Culture, Identity, and App Genres in Mobile Application Permission Decisions

CS 6501 HCI Project Report

Andrew Balch
xxv2zh@virginia.edu
University of Virginia
Charlottesville, USA

Sabit Ahmed
bcw3zj@virginia.edu
University of Virginia
Charlottesville, USA

Archit Uniyal
deu9yh@virginia.edu
University of Virginia
Charlottesville, USA

Shaina Kumar
mzm2cj@virginia.edu
University of Virginia
Charlottesville, USA

Szu Yuan Cheng
skf7kn@virginia.edu
University of Virginia
Charlottesville, USA

Abstract

Privacy and security are crucial functions in most user applications, especially smartphone apps. An individual's privacy needs, concerns, and preferences can vary based on demographic factors as well as the type of application. Precise understanding of how age, race, ethnicity, nationality, and app genre plays a role in allowing those permissions is important for inclusive and diversity-oriented application and system designs. In this study, we conducted a survey to understand how demographic factors and app genre affect the data permissions users grant to an app. Our survey design includes a series of hypothetical scenarios involving different mobile app genres, where participants were instructed to indicate the permissions they would allow, to what degree they would allow them, and how they expected the app to use this information. We conducted a preliminary study to identify the most effective scenario description approach, which we later employed in our main survey design. Our main survey collected users' demographic information, privacy and security preferences across different app genres, and their expectation on the collected app data usage. The demographics of our sample was diverse, with a total number of 81 participants in our study. We found race and ethnicity had significant effects on permission-granting, while age, gender and nationality had not. We also found that app genres had a pivotal role in privacy decisions of participants. Last, we discovered an interaction between nationality and app genre for permissions granted to an app. The findings of our study hold implications for more inclusive and privacy-focused mobile application designs.

ACM Reference Format:
Andrew Balch, Sabit Ahmed, Archit Uniyal, Shaina Kumar, and Szu Yuan Cheng. 2024. Culture, Identity, and App Genres in Mobile Application Permission Decisions: CS 6501 HCI Project Report. In . ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Privacy needs, risks, and behaviors differ significantly across various social groups, and these differences play a critical role in how individuals interact with technology and make privacy-related decisions. Previous research has shown that people from different demographics may view risks, trust, and privacy concerns through different lenses, influenced by their backgrounds and social experiences. Understanding these variations is crucial, particularly in the context of permissions granted to cookies, app tracking, and other privacy-sensitive technologies. In today's world, privacy is no longer a universal concept but is shaped by numerous personal and social factors.

Moreover, the field of Human-Computer Interaction (HCI) has increasingly emphasized the importance of recognizing user demographic differences in design. While there has been substantial research on what design patterns influence the acceptance or rejection of privacy options, such as personality traits or interface usability, relatively few studies have examined the cultural and national influences on privacy preferences. By exploring how gender, race/ethnicity, and nationality affect privacy permissions, this research aims to fill that gap, providing valuable insights for developers and designers seeking to create more inclusive, respectful, and culturally aware systems. Therefore, we will focus on the following research questions:

- *RQ1:* How do one's demographic factors (race/ethnicity, nationality, gender identity, and age) influence users in allowing app permissions and the extent to which they allow these permissions?
- *RQ2:* How does the app genre influence the permissions users choose to allow and to what degree they allow those permissions?
- *RQ3:* Is there any relationship between the demographic and the app genre in terms of app permission settings?

2 Related Work

Several relevant papers illustrate the scope of this problem. Park argues that one's ability to exercise privacy is partially determined by social standing, extending inequalities into the digital world [12]. It is therefore critical to understand how individual social backgrounds (culture and demographics) interact with privacy. They

present evidence from the Health Information National Trends survey, finding that "age, education, and income had significant impacts on one's privacy confidence", and confidence itself positively influenced digital participation [12]. These results demonstrate that studies of privacy are inseparable from studies of inequity and social group differences, and vice versa. Motivated by the lack of literature investigating how race/ethnicity and socioeconomic status (SES) inform privacy, Wang and Metzger sought to understand how online privacy norms "may deter people of color from active participation" [14]. Their approach involved an online survey of privacy concerns and privacy management behaviors on social media. The responses showed that people of color exhibited more privacy concerns and privacy management behavior, indicating that these groups more actively seek online privacy protection. SES was also indicative of such behavior, but rarely interacted with race/ethnicity. Therefore, the authors concluded that marginalized groups have varied interpretations of and approaches to online privacy protection, and further work is necessary to study how these groups manage their privacy specifically with respect to avoidance strategies. Since women are often seen as a disadvantaged user segment in technology, gender has also been explored as a factor impacting privacy confidence and behavior [11]. Again, an online survey was used to gather data on online privacy management and release, confidence in privacy protection, and general internet use. The author concluded that there were gendered differences in privacy protection behaviors and confidence. Importantly, 77% of study participants were non-hispanic white, failing to account for the dimension of race/ethnicity as argued by [14]. Our study builds upon these existing works by accounting for both race/ethnicity and gender in the context of privacy. At the same time, we expand to consider non-binary gender identities as well as the impact of age and nationality.

Building on recent work by Hutton [6], who used surveys to explore the relationship between app permissions, privacy concerns, and personality traits, it was found that individuals with greater privacy concerns tend to make more calculated decisions, though no significant correlation between privacy and personality was established. In our study, we adopt a similar survey-based approach to examine the correlation between app permissions and demographic factors such as gender, race/ethnicity, and nationality. Additionally, Hamed [4] conducted a privacy risk assessment to evaluate users' awareness of mobile app permissions, proposing metrics like PrivacyScoreapp, which measures both the number of permissions granted to all applications on a device and those requested by specific apps like Facebook, Yahoo, and Twitter, assessing their impact on user privacy. Another relevant study by Marsch [9] utilized online questionnaires to analyze how app permissions relate to interdependent privacy. This research revealed that users often overlook the privacy of others when granting permissions to the applications. Lavranou et al. [8] investigated a comprehensive list of permissions required by the seven most popular mobile applications to offer distinct functionalities. They found that though some applications had restricted access to some sensitive permissions and offered end-to-end encryption, all applications share common concerns regarding potential audio and video surveillance, audio manipulation, and data privacy implications. Based on these findings, they developed a user education platform so that users can better scrutinize permission settings from the vast number of permissions. In contrast to this study, we aim to analyze the perceived privacy or security risks and benefits of different permission settings that varied by application genres. To our knowledge, it is the first study that seeks to extend this analysis by studying how these effects vary across demographics through targeted surveys.

3 Methods

3.1 Preliminary Study

To assess our research questions, we designed a survey which was carefully curated to understand preferences in mobile app permissions across demographics and app genre. First, we conducted a preliminary study to identify the most effective scenario description technique (*vague*, *contextual*, or *vignette* [9]) for the main survey. This was crucial to understand and navigate the potential response bias induced by describing perceived benefits or motivation in the *contextual* and *vignette* techniques compared to the *vague* question framing. In the *vague* model, we asked the participants to provide permissions for a new app (e.g. a social media app) without any context or specific information. In the *contextual* scenario, we provided a motivation for why the respondent is using the app (e.g. "Your friends are asking you to join a new social media app so you can stay connected over the summer"). For the *vignette* model, we presented a scenario using a third party, where a graduate student, rather than the respondent, wants to post in a new social media app for the first time. Given this information, we asked participants which mobile app permissions the graduate student should provide, accounting for common degrees of access requested by mobile apps (Allow always, Allow only while using the app, Allow once/Ask every time, or Never allow). For all scenario types, we provided the same set of permissions, common to all smartphones: Location, Microphone, Camera, Bluetooth, Wifi, Contacts, and Photos. Next, we asked the participants if they use Instagram, a popular social media app. If they did, then we asked which permissions they allow to this particular social media app to understand how well the scenario description responses represent the permissions granted for a real-world app. We performed the preliminary survey in a graduate level Human Computer Interaction course. A copy of the preliminary survey can be found at this link.

We received 14 responses and performed analysis based on the collected data, comparing the means and deltas of different permissions across different scenarios to the "true" permissions granted to Instagram. We primarily observed no variations in different scenario description methods, which suggests that different scenario types had little to no effect on the response of the participants. Therefore, we proceeded with *vague* scenario framing for simplicity in our main survey design.

3.2 Main Survey Design

The main survey design collected demographic information including age, race/ethnicity (select all that apply), the continent lived on for most of one's life (select all that apply), and gender (select all that apply). Additionally, participants were asked to indicate the degree to which various mobile app permissions (listed above) were related to their privacy and security. Using the vague question framing, we proceeded to asked participants about their privacy

and security preferences related to each mobile app permission across five genres: Social Media, Health, Dating, Food Delivery, and Finance. For each app genre, participants were asked which permissions they would allow (Allow always, Allow only while using the app, Allow once/Ask every time, or Never allow). Additionally, respondents were asked how they expected each type of data would be used by the app: for security/fraud prevention, app features, tracking, or personalization (select all that apply). The survey concluded with a question about any specific stipulations or limitations regarding how each permission should be accessed, used, or shared by the app. To minimize order effects, the order in which each app genre was presented to the participant was fully randomized. A copy of the main survey can be found at this link.

3.3 Sampling Approach

With a primary value of our research being diversity, we set out to gather responses that would represent a broad set of backgrounds, cultures, and identities. Given the scope of this project, we targeted members of the UVA community. Relevant studies examining racial/ethnic differences in privacy concerns, socioeconomic disparities in social media privacy attitudes and behaviors, and gender differences in privacy-related measures among young adults use purposive sampling [14], quota sampling [3], and snowball sampling [5], respectively. Our approach builds upon those, following a multistage purposeful sampling design, integrating maximum variation sampling and quota/criterion sampling [1, 10]. The rationale was that maximum variation purposeful sampling means that the populations sampled from are intentionally as different from each other as possible, so as to exhibit differences. Quotas ensure that each population is approximately equally represented. Stage one, maximum variation, consisted of directly reaching out to student groups that represent each of the cultural identities we are investigating. To ensure an inclusive response pool, we aimed to collect 10 samples from each demographic subgroup (i.e. 10 international students and 10 domestic students). If this predefined quota was not reached at the end of stage one, stage two involved convenience sampling to fulfill the quota.

During stage 2, the following UVA student organizations were contacted through email and/or Instagram private message: International Center (nationality: international), Inter Fraternity Council (gender: male), Inter Sorority Council Public Relations Chairs (gender: female), LGBTQ Center (gender: all), Queer Student Union (gender: queer), National Society of Black Engineers (race/ethnicity: Black), Asian Student Union (race/ethnicity: Asian), Latinx Student Center (race/ethnicity: Latinx), Afro Latinx Student Organization (race/ethnicity: Latinx), and the Native American Student Union (race/ethnicity: Pacific Islander/American Indian/Alaskan Native). We faced a low response rate from the purposive sampling efforts despite reaching out to several organizations. Therefore, we resorted to convenience sampling, where the researchers had people they knew directly complete the survey. Following the convienence sample, the survey was also distributed through the UVA Graduate Computer Science mailing list, the UVA Link Lab mailing list, and physical flyers were posted throughout campus, especially in diversity centered areas. Through this multistage approach, we were able to collect 81 survey responses.

3.4 Participants

There was no particular inclusion or exclusion criteria for the participants, as a diverse sample was favorable. Of the 81 participants, 75.3% (n=61) were 18-24 years old, 21.0% (n=13) were 25-34 years old, and 3.7% (n=3) were in the 35-44 years old category. Geographically, the majority of participants (83.95%, n=68) hailed from North America, followed by 23.46% (n=19) from Asia, and a small portion (1.23%, n=1) from Africa. There were no participants from Europe, South America, or Australia. In terms of ethnicity, the sample was diverse relative to the UVA student body as a whole [7], with 45.68% (n=37) identifying as White or Caucasian, 33.33% (n=27) as Asian American or Asian, and 13.58% (n=11) as African American or Black. We did not meet our quota for the remaining ethnic groups: 11.11% (n=9) Middle Eastern or North African, 2.47% (n=2) Hispanic or Latinx and self-identified, with smaller percentages of Pacific Islander, American Indian or Alaskan Native, and self-identified participants, each making up 1.23% (n=1). Additionally, in terms of gender, the sample was fairly distributed, with 51.85% (n=42) identifying as female, 43.21% (n=35) as male, and 11.11% (n=9) identifying as other (also falling below the quota).

4 Results

Due to time constraints and in an effort to simplify our analysis, we opted to aggregate response data across permission types rather than treat permission type as a third IV. For statistical analysis, we utilized Kruskal-Wallis (KW) tests for ordinal data (Likert scales) and χ^2 tests for binomial data (anticipated data uses). Post hoc tests were conducted using Dunn's test and pairwise χ^2 tests, respectively. A significance level of 0.05 was chosen when interpreting our results and determining the necessity of post hoc tests. In line with our research questions, our results emphasize differences between demographics and app genres over discovering general trends and tendencies in mobile app privacy behaviors. Data and analysis code can be found at this link.

4.1 Demographic Differences

Responses from those who aligned with multiple identities (e.g. male and non-binary) were extrapolated across each of these identities. While not an ideal representation, this decision was made based on the relatively small number of responses that selected multiple identities. If we were to treat 'mixed identity' as a bespoke level, it would contain few samples and impede our analysis. Other notable decisions made during analysis for the sake of maintaining sample sizes include combining non-male and non-female gender identities (denoted as 'other or third gender'), combining racial/ethnic identities with less than 9 samples (n=4; denoted as 'other race/ethnicity'), combining respondents from Africa and Asia into an 'international' group (n=20), and combining age groups older than 18-24 (n=20; denoted as '≥25 years old')

4.1.1 Relevance of Permissions to Privacy and Security. Demographic factors did not have a statistically significant effect on how mobile app permissions were perceived to be relevant to a participant's privacy and security. However, both age (KW = 3.2, p = 0.07) and ethnicity (KW = 8.3, p = 0.08) approached our significance level.

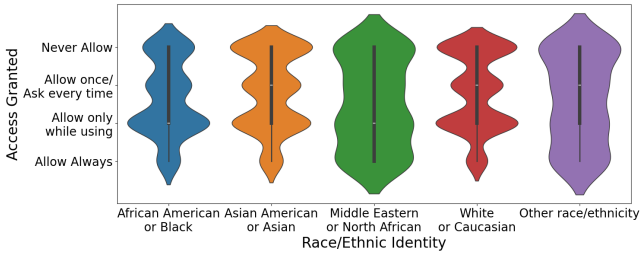

Figure 1: Distribution of data access granted to a hypothetical app via permissions, by race and ethnic identity. Generally, participants most often opted for 'Allow only while using' or 'Never allow'.

4.1.2 Permissions Granted to Apps (RQ1). The degree to which participants would grant permissions to an app varied strongly with their race and ethnicity (KW = 19.5, p < 0.001). Nationality (North American or International) was determined to be not statistically significant, but just barely so (KW = 3.6, p = 0.058). Within race and ethnicity, it was found that Black and Middle Eastern/North African participants differed significantly from their Asian (p = 0.04 and 0.01) and White (p = 0.4 and 0.01) counterparts. This tendency is illustrated in Figure 1, where the medians and overall response distributions for these groups are clearly distinct. Responses from those who identified as Asian or White skewed towards relatively stricter permissions, with the median response being "Allow once/Ask every time". Black and Middle Eastern/North African participants most frequently opted for "Allow only while using". Also, those who identified as Middle Eastern or were grouped into 'Other' tended to grant permissions in a more uniform manner.

4.1.3 Anticipated Uses of Data. All demographic features influenced how participants expected their data to be used by an app, aggregated across all 4 usage types (Figure 2). Starting with race and ethnicity (χ^2 = 65.1, p < 0.001), pairwise differences were observed for all combinations except Asian vs Middle Eastern/North African (χ^2 = 7.6, p = 0.054). Significant observations from these data include that 'Other' and Asian participants anticipated data use for tracking purposes at a greater rate than other groups (62.4% and 51.2%). Black participants anticipated personalization the most often (50.3%), while White participants did so the least (34.1%). Middle Eastern participants selected security-related uses more often (29.1%) than the 'Other' group (7.2%).

Looking next to gender (χ^2 = 17.7, p = 0.007), only the Male vs Female comparison was statistically significant in post hoc testing (χ^2 = 11.6, p = 0.009). Differences in feature-related and tracking-related uses were highlighted by these tests. Those who identified as male expected their data to be used more often for features (73.4%) and tracking (48.1%) than those who identified as female (70.1% and 44.7%). Although gender identities grouped into 'Other' exhibit differences in proportion (especially for uses related to features and tracking), the observed frequencies of anticipated data uses versus male (χ^2 = 5.8, p = 0.12) and female (χ^2 = 6.9, p = 0.07) participants were not statistically significant at our p-value.

International vs North American participants (χ^2 = 17.1, p < 0.001) as well as participants from our either side of our 25-year-old age boundary (χ^2 = 39.4, p < 0.001), also differed. Participants from Africa or Asia expected data to be used by an app for security purposes to a greater degree (20.2% vs 13.3%), while participants from North America anticipated more use for app personalization (40.6% vs 36.7%), tracking (49.7% vs 38.9%), and features (75% vs 63.3%). Lastly, participants over the age of 25 anticipated data use for tracking (57.9%) and personalization (47.4%) at a proportion far greater than those 18-24 years old (42.6% and 36.4%), but expected their data to be used less for security (10.1% vs 16.6%) and features (60.6% vs 75.2%).

4.2 App Genre Differences

4.2.1 Permissions Granted to Apps (RQ2). The app genre provided in the scenario was found to significantly affect the degree of permissions allowed to the app (KW = 120.7, p < 0.001). Overall, Social Media and Health apps were allowed the more relaxed permissions (median = 'Allow only while using') while Delivery and Finance apps were given more strict access (median = 'Allow once/Ask every time'; Figure 3). Post hoc tests uncovered several pairs that differed significantly. The Delivery app scenario differed from the Social Media (p < 0.001) and Dating app (p = 0.02) scenarios. Likewise, the Finance app scenario differed from the Social Media (p < 0.001) and Dating app (p < 0.001) scenarios, but was also distinct from the Health app scenario (p = 0.002). This is likely because participants more frequently selected 'Never allow' and 'Allow once/Ask every time' for the Finance scenario (n = 185 and 73) compared to the Delivery scenario (n = 174 and 58). However, it should be noted that the significance level for Delivery vs Health was very close to our cutoff (p = 0.06).

4.2.2 Anticipated Uses of Data. Anticipated uses indicated by the participants was also affected by the app genre in the scenario (χ^2 = 58, p < 0.001). Overall, a Social Media app was most frequently expected to use data for personalization (49%). Dating, Social Media, and Health app scenarios were more often anticipated to use data for app features (79.2%, 79.1%, and 75.9%). The Finance app scenario was once again found to be statistically distinct from the Social Media (χ^2 = 40.6, p < 0.001), Health (χ^2 = 29, p < 0.001), and Dating (χ^2 = 21.7, p < 0.001) app scenarios, but also from the Delivery app scenario (χ^2 = 10.2, p = 0.02). Feature (55.9%) and personalization-related (30.1%) uses were selected less frequently for the Finance app scenario compared to the Social Media, Health, and Dating scenarios, while security-related uses (24.5%) were anticipated more often relative to all of the statistically significant pairings.

The Delivery app scenario also differed from the Social Media app scenario (χ^2 = 11.3, p = 0.01), with significant decreases in the frequency of anticipated feature (67.2% vs 79.1%) and personalization-related (33.4% vs 49%) uses. Interestingly, no app genre pair exhibited significant differences in the frequency of tracking-related uses, and participants anticipated each app genre to use data for tracking purposes at about the same rate (42.5% to 47.8%).

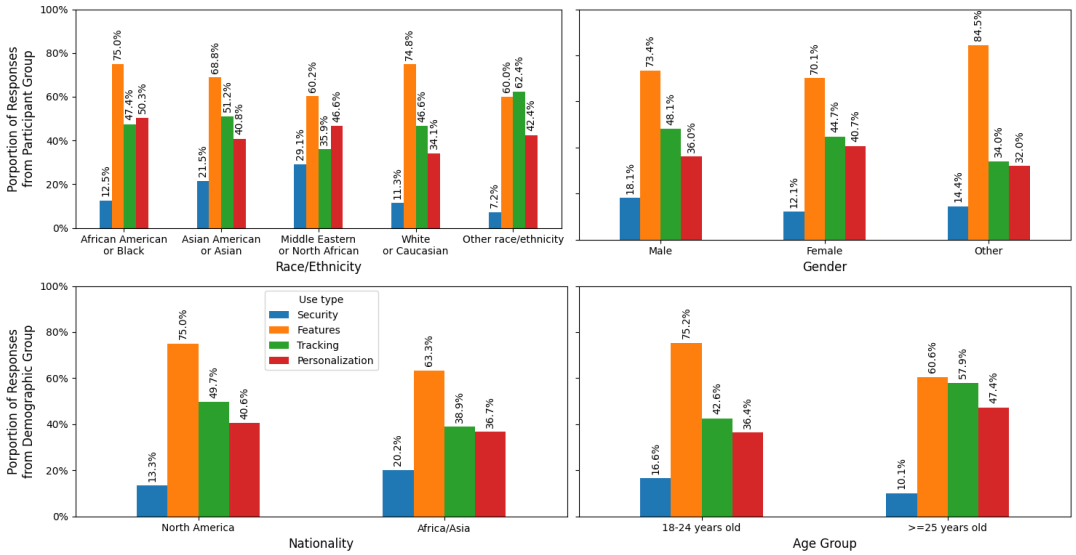

Figure 2: The rate that a participant expected their data to be used by an app for security, features, tracking, and personalization, by demographic identity. Generally, participants expected data to be used for app features the most, followed by tracking, personalization, then security.

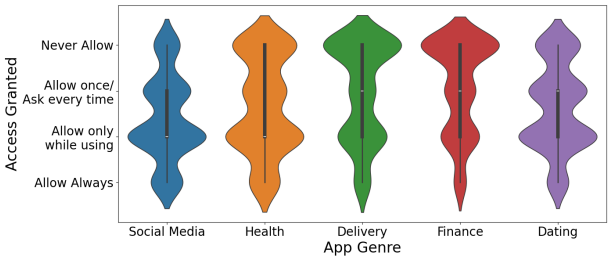

Figure 3: Distribution of data access granted to a hypothetical app via permissions, by app genre.

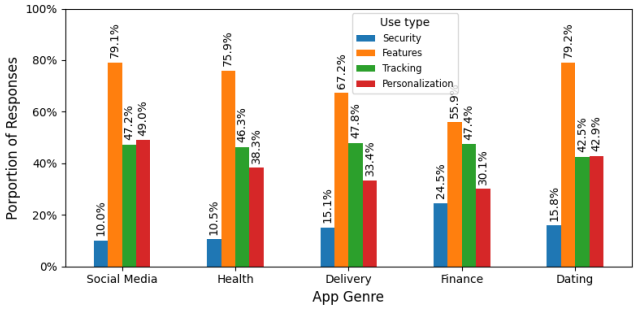

Figure 4: The rate that a participant expected their data to be used by an app for security, app features, tracking, and personalization, by app genre.

4.3 Interaction Between Demographics and App Genre

To test for interactions between demographics and app genre, we utilized two-way analysis of variance (ANOVA) on rank-transformed ordinal data with a Tukey's HSD post hoc analysis. Therefore, we note that 1) interaction effects should be interpreted as an effect on the rank-ordered data, not the raw data itself and 2) this approach could result in inflated Type I error. Despite these limitations, we proceed with this approach in light of recent work that demonstrates its validity, compared to common alternatives [13]. Furthermore, we did not encounter results that were symptomatic of inflated Type I error. Cramér's V test was applied to the binomial data for pairwise associations between levels of each independent variable.

4.3.1 Permissions Granted to Apps (RQ3). There was a significant interaction between nationality (North America vs International) and app genre ($F = 2.7$, $p = 0.03$). Compared to participants from North America (NA), participants from Asia or Africa (AA) tended to allow more access in the Dating and Health app scenarios (median = 'Allow once/Ask every time' vs 'Allow only while using' for both), less access for Delivery (median = *between levels 2 and 3* vs 'Allow once/Ask every time') and Finance (median = 'Allow once/Ask every time' vs 'Never allow'), and comparable access for Social Media (median = Allow only while using; Figure 5).

For space, a table reporting the results of the post hoc tests has been omitted and significant comparisons will be summarized instead. Beginning first with comparisons between different nationalities, the interaction between NA and the Finance app scenario (NA x Finance) was distinct from AA x Dating ($p = 0.001$), AA x Delivery ($p = 0.03$), AA x Health ($p = 0.02$), and AA x Social Media ($p < 0.001$). Relative to these pairs, participants selected the most

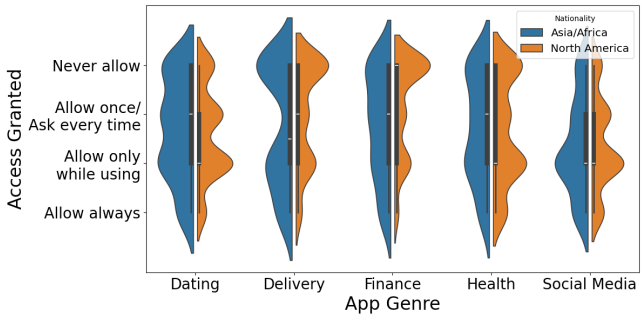

Figure 5: Distribution of data access granted to a hypothetical app via permissions, by app genre and nationality (North America or International).

strict permissions for NA x Finance (median = 'Never allow'). In addition, NA x Delivery was distinct from AA x Social Media ($p < 0.001$) and NA x Social Media was distinct from AA x Delivery ($p = 0.02$). For both NA and AA participants, Social Media permissions were rather liberal, while NA x Delivery and AA x Delivery were more strict.

Within NA participants, the genres of Delivery and Finance once again dominate the significant comparisons. Both differed from the Dating, Health, and Social Media scenarios (all $p < 0.001$). As discussed before, this is likely a result of the comparatively strict permissions grated under the Delivery and Finance scenarios. Finally, NA x Health was found to be statistically different from NA x Social Media ($p = 0.02$). Although both share the same median ('Allow only while using'), the upper quartile of NA x Health is stricter ('Never allow', n = 98 vs 42). No significant interactions were found between genres within AA (e.g. AA x Dating vs AA x Health). Although the main effect for race/ethnicity was confirmed by ANOVA ($F = 5.2$, $p < 0.001$), the interaction with genre did not meet our significance level ($F = 1.6$, $p = 0.07$). The interaction between gender and genre also approached a p-value of 0.05 ($F = 1.9$, $p = 0.051$).

4.3.2 Anticipated Uses of Data. Only 'medium' effect sizes (ES) between demographic features and app genre were found (Cramér's $V \in [0.2, 0.6]$) [2]. Gender identities grouped as 'Other' anticipated data uses that had a medium association with the Social Media and Delivery app scenarios (ES = 0.27 and 0.2). Race and ethnic identities grouped as 'Other' were similarly associated with the Dating app scenario (ES = 0.2).

5 Discussion and Conclusions

The results of our study reveal the different ways demographic factors and app genres influence permissions granted to mobile applications (RQs 1, 2, and 3), as well as the expected use of these permissions. While race and ethnicity exhibited statistically significant effects on permission granting behaviors, others, including gender, nationality, and age did not (RQ1). In particular, Black and Middle Eastern/North African participants were more inclined to grant permissions "only while using" an app, characterizing different privacy considerations compared to White and Asian participants. In contrast, White and Asian participants demonstrated a tendency toward stricter permissions overall. This finding may be a result of the different ways these groups anticipated that data would be used by an app. For example, Black participants expected data to be used for personalization much more often than White participants, which may have led to more relaxed permissions (Fig. 2).

Similarly, the role of the app genres in the privacy decisions of the participants was found to be significant (RQ2). Social Media and Health apps were granted relatively relaxed permissions, reflecting user trust in these categories or their perceived necessity for app functionality. However, the Finance and Delivery app scenarios received more restrictive responses, possibly due to the sensitive nature of financial information and skepticism about data security in delivery services. Anticipated uses may reflect this, as we found there was lower anticipation of data usage for app functionality in delivery and finance apps compared to other genres. Additionally, the finance app scenario had an increased expectation to provide security via data-use.

Differences in the genre of the application did not exhibit substantial variations in tracking-related expectations. This implies that our participants expected to be tracked at similar rates by each type of app. Evidence of an interaction between demographics and app genres further highlights the complexity of user privacy behaviors (RQ3). Participants from different regions, such as North America and Asia or Africa, displayed significant variations in permissions granted to the different app genres, highlighting the influence of cultural and regional factors on privacy preferences and its relation with the app genre.

These findings emphasize the importance of tailoring privacy-related design and communication strategies to diverse user groups. Developers and policymakers should take into account demographic-specific concerns and app genre characteristics when designing privacy settings and permissions interfaces. For example, allowing users to provide their desired level of permission access can help ensure inclusivity. Although this is already standard practice, we observed that our median responses were overwhelmingly split between "Allow once/Ask every time" and "Allow only while using", indicating that more options in between these levels (e.g. "Allow for one hour") should be seriously considered. Additionally, providing explanations related to why a particular permission may facilitate trust between the developers and users. Future work may consider exploring whether anticipated uses change when explanations are provided. Lastly, providing real-time transparency on how user data is being used and to what degree user activities are being tracked may build user literacy of permissions and privacy while encouraging developers to not request unnecessary permissions and to faithfully report data uses.

Future research could address limitations such as the relatively small sample sizes for certain demographic categories (such as African American or Black, Middle Eastern or North African, and Hispanic or Latinx). Additionally, it would be worthwhile to form an understanding of how combinations of identities (e.g. mixed-race or female and Black) inform mobile app privacy decisions. We were also unable to fully consider the impact of different permission types nor the data usage limitations described by our participants. Expanding participant diversity, accounting for intersecting identities, and using more sophisticated methods to analyze privacy behaviors could provide deeper insight into these dynamics. By

understanding and respecting the diversity in the context of mobile app privacy preferences, better systems can be created which align with the expectations and needs of their users.

References

[1] Steve Campbell, Melanie Greenwood, Sarah Prior, Toniele Shearer, Kerrie Walkem, Sarah Young, Danielle Bywaters, and Kim Walker. 2020. Purposive sampling: complex or simple? Research case examples. *Journal of Research in Nursing: JRN* 25, 8 (Dec. 2020), 652–661. https://doi.org/10.1177/1744987120927206

[2] IBM Corporation. 2024. Cramér's V. https://www.ibm.com/docs/en/cognos-analytics/11.1.0?topic=terms-cramrs-v

[3] Dmitry Epstein and Kelly Quinn. 2020. Markers of Online Privacy Marginalization: Empirical Examination of Socioeconomic Disparities in Social Media Privacy Attitudes, Literacy, and Behavior. *Social Media + Society* 6, 2 (April 2020), 2056305120916853. https://doi.org/10.1177/2056305120916853 Publisher: SAGE Publications Ltd.

[4] A. Hamed and H. K. Ben Ayed. 2016. Privacy risk assessment and users' awareness for mobile apps permissions. In *2016 IEEE/ACS 13th International Conference of Computer Systems and Applications (AICCSA)*. Agadir, Morocco, 1–8. https://doi.org/10.1109/AICCSA.2016.7945694

[5] Mariea Grubbs Hoy and George Milne. 2010. Gender Differences in Privacy-Related Measures for Young Adult Facebook Users. *Journal of Interactive Advertising* 10, 2 (March 2010), 28–45. https://doi.org/10.1080/15252019.2010.10722168 Publisher: Routledge _eprint: https://doi.org/10.1080/15252019.2010.10722168.

[6] Hannah J. Hutton and David A. Ellis. 2023. Exploring User Motivations Behind iOS App Tracking Transparency Decisions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3544548.3580654

[7] Institutional Research & Analytics. 2024. Enrollment. https://ira.virginia.edu/university-data-home/enrollment

[8] Rena Lavranou, Stylianos Karagiannis, Aggeliki Tsohou, and Emmanouil Magkos. 2023. Unraveling the Complexity of Mobile Application Permissions: Strategies to Enhance Users' Privacy Education. *European Journal of Engineering and Technology Research* (Dec. 2023), 87–95. https://doi.org/10.24018/ejeng.2023.1.CIE.3141

[9] Maximilian Marsch, Jens Grossklags, and Sameer Patil. 2021. Won't You Think of Others?: Interdependent Privacy in Smartphone App Permissions. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (Oct. 2021), 437:1–437:35. https://doi.org/10.1145/3479581

[10] Lawrence A. Palinkas, Sarah M. Horwitz, Carla A. Green, Jennifer P. Wisdom, Naihua Duan, and Kimberly Hoagwood. 2015. Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and policy in mental health* 42, 5 (Sept. 2015), 533–544. https://doi.org/10.1007/s10488-013-0528-y

[11] Yong Jin Park. 2015. Do men and women differ in privacy? Gendered privacy and (in)equality in the Internet. *Computers in Human Behavior* 50 (Sept. 2015), 252–258. https://doi.org/10.1016/j.chb.2015.04.011

[12] Yong Jin Park. 2021. Why privacy matters to digital inequality. In *Handbook of Digital Inequality*. Edward Elgar Publishing, 284–295. https://www.elgaronline.com/edcollchap/edcoll/9781788116565/9781788116565.00028.xml Section: Handbook of Digital Inequality.

[13] Theophanis Tsandilas and Géry Casiez. 2024. The illusory promise of the Aligned Rank Transform. https://statransform.github.io/jovi/

[14] Laurent H. Wang and Miriam J. Metzger. 2024. The Online Privacy Divide: Testing Resource and Identity Explanations for Racial/Ethnic Differences in Privacy Concerns and Privacy Management Behaviors on Social Media. *Communication Research* (Aug. 2024), 00936502241273157. https://doi.org/10.1177/00936502241273157 Publisher: SAGE Publications Inc.

Received 10 December 2024