inadvertently shifted some users' focus towards concerns about data tracking and recording by machine learning models.

Help users pinpoint the source of privacy protection. Our participants often find it challenging to identify the exact mechanisms protecting their privacy, whether through data modification and sharing the modified data (LDP), sharing machine learning models without transferring actual data (FL), models trained on modified data (FL+LDP), or sharing a selection of inferred topics that may include some random ones (GT). A clearer explanation in the implications statement would help users understand these mechanisms.

Provide more specificity about user data. Common confusions and misconceptions across PETs we observed in our participants' survey responses were often associated with their desire to understand the nuances of data collection and processing—what data is collected, where it is analyzed, and by whom. This understanding is crucial because the protective mechanisms of certain PETs, such as FL and FL+LDP, hinge on the fact that data processing occurs on the user's device rather than a centralized system accessible to the organization. Clearly stating that data remains on the user's device when not shared with the organization can help clarify these privacy measures.

References

[1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[2] Apple. Apple differential privacy technical overview, n.a. Accessed on 14 February 2024.

[3] Engineering at Meta. Applying federated learning to protect data on mobile devices, 2022. Accessed on 14 February 2024.

[4] 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.

[5] Matt Burgess. Google's cookie ban and floc, explained, 2021. Accessed on 14 February 2024.

[6] Alisa Chang and Pritish Kamath. Practical differentially private clustering, 2021. Accessed on 14 February 2024.

[7] Google Chrome. More about ad topics, 2024. Accessed on 14 February 2024.

[8] 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.

[9] Apple Differential Privacy Team. Learning with privacy at scale, 2017. Accessed on 14 February 2024.

[10] David Eliot and David Murakami Wood. Culling the floc: Market forces, regulatory regimes and google's (mis) steps on the path away from targeted advertising 1. *Information Polity*, 27(2):259–274, 2022.

[11] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067, 2014.

[12] 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.

[13] Vinay Goel. Get to know the new topics api for privacy sandbox, 2022. Accessed on 14 February 2024.

[14] Google. Distributed differential privacy for federated learning. Google Research, 2023. Accessed on 14 February 2024.

[15] Miguel Guevara. Expanding access to differential privacy to create a safer online ecosystem, 2022. Accessed on 14 February 2024.

[16] Nadine Guhr, Oliver Werth, Philip Peter Hermann Blacha, and Michael H Breitner. Privacy concerns in the smart home context. *SN Applied Sciences*, 2:1–12, 2020.

[17] Karen Hao. How apple personalizes siri without hoovering up your data. MIT Technology Review, 2019. Accessed on 14 February 2024.

[18] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018.

[19] IBM. Ibm launches new watson capabilities to help businesses build trustworthy ai, 2021. Accessed on 14 February 2024.

[20] Farzaneh Karegar and Simone Fischer-Hübner. Vision: A noisy picture or a picker wheel to spin? exploring suitable metaphors for differentially private data analyses. In *Proceedings of the 2021 European Symposium on Usable Security*, pages 29–35, 2021.

[21] 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.

[22] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. Morgan Kaufmann, 2017.

[23] Brendan McMahan and Daniel Ramage. Federated learning: Collaborative machine learning without centralized training data, 2017. Accessed on 14 February 2024.

[24] Eric Miraglia. Privacy that works for everyone, 2019. Accessed on 14 February 2024.

[25] Priyanka Nanayakkara, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers. Visualizing Privacy-Utility Trade-Offs in Differentially Private Data Releases, January 2022. arXiv:2201.05964 [cs].

[26] 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.

[27] Ismini Psychoula, Deepika Singh, Liming Chen, Feng Chen, Andreas Holzinger, and Huansheng Ning. Users' privacy concerns in iot based applications. In *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*, pages 1887–1894. IEEE, 2018.

[28] Swaroop Ramaswamy, Rajiv Mathews, Kanishka Rao, and Françoise Beaufays. Federated learning for emoji prediction in a mobile keyboard. *arXiv preprint arXiv:1906.04329*, 2019.

[29] The Privacy Sandbox. Topics api: Relevant ads without cookies, 2023. Accessed on 14 February 2024.

[30] ADP Team et al. Learning with privacy at scale. *Apple Mach. Learn. J*, 1(8):1–25, 2017.

[31] 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, May 2020. ISSN: 2375-1207.

[32] 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.

5 Appendix

5.1 PETs Descriptions

We list the 10 descriptions shown to participants in the survey experiment in Table 1.

5.2 In-the-wild descriptions of other PETs

Table 2 provides examples of in-the-wild descriptions of FL, FL+LDP, and GT provided by major tech companies. The sources include the companies' research webpages, announcements, and product settings. Among these, the description of Topics was incorporated by Google in the Ad Topics settings of Chrome [7] after we designed the survey and collected our data.

5.3 Survey Instrument

We use the *control* condition as an example to list the survey instructions and questions. We provide additional information and mark the variations in survey questions shown to respondents assigned to other conditions in italicized texts.

Survey instruction.

In this survey, we will ask you a series of questions about a hypothetical scenario. Please do your best to imagine yourself in this scenario and answer the questions.

Scenario description.

Imagine that you came across the following description of a social platform.

The platform makes revenue by showing users personalized ads via inferring users' interests from their online activity tracked on the platform and other businesses' websites/apps. To protect your information, the organization stores all of your behavioral data for targeting ads (e.g., your interaction with the platform and with other apps/websites) securely on their servers. *Please see Table 1 for the PET description under other experiment conditions.*

Imagine you are trying to decide whether you would like to use this platform.

Confidence in platform use question.

- How confident are you about deciding whether to use this platform?
  - Very confident
  - Confident
  - Moderately confident
  - Slightly confident
  - Not at all confident

Perceived protection of user data.

- How would you explain to other people how the platform protects users' data? Please write at least two clear sentences. [Space for open-ended responses was provided.]

Objective comprehension questions.

For each of the following statements, please indicate if you expect the following to be true or false if you would use the platform described above.

- An employee working for the platform, such as a data analyst, could be able to see my exact behavioral data (e.g., my interaction with the platform and with other apps/websites).
  - True
  - False
  - I don't know
- A criminal or foreign government that hacks the platform could learn my behavioral data (e.g., my interaction with the platform and with other apps/websites).
  - True
  - False

Table 1: PETs descriptions in the survey experiment, with implications statements in square brackets.

Condition	Description
Control	To protect your information, the organization stores all of your behavioral data for targeting ads (e.g., your interaction with the platform and with other apps/websites) securely on their servers.
LDP-Xiong	To respect your personal information privacy and ensure best user experience, the behavioral data (e.g., your interaction with the platform and with other apps/websites) shared with the company will be processed via an additional privacy technique. That is, your behavioral data will be randomly modified before it is sent to the company. Since the company stores only the modified version of your personal information, your privacy is protected even if the company's database is compromised.
LDP[-impl]	To protect your information, the organization adds noise to your behavioral data (e.g., your interaction with the platform and with other apps/websites) before being sent to the organization for targeting ads. This means that your data is randomly modified, so that some of your actual data is used whereas some of it is random and not representative of your behavior. Your exact behavioral data is never sent to the organization, instead a subset of your noisy data is randomly selected and sent. The organization can still infer patterns from the noisy data across a large number of users. [This way, the organization still learns aggregated interests across users but not your exact behavior, which protects your privacy against the organization's employees or if the organization's database is compromised.]
FL[-impl]	To protect your information, the organization uses machine learning on your device to infer interests from your behavioral data (e.g., your interaction with the platform and with other apps/websites) for targeting ads. Your exact behavioral data is never sent to the organization and only a machine learning model representing your inferred interests will be sent. Then, to infer patterns across a large number of users, your model is merged with other users' models. [This way, the organization still learns your interests but not your exact behavior, which protects your privacy against the organization's employees or if the organization's database is compromised.]
FL+LDP[-impl]	To protect your information, the organization uses machine learning on your device to infer interests from your behavioral data (e.g., your interaction with the platform and with other apps/websites) for targeting ads. Noise will be added to your behavioral data so that it is randomly modified before being used for training a machine learning model representing your inferred interests. This means that, for training the model, some of your actual data is used whereas some of it is random and not representative of your behavior. Your exact behavioral data is never sent to the organization and only the model representing your inferred interests will be sent. Then, to infer patterns across a large number of users, your model is merged with other users' models. [This way, the organization still learns aggregated interests across its users but not your exact behavior, which protects your privacy against the organization's employees or if the organization's database is compromised.]
GT[-impl]	To protect your information, the organization uses machine learning on your device to infer interests from your behavioral data (e.g., your interaction with the platform and with other apps/websites) for targeting ads. This means that the technology records inferred topics you may be interested in from your behavioral data only on your device. Your exact behavioral data is never sent to the organization, instead from your top topics of the last week, a small number are randomly selected and sent; there is also a small chance a random topic will be selected instead of one of yours. [This way, the organization still learns some of your interests but not your exact behavior, which protects your privacy against the organization's employees or if the organization's database is compromised.]

Table 2: Examples of industry descriptions of FL, FL+LDP, and GT. The last description of Topics [7] (last row) in the table was incorporated in the Ad Topics settings of Chrome by Google after our data collection was completed. We list the description of Topics from Chrome Version 122.0.6261.69 [7] in the table.

PET	Description
FL [23]	Federated Learning enables mobile phones to collaboratively learn a shared prediction model while keeping all the training data on device, decoupling the ability to do machine learning from the need to store the data in the cloud. This goes beyond the use of local models that make predictions on mobile devices (like the Mobile Vision API and On-Device Smart Reply) by bringing model training to the device as well. It works like this: your device downloads the current model, improves it by learning from data on your phone, and then summarizes the changes as a small focused update. Only this update to the model is sent to the cloud, using encrypted communication, where it is immediately averaged with other user updates to improve the shared model. All the training data remains on your device, and no individual updates are stored in the cloud.
FL [9]	Differential privacy provides a mathematically rigorous definition of privacy and is one of the strongest guarantees of privacy available. It is rooted in the idea that carefully calibrated noise can mask a user's data. When many people submit data, the noise that has been added averages out and meaningful information emerges.
FL+LDP [3]	Federated learning with differential privacy (FL-DP) is one of the latest privacy-enhancing technologies being evaluated at Meta as we constantly work to enhance user privacy and further safeguard users' data in the products we design, build, and maintain. FL-DP enhances privacy in two important ways: 1. It allows machine learning (ML) models to be trained in a distributed way so that users' data remains on their mobile devices. 2. It adds noise to reduce the risk of an ML model memorizing user data.
GT [13]	With Topics, your browser determines a handful of topics, like "Fitness" or "Travel & Transportation," that represent your top interests for that week based on your browsing history. Topics are kept for only three weeks and old topics are deleted. Topics are selected entirely on your device without involving any external servers, including Google servers. When you visit a participating site, Topics picks just three topics, one topic from each of the past three weeks, to share with the site and its advertising partners. Topics enables browsers to give you meaningful transparency and control over this data, and in Chrome, we're building user controls that let you see the topics, remove any you don't like or disable the feature completely.
GT [7]	Chrome notes topics of interest based on your browsing history from the last few weeks. Later, a site you visit can ask Chrome for your topics to personalize the ads you see. Chrome shares up to 3 topics while protecting your browsing history and identity. Chrome auto-deletes topics that are older than 4 weeks. As you keep browsing, a topic might reappear on the list. Or you can block topics you don't want Chrome to share with sites. Learn more about managing your ad privacy in Chrome.

- I don't know

- A law enforcement organization could access my behavioral data (e.g., my interaction with the platform and with other apps/websites) with a court order requesting this data from the company.
  - True
  - False
  - I don't know
- Graphs or informational charts created using information given to the platform could reveal my behavioral data (e.g., my interaction with the platform and with other apps/websites).
  - True
  - False
  - I don't know
- Data that the platform shares with its partner organizations could reveal my behavioral data (e.g., my interaction with the platform and with other apps/websites).
  - True
  - False
  - I don't know

Interpretation of PET description segments.

- In your own words, describe what "behavioral data (e.g., your interaction with the platform and with other apps/websites)" in the above description means. Please write at least two clear sentences.
- In your own words, describe what "To protect your information, the organization stores all of your behavioral data for targeting ads (e.g., your interaction with the platform and with other apps/websites) securely on their servers" in the above description means. Please write at least two clear sentences.
  This question is different under different experiment conditions; this question is specific to the PETs description provided in the scenario description section of the survey (see Table 3).

Subjective comprehension questions.

- How confident are you in your understanding of the privacy technology used by the platform's company?
  - Very confident
  - Confident
  - Moderately confident
  - Slightly confident
  - Not at all confident
- You indicated that the description of privacy technology used by the platform was not easy to understand. Please indicate which words or sentences were hard to understand, or you wished you had more details about. [This question is asked to only respondents who gave a rating less than 4.]

Prior PET familiarity

- Have you ever heard of the following technologies? (select all that apply)
  - Differential privacy
  - End-to-end encryption
  - Secure multi-party computation
  - Deliquescent security
  - Federated learning
  - Topics
  - FLoC (Federated Learning of Cohorts)
  - None of the above

PET identification

- Which of these technologies do you think was described in the scenario?
  - Differential privacy
  - End-to-end encryption
  - Secure multi-party computation
  - Deliquescent security
  - Federated learning
  - Topics
  - FLoC (Federated Learning of Cohorts)
  - None of the above

Demographic questions

- In what year were you born? (four digits please)
- What is your gender?
  - Man
  - Woman
  - Non-binary
  - Prefer to self-describe
  - Prefer not to answer
- Please specify your race/ethnicity (select all that apply)
  - Hispanic, Latino, or Spanish
  - Black or African American
  - White
  - American Indian or Alaska Native
  - Asian, Native Hawaiian, or Pacific Islander
  - Prefer to self-describe
  - Prefer not to answer
- What is the highest level of school you have completed or the highest degree you have received?
  - Less than high school degree

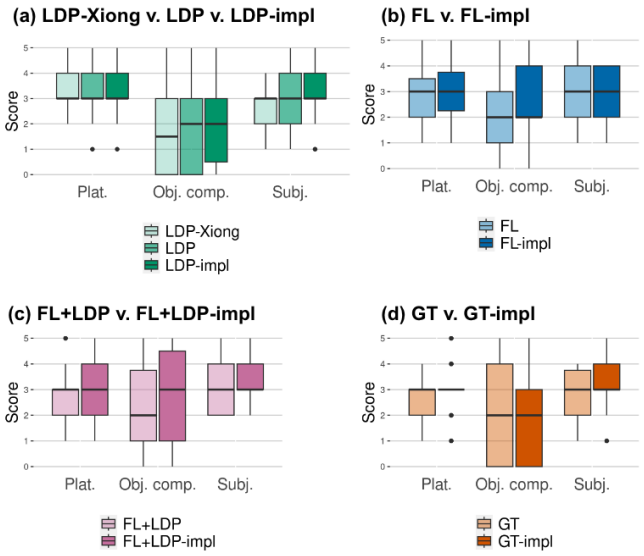

Figure 1: Grouped boxplots compare confidence in platform use (Plat.), objective comprehension (Obj.), and subjective comprehension (Subj.) across experimental condition pairs in RQs 1–3. Apart from the confidence in platform use comparison in RQ1, none of the other comparisons reach significance at the 0.05 level in the Mann-Whitney tests.

- High school graduate (high school diploma or equivalent including GED)
- Some college but no degree
- Associate's degree
- Bachelor's degree
- Advanced degree (e.g., Master's, doctorat)
- Prefer not to answer

- Which of the following best describes your educational background or job field?
  - I DO NOT have an education in, nor do I work in, the field of computer science, computer engineering or IT.
  - I have an education in, nor do I work in, the field of computer science, computer engineering or IT.

5.4 Quantitative Analysis Results

Quantitative analysis results for RQ1-3 are shown in Figure 1.

Table 3: Open-ended questions about specific segments of the PETs descriptions, which vary across experimental conditions.

Condition	**Survey question**
Control	In your own words, describe what "To protect your information, the organization stores all of your behavioral data for targeting ads (e.g., your interaction with the platform and with other apps/websites) securely on their servers" in the above description means. Please write at least two clear sentences.
LDP-Xiong	In your own words, describe what "... the behavioral data (e.g., your interaction with the platform and with other apps/websites) shared with the company will be processed via an additional privacy technique. That is, your behavioral data will be randomly modified before it is sent to the company" in the above description means. Please write at least two clear sentences.
LDP	(1) In your own words, describe what "To protect your information, the organization adds noise to your behavioral data (e.g., your interaction with the platform and with other apps/websites) before being sent to the organization for targeting ads. This means that your data is randomly modified, so that some of your actual data is used whereas some of it is random and not representative of your behavior" in the above description means. Please write at least two clear sentences. (2) In your own words, describe what "Your exact behavioral data is never sent to the organization, instead a subset of your noisy data is randomly selected and sent. The organization can still infer patterns from the noisy data across a large number of users" in the above description means. Please write at least two clear sentences.
FL	In your own words, describe what "Your exact behavioral data is never sent to the organization and only the model representing your inferred interests will be sent. Then, to infer patterns across a large number of users, your model is merged with other users' models" in the above description means. Please write at least two clear sentences.
FL+LDP	(1) In your own words, describe what "Noise will be added to your behavioral data so that it is randomly modified before being used for training a model representing your inferred interests. This means that, for training the model, some of your actual data is used whereas some of it is random and not representative of your behavior" in the above description means. Please write at least two clear sentences. (2) In your own words, describe what "Your exact behavioral data is never sent to the organization and only the model representing your inferred interests will be sent. Then, to infer patterns across a large number of users, your model is merged with other users' models" in the above description means. Please write at least two clear sentences.
GT	In your own words, describe what "... the technology records inferred topics from your behavioral data only on your device. Your exact behavioral data is never sent to the organization, instead from your top topics of the last week, a small number are randomly selected and sent; there is also a small chance a random topic will be selected instead of one of yours" Please write at least two clear sentences.