Who Am I? A Design Probe Exploring Real-Time Transparency about Online and Offline User Profiling Underlying Targeted Ads

NATÃ M. BARBOSA, University of Illinois at Urbana-Champaign, USA
GANG WANG, University of Illinois at Urbana-Champaign, USA
BLASE UR, University of Chicago, USA
YANG WANG, University of Illinois at Urbana-Champaign, USA

To enable targeted ads, companies profile Internet users, automatically inferring potential interests and demographics. While current profiling centers on users' web browsing data, smartphones and other devices with rich sensing capabilities portend profiling techniques that draw on methods from ubiquitous computing. Unfortunately, even existing profiling and ad-targeting practices remain opaque to users, engendering distrust, resignation, and privacy concerns. We hypothesized that making profiling visible at the time and place it occurs might help users better understand and engage with automatically constructed profiles. To this end, we built a technology probe that surfaces the incremental construction of user profiles from both web browsing and activities in the physical world. The probe explores transparency and control of profile construction in real time. We conducted a two-week field deployment of this probe with 25 participants. We found that increasing the visibility of profiling helped participants anticipate how certain actions can trigger specific ads. Participants' desired engagement with their profile differed in part based on their overall attitudes toward ads. Furthermore, participants expected algorithms would automatically determine when an inference was inaccurate, no longer relevant, or off-limits. Current techniques typically do not do this. Overall, our findings suggest that leveraging opportunistic moments within pervasive computing to engage users with their own inferred profiles can create more trustworthy and positive experiences with targeted ads.

CCS Concepts: • **Human-centered computing → Empirical studies in HCI**.

Additional Key Words and Phrases: transparency, online behavioral advertising, profiling, technology probe

ACM Reference Format:
Natã M. Barbosa, Gang Wang, Blase Ur, and Yang Wang. 2021. Who Am I? A Design Probe Exploring Real-Time Transparency about Online and Offline User Profiling Underlying Targeted Ads. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3, Article 88 (September 2021), 32 pages. https://doi.org/10.1145/3478122

1 INTRODUCTION

Online Behavioral Advertising (**OBA**), which we define as the targeting of online ads based on a user's online (and potentially also offline) activities, is becoming more sophisticated and ubiquitous. Over the years, Internet companies – whose main business models rely on ad targeting – have stretched the boundaries of user tracking and profiling. Today, data about users' behaviors, interests, and identities are captured from computers, phones, wearables, and household devices [40]. This data is collected, processed, and used to target ads. While such ads may be more relevant, OBA engenders surveillance capitalism [56], where personal data is a commodity [12].

Authors' addresses: Natã M. Barbosa, University of Illinois at Urbana-Champaign, USA, natamb2@illinois.edu; Gang Wang, University of Illinois at Urbana-Champaign, USA, gangw@illinois.edu; Blase Ur, University of Chicago, USA, blase@uchicago.edu; Yang Wang, University of Illinois at Urbana-Champaign, USA, yvw@illinois.edu.

2474-9567/2021/9-ART88 $15.00
https://doi.org/10.1145/3478122

The evolution of OBA has created information asymmetries and power imbalances. Few regulations mandate transparency about tracking or profiling. As a result, ad networks expand user-profiling techniques while keeping the process largely opaque to users. While ad networks give prospective advertisers increasingly extensive options for micro-targeting users, they give users only limited visibility into their own ad-targeting profiles [50]. Ad networks like Google and Facebook tend to disclose only the most generic ways in which users are tracked and profiled [4, 17]. Existing transparency mechanisms are often inaccurate, incomplete, and insufficient [4, 50]. These transparency mechanisms are also typically post-hoc, not providing any visibility until long after information has been collected, processed, and used to target ads. As a result, users are left in the dark about how their present actions might affect what ads they will see in the future. Users therefore develop folk theories, such as that apps eavesdrop on their conversations [7, 34]. Lacking the knowledge and opportunities to act upon their concerns and preferences, users feel privacy invasion, helplessness, resignation [49], and distrust of companies.

Prior studies of OBA have uncovered four key user-facing problems: (1) a lack of user awareness and participation [2, 48]; (2) insufficient user understanding of how profiling works [38, 48, 55]; (3) profiles that can misrepresent a user [2, 15, 17]; and (4) a lack of direct control over profiling [13, 38]. These problems partially stem from a lack of transparency about how profiles are created and used. Increased transparency may have mixed effects [14]. While it can increase awareness and inspire action [51], it can also add unnecessary friction [17, 41]. Much remains to be learned about designing meaningful transparency mechanisms concerning user profiling.

To explore speculative, novel mechanisms for improving transparency and control of ad-targeting profiles in a world of pervasive computing, we designed and deployed a technology probe. Our probe differed from existing transparency mechanisms in two key ways. First, while existing mechanisms provide transparency long after data has been collected and used, our probe provides real-time data about the incremental evolution of a profile. It also highlights the specific user behaviors that prompted those changes. Second, while existing mechanisms focus on activities online and give only vague hints about the provenance of profile data, our probe delineated how specific activities, both online and offline, might feed into profiles. We built a browser extension and a smartphone app to instrument a user's primary web browser (computer) and smartphone, respectively. Our probe thus fuses detailed information about the user's web browsing, movements in the physical world, and other offline activities. Such combined data promises to be the next frontier in pervasive ad targeting [24]. It represents the application of techniques from ubiquitous and pervasive computing, such as context sensing, to advertising.

In a longitudinal field study involving daily diary activities, 25 participants used our technology probe over the course of two weeks. As a result, participants were given transparency about their activities and profiles displayed in novel ways — focused on the real-time evolution of a profile and the precise sources of both online and offline data. We found that increased transparency can lead to meaningful revelations, opportunities for identity performance and boundary regulation, and elucidation about profiling and ad-targeting processes. The study also revealed nuances about implementing transparency. For instance, participants' desired level of engagement with their profiles varied in part based on their overall attitudes toward ads. Nonetheless, some participants engaged with their profiles to prevent irrelevant ads, prevent ads triggered by private activities, and remove activities they did not expect to be used to target ads. These results also showed the potential for more active participation in profiling in order to correct inaccuracies and benefit more from ads. The desire to perform identities and prevent undesirable experiences can further motivate user engagement. We present the ensuing design implications.

The key contributions of this paper are: (1) the design, implementation, and deployment of a novel technology probe that fuses both online and offline activities and enables real-time awareness and control of user profiles; (2) insights about how users react to real-time transparency and control about profiling processes, including clear indications of data provenance; and (3) new, speculative ideas inspiring future transparency designs for OBA.

2 BACKGROUND AND RELATED WORK

OBA relies on tracking users' activities to target them with relevant and personalized ads. Over the years, its sophistication and ubiquity have grown significantly. Today, user profiles are based not just on browser and search history, but also email and social media activity, behavior tracked on smartphones (e.g., precise location data, app installations), and even offline activity like in-store purchases [3, 50]. Advertisers can thus target users based on an extensive list of demographics, interests, and behaviors. OBA has been found to track sensitive topics, such as religion, health, and sexual orientation [10]. It has proven difficult for even advanced users to escape [1].

2.1 Tracking, Profiling, and Privacy in Ubiquitous Computing

The tracking mechanisms that underpin user profiling for ad targeting embody a type of ubiquitous computing system. Through a variety of sensors on different consumer devices and platforms – some of which have arguably *"woven themselves into the fabric of everyday life"* [53] – a progressively enhanced picture of people's behaviors, likes, and needs can be created by triangulating different signals across personal, wearable, and home devices, as well as websites and apps. As Weiser himself put it in his vision for ubiquitous computing, computers are now made "indistinguishable" in everyday life. Indeed, it can be said that ubiquitous tracking and profiling can be characterized by their invisibility and memory amplification [33].

Users thus lack clarity about what activities are being tracked and used for ad targeting [15, 17, 55]. Moreover, tracking, profiling, and targeting technologies are so entrenched in daily life that users may be unable to distinguish between when they have unwittingly generated explicit signals that led to a given advertisement and when the technology is so attuned to their inner thoughts that it accurately predicts a hidden interest. Privacy is an important aspect of realizing a well-implemented version of ubiquitous computing [8, 52]. Compounding the importance of this issue, people are growing ever more dependent on technologies, personal devices are gaining more powerful sensing and inferencing capabilities, and new consumer devices and apps are being introduced by companies whose core business relies on targeted ads (e.g., Facebook's Portal, Oculus, and rumored SmartWatch [46], as well as Google's Android OS and Nest devices). Together, these trends create a reality of ubiquitous advertising [32] and pose serious and meaningful challenges to user privacy and trust. Therefore, promoting transparency and control of underlying tracking and profiling systems may allow users to control and optimize their ad experiences, as well as to build greater trust and understanding. Such efforts could also increase user participation in engaging with their profiles, helping to prevent an undesirable future where powerful sensing systems are concealed from users while the limits of what is acceptable are pushed as a result of power imbalances – when providers of free services benefit from a given information flow more than their users [27].

Many notable papers about tracking focus on location data. These include both Beresford and Stajano's [9] and Lederer's [35] frameworks for tracking location in ways that balance convenience with user privacy. However, modern technologies enable not only extremely accurate location tracking, but also inferences about the associated offline contexts, including physical activities, user intentions (e.g., searching for nearby points of interest on a phone app [39]), human mobility [31], sleep patterns [30], and even mood [5]. As a result, improving the transparency and control of ubiquitous tracking and profiling in the physical world can mutually benefit users and service providers, especially since users' mental models of OBA are heavily rooted in the tracking of online shopping, web browsing, and social media activity [55]. Increasingly, however, data for ad targeting is being gathered from ubiquitous computing contexts and the physical world [3, 24, 50].

2.2 Existing Transparency and Control in OBA

Inferences made in the profiling process remain largely invisible to users and out of their control. Such inferences are frequently inaccurate, making users feel embarrassed, misjudged, and powerless [2, 17, 42]. Nonetheless, when hidden inferences beget extremely accurate ads, users may incorrectly blame eavesdropping [7]. Not knowing

how and when inferences are made, users develop folk theories and act based on potential misconceptions [16, 55]. Prior work has found that users demand more transparency and control over what is happening "under the hood" [15]. While more transparency is generally welcome, its effects can be mixed. Users may want more transparency, but they also want to keep their experiences convenient; technical terms may get in the way [41]. "Full transparency" can also be undesirable. Both highly specific and overwhelmingly vague explanations of OBA can be considered creepy [17, 50]. Some researchers argue that users prefer "medium" levels of transparency [14].

Ad networks commonly provide some visibility into how ads are personalized. On Facebook, ad interests can be managed via the Ad Settings page.[1] This page lets users visualize interests inferred across different categories. Clicking an interest provides a brief and vague explanation of why it was added. The page also lets users remove interests. Google provides a similar page: Ad Personalization settings.[2] Users can see a list of factors considered in personalizing ads. Clicking a factor explains generally why that factor may have been added, with a link redirecting users to Google's user activity dashboard showing a user's entire list of actions tracked by Google across products and services.[3] While this list gives users some idea of the extent of Google's tracking, it does not link back to the factors shown on the previous page. Ad networks sometimes provide explanations for individual ads, usually by clicking a nearby *"Why am I seeing this ad?"* icon on the ad placement [50]. In both cases, information about ad profiles is only accessible post-hoc. That is, these tools give users (limited) visibility long after data has been collected, processed, and potentially used for targeting. Currently, users cannot control profiling *as it happens*, partially inspiring our technology probe. The information provided in dashboards and explanations often does not enable users to see a full picture of their ad-targeting profile at any point in time.

Researchers have studied these existing transparency mechanisms, finding them incomplete, vague, ineffective, and often misleading. Andreou et al. showed that ad-interest dashboards do not enable users to trace parts of their profiles back to the actions that originated them [4]. They also found that such tools do not provide temporal awareness into how profiles change over time. Broadly, there is a discrepancy between the level of transparency provided to users and the targeting tools given to advertisers. For example, Facebook's ad explanations often state that *"there could also be more factors not listed here."* Further, ad explanations explicitly exclude data obtained from "data partners" who obtain data elsewhere, including offline [3]. Users thus struggle to understand how their actions impact ad targeting and what steps they can take to prevent undesirable inferences out of profiles.

Most recently, Rader et al. conducted interviews investigating reactions to Facebook and Google's inference dashboards [42]. They found that participants' interpretations of items in their profiles were largely guided by participants being able to recall and justify their past behavior in relation to the underlying platforms. Most importantly, their study unearthed a promising shift in OBA transparency mechanisms from presenting a descriptive picture of users' past actions to helping users anticipate future targeting.

2.3 Prior Attempts to Make OBA More Transparent

Prior work has built novel OBA transparency tools that visualize information differently than our technology probe. Weinshel et al.'s browser extension "tracked the trackers" locally, letting users visualize the full set of data an advertiser may have collected about them over time, as well as what particular interests that data might suggest [51]. Their field deployment found that increased visibility of tracking increased user intentions to take protective actions. Jin et al. showed participants a workflow visualization of a simulated ad-bidding process, finding that participants preferred workflow-based visualizations linking recognizable actions to the selection of ads [28]. They posited that users may become more accepting of OBA if they can examine, control, and understand ad targeting more directly. A recent qualitative study by Shah mocked up a "forget" feature on tracking data

[1] https://www.facebook.com/adpreferences/ad_settings, then click "categories used to reach you" > "interest categories" (must be logged in)
[2] https://adssettings.google.com/authenticated (must be logged in)
[3] https://myactivity.google.com (must be logged in)

associated with a given ad [44]. They found that this feature increased user comfort and acceptance. Participants planned to use such a feature for ads perceived as creepy or invasive. The feature gave them a sense of control.

Researchers have explored the design of explanations of why a single ad was shown, finding benefits to incorporating explanations into ads themselves to promote "algorithm engagement" [4, 15, 17]. Increased visibility of inferences can lead to "algorithm disillusionment," demonstrating that algorithms are far from perfect [17].

Increased transparency can help users better participate in the composition of their own "algorithmic selves" [17]. Collectively, findings from past work also hint that increased user participation could bolster trust and lead to more realistic impressions of the accuracy and complexity of ad-targeting profiles.

Differences from prior work: To the best of our knowledge, no prior work has studied users' reactions to an incremental, longitudinal, and comprehensive view of the "knowledge surface" generated from users' online and offline activities that could feed into targeted ads. Instead, prior work mostly focused on web browsing behavior. Providing incremental, in-the-moment, and both online (browser-based) and offline (smartphone-based) transparency of OBA mechanisms is the main difference, and advantage, of our work compared to prior attempts to make OBA more transparent. This is advantageous because browsing behavior may account for only a small fraction of ad targeting, with increasingly more targeting from contexts beyond online browsing, such as offline contexts (e.g., physical store visits) and "data partners" [4, 6, 50]. Such an approach is thus more aligned with the reality of ubiquitous advertising on personal devices moving beyond online, browser-based contexts [32]. Our work can lead to greater awareness and control of profiling, as well as improve the accuracy of users' understanding, expectations, and perceptions of OBA. It could also improve user engagement with control mechanisms and surface new facets of those mechanisms in a world of ubiquitous sensing and tracking.

2.4 The Technology Probe Method

Technology probes are research artifacts used to "find out about the unknown" and "hopefully return with meaningful data" [26]. They are "simple, flexible technologies with three goals: the social science goal of collecting in-context information about the use and the users, the engineering goal of testing the technology, and the design goal of inspiring users and researchers to envision future technologies" [26]. Technology probes are used as situated research artifacts to inform future design directions. For example, in a study of privacy tools for smart homes, Seymour et al. showed that the longitudinal deployment of a privacy assistant technology probe prompted users to "co-adapt," "find new control mechanisms," and "suggest new approaches to address the challenge of regaining privacy in the connected home" [43]. Kaur et al. explored whether interpretability tools helped data scientists understand machine learning models, uncovering recurring problems [29]. Their probe showed that data scientists "over-trust" and misuse machine learning models, influencing the design of future tools.

A technology probe is suitable for our work because there is currently no way for users to directly monitor and control the composition of their actual ad-targeting profiles in real time, necessitating a realistic simulation. Because user acceptance, perceptions, and desired engagement with OBA can be mixed [14, 41], it was unknown whether increased transparency about how a profile evolves over time would be beneficial or bring unexpected results. Accordingly, our technology probe aimed to capture: (1) in-context reactions to the evolution of a simulated ad-targeting profile (a social science goal); (2) the feasibility of providing transparency about a profile (an engineering goal); and (3) the act of increasing transparency and user participation (a design goal).

3 THE PROBE

We named our technology probe WHO AM I, which was the user-facing name given to both our browser extension and mobile app. First, we introduce the probe's main features in Section 3.1 and give example scenarios in Section 3.2. We then detail the probe's technical implementation in Section 3.3, with further details in Appendix B.

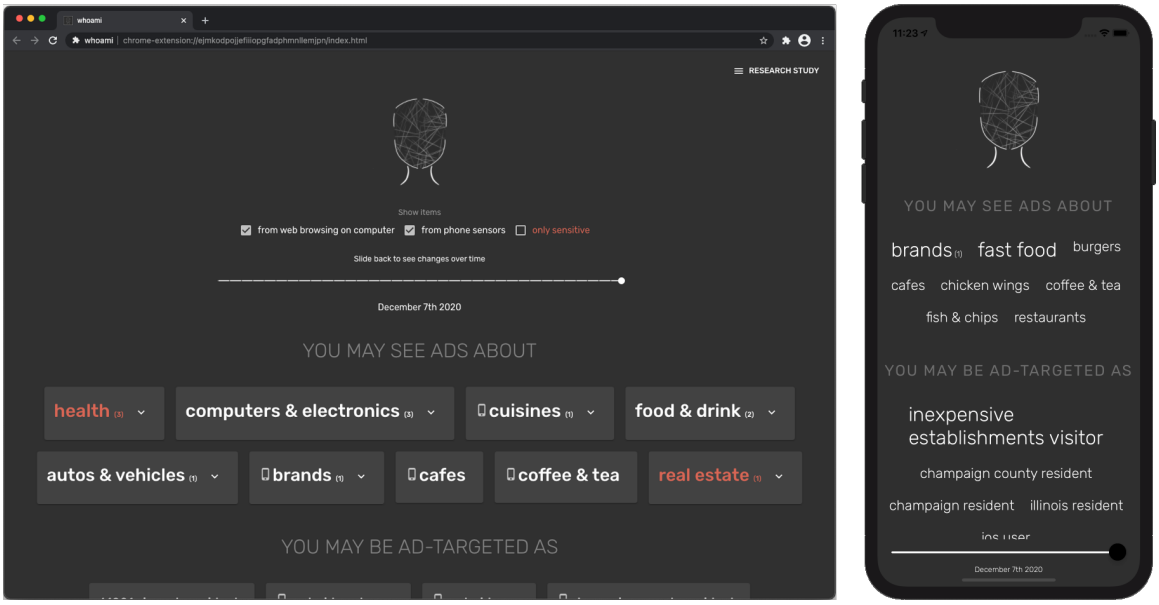

Fig. 1. The WHO AM I probe. Participants used both the browser extension (left) and mobile app (right) for two weeks.

3.1 Probe Features

In Hutchinson et al.'s definition [26], technology probes are distinct from prototypes in five aspects: *functionality* (probes have at most a few features); *flexibility* (probes enable use in unexpected ways); *usability* (probes do not emphasize usability); *logging* (probes focus on idea generation); and *design phase* (probes are used "early in the design process…for challenging pre-existing ideas"). Following these principles, our probe (1) notifies users about additions to their profile as they carry out daily activities and (2) lets users review their profile at any time.

In-the-moment notifications. Notifications are sent as soon as new items are added to the user's profile. For example, if the user visits `cvs.com`, a notification says they may see ads about "pharmacy" in the future. If they walk into a Starbucks store with their phone on them, a notification tells them they may see ads about "coffee & tea" in the future. When multiple profile items are added from a single activity, they are grouped into a single notification. For example, visiting a Target store adds "department stores" and "grocery stores" to a profile.

Profile visualization. Users can always see all items cumulatively included in their profile. Items can be filtered based on their potential sensitivity and their source. Items added from web browsing follow the hierarchical structure of Google AdWords,[4] grouped by root level. Profile items are grouped into three main sections: *"you may see ads about"*; *"you may be ad-targeted as"*; and *"hints about your intents today."* The first section shows inferences about the user's potential interests and the second shows items that are relatively stable and certain (e.g., the user's geolocation, the brand and model of their phone). The third section contains speculative items hinting at immediate needs, such as "looking for bakeries" after the user searches for "Panera" on Google Maps, or "had little sleep" after having the probe detect under seven hours of device charging in the evening. Items in the *"hints about your intents today"* section expired (and were removed) at the end of each day, with new items added the next day. Users could move a slider to go back in time, visualizing profile additions from each day.

To spur speculation, the probe incorporated other features detailed in Appendix C, including profile item details, ad examples, ad tags, and profile controls. Showing ad examples was suggested by pilot users to provide a

[4] https://developers.google.com/adwords/api/docs/appendix/verticals

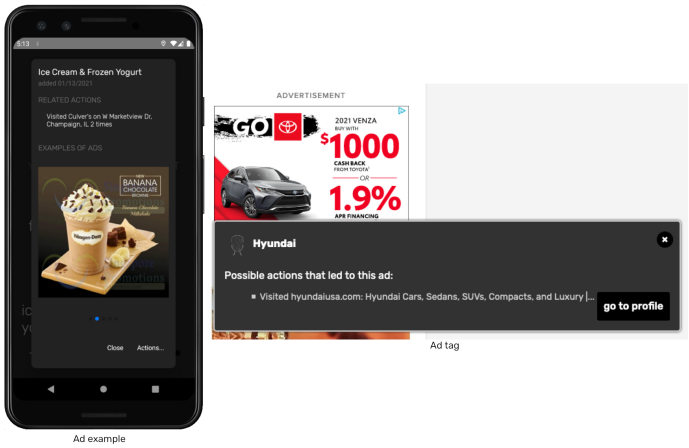

Fig. 2. An ad example (left) and ad tag (right) for a given profile item provided by the WHO AM I probe.

more visual way of interpreting a profile. Ad tags, added near actual ads on the computer, suggested connections between actual ads and participants' profiles. Figure 2 shows an ad example (left) and an ad tag (right).

In summary, the WHO AM I probe abides by the five principles of technology probes [26] in the following ways:

- *Functionality*: The probe has a limited number of features. These key features are the profile visualization and notifications of incremental additions to the profile.
- *Usability*: The emphasis is on information related to a potential ad-targeting profile and its evolution. These items are front and center in the user experience, rather than designing for particular user groups.
- *Flexibility*: We added minor features, such as *removing* actions and *marking them as sensitive*, to let participants modify their profiles in ways we would not have anticipated. Doing so surfaces different ways in which users might want to engage with their ad-targeting profiles.
- *Logging*: The probe implemented comprehensive, fine-grained logging of participants' actions to gauge their interest in their own ad-targeting profiles and create a shared understanding of their experience for idea generation during the exit interview.
- *Design Phase*: The probe challenges pre-existing ideas of post-hoc profile visibility, testing an alternative.

3.2 User Experience Example Scenarios

The following are example scenarios of how users might experience the WHO AM I probe.

- Alice goes to Chipotle for lunch. Moments after entering the restaurant, she gets a notification on her phone saying she might see ads for "Mexican food" and "fast food" in the future. She can tap the notification to see the new items added to her profile, the exact reason why they were added (e.g., "visited Chipotle on 42 W Almond St"), and examples of ads she might expect to see for each interest category in the future.
- Darnell leaves work at 5pm and heads home. As he starts driving, he gets a notification on his phone that he may be targeted as a "vehicle commuter" in the future and that he might see ads about "auto insurance." When he arrives home, he goes for a run while carrying his phone. Shortly after he starts running, the phone displays that he may be targeted as a "runner" and that he may see ads for "sports apparel."
- Layla does a Google search for "Apple credit card." Upon landing on the search results, she gets a notification on her browser saying that she might see ads for "finance" and "credit & landing" in the future.
- Miguel is an Android phone user. He keeps in touch with his family in Guatemala via instant messaging apps. Upon sending texts in Spanish to his family group chat, he gets a notification saying that he may be

targeted as a "Spanish speaker" in the future. Later, he sends his close friend a message saying, *"Happy Friday, buddy! I'm looking forward to going to the game together tonight!!"* The phone sends him a notification saying that he may be targeted as "multilingual" and that he seems to be in a "good mood" today.

In each scenario, the user can click the browser extension icon or open the app on their phone to see each item in their ad-targeting profile (Figure 1). They can see the specific actions they performed, online or offline, that spawned each profile item, as well as example ads for each (Figure 2, left). Users can remove profile items or mark an item as sensitive, which will notify them every time a future activity triggers the same profile item.

3.3 System Implementation

We implemented our technology probe via a browser extension and a mobile app. Both capture signals to determine possible additions to the user's ad-targeting profile while also giving users continuous and real-time visibility of such additions. This approach simulates advertising networks' inferencing techniques and their categorization of user attributes into *potential* interests, demographics, and behaviors for targeting personalized ads. We emphasize the word *"potential"* because we do not have access to actual user profiles or ad networks' actual mechanisms, which are trade secrets. Instead, the probe tries to construct parallel profiles using state-of-the-art inferencing and machine intelligence technologies. For example, to connect text fragments to ad interests from Google's AdWords, we used Google's Universal Sentence Encoding [11] algorithm. In many ways, the probe reveals possible ways of profiling that *could* be used to target ads.

Browser extension. The WHO AM I browser extension serves two purposes: (1) making inferences from browsing activities and (2) providing a comprehensive visualization of the user's profile. For example, interests are inferred based on the content of the pages visited when using the browser extension. Users are notified about new additions to their profile when such inferences introduce a new, potentially sensitive item to their profile. We define sensitive items fully in Appendix A. The browser extension's visualization (Figure 1, left) shows all profile items originating either from the computer (web browsing data) or phone sensors. It also offers filters for subsetting the profile items. The visualization on the phone (Figure 1, right) does not include filters or items originating from the computer. This reflects a conscious choice not to overwhelm users with a long list of items and buttons on the phone's small screen. A sync feature was added so that new items from the phone could be visualized on the computer browser by simply opening both the app and the browser extension at the same time, after having initially paired them. Ad tags are only shown on the computer browser since the mobile app does not have permission to access webpages visited on the phone's browser or ads shown inside other apps. The browser extension was written using the React framework and the WebExtensions API. It was published as unlisted in the Chrome Web Store and also signed for distribution by the Mozilla Developer Hub to be used on Firefox. As instrumentation for the study, the browser extension can send the locally stored profile data to our research server (done manually, with the option to exclude individual items) and can also submit diary notes.

Mobile app. The mobile app captures offline actions, making inferences from physical locations visited, physical movement (e.g., driving, cycling, running), apps used, messages exchanged, and device charging activity. Notifications are again sent to users when new items are added to their profiles. On the phone app, just as on the computer, users can also see their profiles, details of profile items, and ad examples, but only for items originating from the phone. Filters are not available on the phone. The mobile app was developed on top of the React Native and AWARE [20] frameworks. The AWARE Framework is an open-source mobile context instrumentation framework that processes sensor data, allowing us to create rich data about the user's context as they use their phone and move through the world. For example, it provides data on locations visited, communications exchanged, activities, ambient noise, and device usage. Table 2 in the appendix details the inferences made through the AWARE framework. For the field study, we published our mobile app to both Apple's App Store and Google's Android Play Store using their beta-testing channels (TestFlight on iOS and Internal Testing on

Android). Due to restrictions of the iOS operating system about what data apps can or cannot capture, some speculative profile attributes (e.g., the user's mood from keystrokes in messaging apps, apps used on the phone, and searches performed on Google Maps) can only be captured on Android.

As part of the probe's development, an informal pilot was conducted with 10 friends and acquaintances. Pilot participants suggested adding ad examples, prioritizing notifications for sensitive items, showing connections between actual ads and profile items, and adding the ability to remove or mark as sensitive entire groups of items.

4 METHOD

In this section, we describe the methodology of our field study.

4.1 Probe Field Deployment

In our field deployment, participants installed the WHO AM I browser extension on their computer's primary web browser and the WHO AM I mobile app on their smartphone. They used both for a period of two weeks. Participants were divided into two groups: the *awareness* and *direct control* groups. The experience of both groups was the same except that participants in the *awareness* group could not make any changes to their profiles during the first week (e.g., removing items or marking them as sensitive). This was done to elicit the feeling of having no control over profiles during the first half of the study, letting us examine whether participants who saw their profiles being composed without being able to modify them would later change their profiles more than participants who could make changes from the beginning. The field deployment was conducted with US-based participants from October 21 to November 21, 2020. The study was approved by our IRB.

The field deployment started with a semi-structured entry interview to gather participants' initial perspectives on concepts of privacy, experiences with OBA, and the extent to which they feel in control of OBA. We also assisted them in installing the technology probe onto their computer's web browser and their smartphone.

At the middle and at the end of the study (after the first and second weeks, respectively), participants completed short questionnaires related to the four key user-facing problems of OBA identified in prior works: *engagement*, *interpretation*, *representation*, and *action*. These questionnaires let participants share their impressions related to the four problems once they had used the probe for a week or two. Appendix F contains our survey instruments, as well as additional detailed results.

The study also included a diary component in which we asked participants to log their thoughts and feelings as they experienced the technology probe, such as things they found surprising or interesting, or when particular OBA-related events occurred. The latter included seeing privacy-invasive, surprising, or inconsistent tracking and ads, or carrying out actions during which they wished they were not tracked. Participants submitted diary entries from the WHO AM I browser extension. Participants were not required to do so, but we offered an incentive of $1 (USD) per note submitted, up to $15 per participant. These notes served as a way of experience sampling.

During the two weeks of use, the probe logged events related to the composition of participants' profiles, along with their interactions with the probe. Logging was implemented in a privacy-preserving way. For instance, when new additions to participants' profiles were made, the logs did not indicate which items were added, but just the number of items added. At the end of the study, participants were required to submit their profile data, which remained stored on participants' devices for the duration of the study. At that time, they could exclude from the transfer any profile items they were not comfortable sending. The logs also captured fine-grained interactions with the probe, such as time spent, profile views, items investigated, filtering actions, and other user interactions.

Finally, participants completed an exit interview (Appendix E) aimed at gathering insights from their situated experience with the probe, understanding changes in behavior or attitudes that the probe may have motivated, and further explaining any behavior or diary notes captured during the study. Most importantly, the exit interview's main focus was to derive concrete design directions for future transparency and control mechanisms focused on

OBA profiling. At this point in the study, participants had gone through the two-week probe and were situated enough to provide meaningful accounts and suggestions. During the exit interview, participants were also invited to compare their probe experience with ad dashboards and explanations provided by Facebook [18] and Google [22] (even-odd counterbalancing). The interview concluded with several co-design questions about dealing with ephemeral items, inaccuracies, interests versus identities, missing pieces, and sensitive items in OBA profiling. For example, one such question was, *"You have probably noticed that there are mistakes on your profile. For example, inaccuracies and/or conflated identities, such as when doing something for work or sharing a device with others. Would you be interested in correcting those? Why or why not? If so, how might you want to do it?"*

Participants were recruited via several channels, including Craigslist (Chicago and Urbana-Champaign areas), Prolific, and Amazon Mechanical Turk. Purposive sampling was conducted in order to ensure the diversity of participant demographics, technology experience and expertise, and privacy concerns and considerations. A total of 27 participants were recruited for the study, and 25 completed the study. One of the two dropouts could not install the mobile app at the entry interview and was set to install it later, but never did. The other dropout was inactive and unresponsive for a few days, later reaching out to say they had gotten a new phone in the middle of the study. In both cases, we decided to compensate the two participants for the activities completed, yet discontinue their participation given the period of inactivity and non-compliance with the study requirements.

Appendix D shows participant demographics, along with the number of times they visited their probe profiles during the two-week study. The 25 participants came from diverse backgrounds in terms of age, education, gender, occupation, weekly internet use, mobile operating system used, and attitudes toward online privacy. Among participants, 14 used an Android phone and 11 used an iPhone. In the screening survey, participants reported using the Internet an average of 40.69 hours per week (Mdn=30, SD=28.67, Min=2, Max=135). Occupations were very diverse, with some examples being healthcare administrator, package handler, substitute teacher, librarian, and yoga teacher. Seven participants reported using ad blockers.

Participants' IUIPC scores [37] also varied. Summing answers to create a composite score, the minimum value was 37, Q1=50, Q2=55, Q3=63, and the maximum was 68, with the average being 55.72 (SD=8.38). Relevant to the topic of our research, the statement with the most variance in responses was *"I'm concerned that online companies are collecting too much personal information about me."* In contrast, *"Companies seeking information online should disclose the way the data are collected, processed, and used"* was the statement with the least variance. Participants were compensated with up to $95 (USD) for their participation. Participants who completed the whole study received $80, in addition to $1 for each diary note submitted, up to $15. Compensation was made via electronic Amazon gift cards sent to participants' preferred email addresses upon study completion.

4.2 Log Data

Logging of the participant's activities started upon successfully installing the browser extension and the mobile app. The purpose of collecting log data was to contextualize participants' qualitative responses and capture a more comprehensive picture of each user's participation in the study by triangulating their qualitative accounts with the quantitative signals obtained by the probe. In a way, the logs served as a rich grounding mechanism

Table 1. Events logged during the field study, ordered from most frequent to least frequent (top to bottom and left to right).

website visited on computer	left profile	entered item group on phone	ad tag shown
items added to profile	dragged time slider	opened item	hovered ad tag
items updated in profile	filtered out browser items	changed ad example	clicked ad tag
notification shown on computer	filtered out phone items	opened item actions	sent profile for sync
clicked notification on computer	filtered out non-sensitive items	removed item	sync successful
notification sent to phone	swiped ad example card on phone	marked item sensitive	
notification opened on phone	hovered ad example strip item on computer	marked item not sensitive	
entered profile	expanded item group on computer	ad detected on computer	

for each participant, providing evidence of participants' behaviors during the study. We designed the logging mechanisms to preserve participants' privacy, including only storing profiles locally until the end of the study, allowing participants to exclude profile items before submitting their final data, and not logging any keystrokes or actual profile items throughout the study. We communicated these aspects in the consent form and verbally at the entry interview. Table 1 shows the list of all events logged.

Accuracy influences perceptions of ad-targeting profiles [15]. Users may discredit algorithms if they find their profiles inconsistent with their actual identities or behaviors [17]. Thus, we recorded a proxy value for the confidence or accuracy of profile items whenever possible to track the consistency of our probe's profiles with the ads participants saw during the field deployment. For example, when our probe added new profile items from the similarity of webpage content with AdWords categories, we stored the underlying cosine similarity value. We also stored the cosine similarity for ads and ad tags shown, helping us gauge the accuracy of our profiles objectively, complementing participants' self-reports of perceived accuracy in the midterm and exit questionnaires.

4.3 Data Analysis

We approached data analysis using mixed methods. The anchor and emphasis of our analysis was on the qualitative data from exit interviews and diary notes. The quantitative analysis of log data and questionnaires provided contextual grounding for the qualitative evidence, such as showing participants' engagement with the probe. We also compared responses to the entry, midterm, and exit questionnaires via descriptive statistics. Diary notes were taken as situated accounts of participants' experiences in reacting to transparency and controls afforded by the probe. Whenever possible, we triangulated the main qualitative findings with quantitative data to highlight when participants' attitudes were particularly well-aligned with their behaviors, as indicated in the logs.

One researcher conducted all interviews. The interview sessions were recorded, professionally transcribed using Rev, and analyzed for recurring themes. During each interview, the interviewer took notes that captured the main point from participants' responses, in addition to attempting to capture what they said word-for-word. We conducted a thematic analysis by first reviewing all interview notes, then identifying recurring themes, grouping them, and ordering them by their frequency of occurrence. We also conducted thematic analysis for the 185 diary notes submitted. The same researcher who conducted the interviews first read all the notes to become situated with their content, then mapped recurring themes and identified how many times they were reinforced, uncovering the most reinforced themes. A summary of all codes can be found in Appendix G.

5 RESULTS

In this section, we present the results of our field study.

5.1 Probe Engagement

Engagement. In total, 92,675 system log entries were created. The average length of a profile visit was 44.27 seconds, and participants spent a collective total of about 35 hours engaging with their profiles. Over the two weeks, the average total time spent visiting profiles was 83.91 minutes per participant (Mdn=77.33, SD=31.07, Min=33.15, Max=159.7). The average number of minutes spent daily per participant was 8.43 (Mdn=7.74, SD=3.77, Min=2.55, Max=20.38). Participants visited their profiles 89.32 times on average (Mdn=94, SD=37.75, Min=20, Max=154). Each participant opened an average of 63.92 profile items (Mdn=51, SD=59.17, Min=5, Max=214), while the average number of items removed from profiles per participant was 7.36 (Mdn=2, Min=0, Max=46). Participants visited a total of 2,703 distinct websites on their computers (in the browsers instrumented with our extension) over the course of the study.

Figure 3 shows participants' daily engagement with their profiles. In this figure, *engagement* means the participant interacted with their profile visualizations (e.g., opened items, used filters), but did not remove or

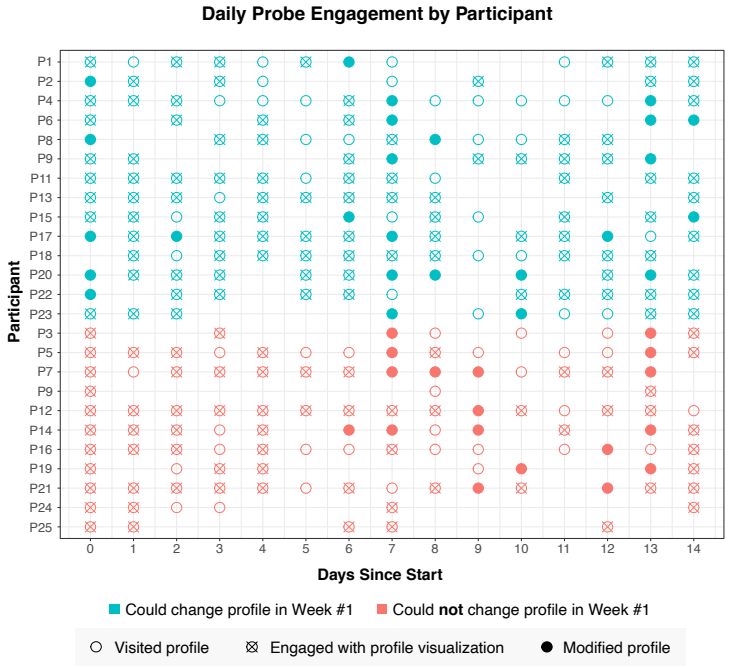

Fig. 3. Daily probe engagement. Profile modifications include both removals and sensitive/not-sensitive marks. Touchpoints were on Day 7 (a reminder about the midterm questionnaire and telling participants in the second group they could now make changes) and Day 13 (asking participants to answer the exit questionnaire and submit their profile, as well as reminding them of the exit interview happening the next day).

mark items. *Modification* means the participant also removed items, marked items as sensitive/not sensitive, or both. The logs indicate that most participants were very engaged with the probe, visiting and engaging with their profiles almost every day. Overall, the logs also indicate that most participants did not change their profile very much, and that participants made a comparable number of profile modifications regardless of their assigned group (whether or not they could make changes from the start).

These figures suggest that the probe's deployment was successful in engaging participants, especially considering that they were not required to visit or engage with their profiles as part of the study criteria. The small number of profile modifications (item removals and sensitive marks) also suggests that the utility of the probe was perhaps more on the side of transparency than on the side of control.

Submitted profiles. All 25 participants submitted their profiles at the end of the study. The average profile size was 257.08 items (Mdn=249, SD=123.67, Min=27, Max=504). Out of the 25 participants, eight participants excluded items from the final profile transfer, with the average number of items excluded being 11.5 (Mdn=8, SD=10.74, Min=1, Max=31). The average number of profile items that originated from participants' phones (as opposed to the web browser on their computer) was 51.84 (Mdn=44, SD=29.8, Min=15, Max=116), with the percentage of items from the phone being 25% on average (Mdn=22%, SD=16%, Min=4%, Max=68%). The average number of profile items spawned from participants' phone locations alone was 10.52 (Mdn=8, SD=10.06, Min=0, Max=32). Final profiles had an average of 49.76 sensitive items (Mdn=50, SD=30.03, Min=0, Max=132) and an average of 9.96 speculative items (Mdn=9, SD=5.11, Min=2, Max=27). The average percentage of speculative profile items was 5% (Mdn=4%, SD=3%, Min=2%, Max=11%).

How accurate were profiles? The log data also contained information about the strength of the relationship between profile items added from the computer browser and participants' activities. Specifically, the cosine similarity value was logged as part of the profile item records. The average value for this relationship across all profile items for each participant was 0.56 (Mdn=0.57, SD=0.01, Min=0.55, Max=0.59). This value is essentially a semantic similarity measure between the sentence embedding value vector from the contents of a given page (e.g., a combination of content keywords, page title, and hyperlink text) and the sentence embeddings of the AdWords category with the strongest similarity (see Appendix B for more details). Another proxy for profile accuracy was the relationship between probe profile items and the text in the ads shown to participants. The average cosine similarity between profile items and the text extracted (using Google's Cloud Vision API) from the actual ads shown to participants was 0.48 (Mdn=0.47, SD=0.08, Min=0.2, Max=0.9). One important note is that even with the ground truth on profiles, perfect similarity may not be attainable due to randomness and churn in ad selection.

A measurement of perceived accuracy was collected in the midterm and exit questionnaires. Specifically, participants were asked whether their probe profiles were an accurate representation of their perceived identities. At the midterm questionnaire, 16 participants responded with *"somewhat"* and three responded with *"to a great extent."* At the exit questionnaire, four participants responded with *"very little,"* 11 responded with *"somewhat,"* and nine participants responded with *"to a great extent."* Combined, these figures suggest the profiles created by the probe were reasonably accurate, and that perceived accuracy mostly increased the longer the probe was used.

Behavior explanations. When asked to explain their engagement with the probe, such as the reason for removing/not removing items or visiting their profiles frequently, some explanations were very elucidating. For example, 12 participants mentioned that they did not find any item on their profile to be sensitive or concerning, so they did not feel like they wanted to remove any items, and two participants reported not removing any items because they thought they were general enough to keep them. Another aspect related to item removals was inaccuracies: five participants said the reason they removed items is because they thought they did not reflect their identities or activities. However, three participants said they did not remove inaccurate items because they were not bothered by potential "mistargets." Finally, when asked to explain why they visited their profiles frequently and engaged with the time slider, filtering, and opening item details, participants reported monitoring the growth of their profile (five participants) and determining whether something sensitive had been added (two).

5.2 Situated Reactions to the Probe

A total of 185 diary notes were submitted by 23 participants. The average number of notes submitted per participant was 7.4 (Mdn=6, SD=6.1, Min=0, Max=22). The notes contain rich accounts of reactions to increased transparency about inferences made from participants' activities. Among the most common participant reactions to the increased transparency were instances of **identity reflection** (19 notes), such as disputing inaccuracies and seeing merged identities, and **algorithm disillusionment** (17 notes), such as participants learning the limitations of algorithms they originally thought to be perfect. Users might react this way when they can "see through" a previously opaque system, as noted in prior work [16, 17, 42, 51]. We present novel findings below.

Unexpected tracking. Surprises about profiling, such as phone locations being used to make inferences and target ads, were mentioned in 17 participant diary notes. For example, P14 wrote, *"I did not know that when using the GPS program such as Waze or Google Maps they can send targeted ads based on where you stop. I stopped nearby the In and Out burger and it said I may see ads based on Burgers."* For participants whose mental models of profiling originally centered around computer browsing activities, they were surprised to see that inferences could be made about them from activities captured away from their computers (e.g., locations visited) and for activities not related to online shopping, such as when visiting education-related websites. Diary notes discussing unexpected profiling also mentioned participants' desire for specific explanations. Often these were items that did not align with participants' mental models. P7 noted: *"I'm interested in the results I received when I used my*

navigation app. I would like to know why I got auto insurance results. Was it just because I was driving? Was it because I was driving at a certain speed?" Such surprises were more common for inferences made in seemingly offline contexts (e.g., driving, taking a walk). These surprises often led participants to regulate boundaries, as seen in 12 diary notes, and to remove profile items where the activities were considered too personal, such as activities captured from the phone (68 removals) and sensitive topics like health (31 removals) and finance (13 removals).

Identity performance. 14 diary notes focused on identity performance where the participants wanted their ad-targeting profiles to reflect certain aspects of their identities, such as their gender identity, social roles (e.g., stay-at-home parent), or interests. For instance, P5 wrote about his attempt to see his interests reflected on his profile: *"I've actively been trying to go to websites for my personal hobbies like Magic the Gathering, World of Warcraft, and wrestling, and few of those things are being picked up."* Notes about identity performance often reflected attitudes related to algorithmic authority and omniscience. Participants expected the algorithm to pick up on nuanced aspects of their lives or have a heightened sensory perception. P4 wrote, *"Another thing I noted today and yesterday was that I spent a chunk of time each day in my kitchen, and the app does not notice that, that I am cooking and baking."* Participants' thoughts of algorithms being omniscient were the subject of 15 notes.

Elucidation. 12 diary notes described how interacting with the probe helped elucidate the profiling process. For example, P10 wrote, *"Visited Office Depot to get a tape. Now I see [an] office equipment tag and ad on my Facebook page."* Similarly, 11 diary notes presented moments in which participants were unsure why certain items were on their profiles at first glance, even when seeing the originating actions. This led them to recall their actions and try to justify items on their profiles that could simply have been inaccuracies in profiling. For instance, P12 noted, *"Today's interesting misunderstanding of my identities was that I'm now, it seems, a cyclist, because I held my phone while cycling…I'm not sure what could confuse it–walking speed? The fact that my brother and I talked about bike commuters in cities? That I walked on a forest trail today marked for walkers and cyclists? I'd be really curious to know what goes into make these determinations."* People's tendency to justify algorithmic inaccuracies is a known reaction to increased transparency of opaque systems [17]. These reactions point to opportunities for generating positive outcomes in implementing profiling transparency, enabling users to engage with these processes in order to achieve desired outcomes and prevent undesired ones.

5.3 User Preferences

The user preferences reported in this section are recurring themes that emerged in analyzing the comparison and co-design exercises conducted at the exit interview. The comparison exercise had participants compare our probe against current Facebook [18] and Google [22] profiling transparency dashboards and ad explanations. The co-design exercise focused on participants' views on presenting ephemeral items, inaccuracies, interests *versus* identities, missing pieces, and sensitive items. Our analysis centered on gauging participants' desired engagement with profiling processes. The exit interview questions can be found in Appendix E.

5.3.1 Favorite features. Participants' favorite probe features were the **ability to trace back actions to profile items** (mentioned by 13 participants), **in-the-moment notifications** (8), the **profile structure** (7), and **the timeline** (6). Participants' thoughts when comparing the probe with existing approaches closely reflected these preferences: 13 participants wished existing profiles were more visible and that they were invited to contribute to such profiles in order to correct inaccuracies, prevent undesired experiences, and benefit from serendipitous discoveries. 10 participants wished they could trace profile items back to past activities, 13 participants wanted in-the-moment controls, and six participants wanted to be able to be profiled in a more general way. Users who oppose tracking may prefer more general ad-targeting profiles. For instance, when discussing YouTube ads, P7 noted, *"I'm of two minds there. The cynic in me is like, 'Keep them wrong, so they don't actually know and they can't track you too well,' but not utilitarian. The pragmatic me is like, 'They should be correct. If I'm forced to sit through these, they should at least be useful. Otherwise, they're just wasting [time]."'*

According to participants, these control opportunities should be better placed in the moment, as noted by P23: *"I think an invitation to review your profile, or when it happened, either would work really well. I wouldn't want it to wait until I physically looked at the profile. Because a lot of the time, you forget about having an actual profile like this. Having those notifications, I think, would be very helpful. To be like, 'Hey, we noticed this behavior that's irregular,' or, 'Hey, we noticed that you may want to review this activity. Do you want to include this in your profile, or would you like us to remove it?'"*

Concretely, participants suggested in-the-moment toggles, such as *"do not track this," "this is not for me"* (for conflated identities), *"transaction mode"* (for one-offs), *"work mode,"* and *"blocklisting"* certain topics or websites when carrying out activities that might lead to ad targeting. Participants wanted these choices for potentially sensitive activities, unexpected tracking, and unusual activities they performed. Beyond more direct choice, there was a desire to be profiled in more general ways, expressed by six participants, and best represented by P17's comment: *"I wish that there were more ways to just turn targeted ads off. I know that companies need to make money somehow, but maybe there could be a more generic ad than like these hyper targeted ones."*

These preferences sharply contrast with existing mechanisms, which some participants described as a form of "transparency theater" because dashboards are not easy to access and control is limited. In fact, only one of the 25 participants reported knowing about the dashboards from Google and Facebook they compared our probe to in the study. The alternative of shifting control to users could lead to more engaging experiences. For example, P23 said that controlling her profile more directly would lead to *"more tailored ads specific to things I like, instead of guessing; then they'd know for sure."* However, we do not take participants' stated intentions to engage at face-value, since in reality users may not be as interested in engaging with settings and controls. For this reason, we paid particular attention to what could motivate users to participate more actively in OBA profiling.

5.3.2 Reasons not to engage. Participants noted many reasons not to engage with transparency tools.

Just figure it out. A common expectation, expressed by 14 participants was that the algorithm should be able to "just figure it out" and determine whether something is timely and relevant for them. For example, participants expected that temporary interests or one-off activities (e.g., a small purchase or an unusual search) should not influence ads by default. Another common expectation was that profiling should be aware of world events and ephemeral life circumstances. For example, P11 noted, *"Okay, a great example is, in the US, we just had the election, and I searched things related to the election…My point is that someone who isn't actually interested in politics may make those searches just because an event that's happening then, and it's relevant, but that's not actually represented as a long-term interest."* All in all, participants hoped that profiling would take these nuances into account in order to mitigate undesired experiences, such as when no-longer relevant ads follow them around or world events or ephemeral life circumstances disproportionately influenced the ads they see.

The expectation for algorithms to account for these contextual factors largely came from participants feeling unmotivated or uninterested in regulating their ad-targeting profile when engaging with their primary activities, with a few exceptions. For example, six participants mentioned that they would like to receive a notification when the algorithm detects them making an unusual search, which can give them choice about whether they would want resulting inferences to inform future ad targeting. For example, P11 said, *"I think I guess on the phone, when a notification pops up about something, maybe being able to click on that and say…For example, if it pops up with, 'You're interested in politics because you searched for this,' clicking on it and saying that it was a one-off search or something like that. Or saying that it's confirming or denying I guess."*

Not keen on ads. Perhaps not surprisingly, participants' desired level of engagement with their profiles was influenced by their original attitudes toward, and experiences with, ad targeting. For example, participants who dislike and block ads are not as interested in engaging with their profiles as much as participants who accept them or even like them. For example, P2's reported experiences with OBA in the entry interview were mostly negative. He described being bombarded with irrelevant ads after searching for a product online. Later, when

asked about whether and how he might want to deal with potential inaccuracies, he said, *"I would not. Because information is getting spewed out there… if you can't block it, then perhaps another good way of going about it is just to overwhelm the system with irrelevant information."* He also explained why he would not be interested in manually adding missing pieces to his profile: *"I understand why ads are there, but I'm not necessarily friends with them. If I need something or want something, I'll drill into it myself."*

However, a common desire expressed by participants related to moments in which they would like to engage to prevent or mitigate undesired experiences. Undesired experiences included being "bombarded" with ads about something they did not buy for themselves, seeing an ad they can only trace back to a potentially private activity (e.g., a conversation with a friend), or seeing an ad for which the profiling process does not align with their mental models of what activities are being tracked. In addition, participants in the spectrum of "despising ads to liking them" might engage at different moments and with varying levels of effort. For instance, some would only engage with their profile after seeing an undesired ad, while others would appreciate being able to anticipate what ads they might get, even being willing to go out of their way to "fill in the gaps" in order to get more relevant ads.

5.3.3 Reasons to engage. Participants also mentioned numerous reasons to engage with transparency tools.

Being in the loop. Although there is an expectation for profiling algorithms to independently figure out context and relevance, 19 participants expressed that some aspects of their activities may inform how involved they want to be with the composition of their ad-targeting profile. Mainly, participants wanted to engage when carrying out activities that may not speak strongly to their identities (e.g., a one-off search, looking up or buying something for a friend) or topics and activities they deem sensitive (e.g., topics related to health or finances, private conversations, or activities they do not expect to become part of how they are targeted online). For such activities, participants would want to see a notification with a choice of whether they would like to be targeted based on an inference about that activity. For example, P7 mentioned that they wish to be given a choice whether the present activity might be turned into ads, saying, *"I think in an ideal world, that would be the way I would go about it, because I don't think I can completely get away from ads. I guess if there's something sensitive that's going on my profile, that can alert, because those would be the things that I'd be most likely to be like, 'Don't, don't, don't track this.' I think that would be the only time I'd want an active alert."*

Keeping users in the loop based on the nature and context of their activities and inferences might require profiling algorithms to become effective in detecting such moments. Those moments include anomalies in user activity, as well as activities and topics that might be considered sensitive for either a large number of users or specific users in particular. The qualitative data in this study suggests that mental models of tracking involve mostly online browsing and shopping activities. When inferences are made beyond such activities, users might want to be kept in the loop and exercise choice. The main reasons participants said they wanted to be in the loop and exercise control is to be able to prevent undesired tracking and ads. This finding further contributes to the idea that users want to control OBA experiences based on the information involved [55]. One challenge that might arise is in having algorithms be aware of the "must have" moments where users wish to anticipate and control profiling, given that this may vary based on a user's attitudes, perceived utility, and engagement with ads.

5.3.4 When and how to engage. Participants also noted opportunities for increasing engagement.

Algorithm feedback. Participants highlighted moments in which they would be willing to participate more actively, such as fixing inaccuracies to get more relevant ads or saying they want "more of this type of ad" upon seeing a relevant ad. When asked about whether and how they wish to fix mistakes and inaccuracies on their profile, participants reported mixed intentions. For example, while 13 participants claimed they would fix inaccuracies in order to more accurately reflect their identities or to get more value out of ads, 10 participants said they might to do so only if a particular type of ad starts bothering them.

Interestingly, five participants found inaccuracies on their profile to be an obfuscation mechanism, with P17 expressing concern over having an accurate profile. P17 noted that she might want to keep inaccuracies as a way

of control, saying, *"I like to think that I'm pretty savvy about trying to obfuscate and protect my digital privacy. So I don't know what other things I can do besides just Googling things I'm not interested in, like cycling."*

The main reasons for which participants wanted to correct their profiles related to identity performance, financial gain (e.g., seeing deals they would not know about), receiving more relevant ads, preventing potential embarrassment (in case of shoulder surfing), and when something potentially sensitive, yet inaccurate, was inferred. For example, P3 explained why she deleted an item: *"I guess the thing that it did pick up, I was uncomfortable being labeled as viewing something like that because I don't even remember searching something from that topic. I think it was addiction. I don't know why that came up. So, I removed it, because I was like that's not me."*

The desire to correct potentially sensitive inaccuracies emerged in many diary notes and explanations of past engagement with the probe during the exit interview. This finding aligns with prior studies that suggest users may be uncomfortable with both inaccuracies and sensitive topics [15]. Most importantly, participants' attitudes about inaccuracies and their motivation to engage with inaccuracies can perhaps be predicted by their general attitudes toward, and perceived utility of, OBA.

Opportunistic interventions. Asked whether and how they would have wanted the probe to pick up pieces of their identities that "flew under the radar," six participants mentioned that they wished their profiles had included key aspects of their identity, such as their gender identity or a strong interest in an activity or topic. For example, P1 noted, *"I attempted to sort of influence some of the information on my WHO AM I to pick up certain aspects that it wasn't picking up on, like specifically the fact that I'm gay. It would not acknowledge that in any way."*

Participants expressed interest in adding such items manually to their profile in order to influence their ad targeting or ensure their profiles reflected their identities accurately. In other instances, expressed by five participants, participants wanted to manually add items to their profile in order to optimize the outcome of ads. For instance, P10 explained, *"I would rather have direct communication with more of the ads I want to see, rather than it wanting to track my browsing history and showing me ads."*

The way in which participants wanted to do so varied. Some participants wanted to be able to manually enter new items into their profile, while others wanted to provide feedback directly via ads, such as to "provide an alternative" for an irrelevant ad or answer a questionnaire of interests. Two participants mentioned moments when they would like to be able to input profile items, such as when gaining a new interest or going through a life change. For example, P9 explained, *"If you're able to manually put it in, maybe you gained a new interest that day, or a week later after you started it, you could manually add a new interest."* At the same time, four participants who were mostly negative or neutral about ad targeting said they did not see any incentive to do so. P17 pondered, *"I wouldn't have an interest in doing it unless there was… some incentive to do it. Just intrinsically, I have no desire to be like, 'This is me, everyone, look.'"* While participants differed in how they might want to engage with profiling, an overriding theme is that managing their profiles should be more intertwined with the activities that originate items, as well as the resultant ads. These comments also express how users may prefer to participate more actively in profiling over just being subject to the tracking and guesswork that drives profiling processes.

6 DISCUSSION

Combining our quantitative data and the richness in participants' qualitative reflections, our probe deployment provided valuable insights into directions for designing increased transparency of OBA's underlying profiling in a world of pervasive computing. Using the probe for two weeks, participants were able to share situated perspectives on how they wish to see and control their ad-targeting profiles, as well as which moments or circumstances are the most important for them in doing so.

The main takeaway from this study is that users wish to be invited to review, and at certain moments to control, pieces of their ad-targeting profiles. Most participants were not aware of their inferred identities created and provided by Google or Facebook, and they wished such identities had been more prominent. Determining

opportunistic moments for intervention and exposure of ad-targeting identities is a promising direction. As seen in this study, such moments can be informed by the nature of the activity that spawns inferences (e.g., instant messaging, online search, location visits) or the nature of the inferences themselves that are being made (e.g., health-related, finance-related, hobby-related). Most importantly, when it comes to managing OBA experiences, our collective evidence (logs, questionnaires, interviews) suggest that *transparency is desired more than control, but that control is expected to be exercised at opportunistic moments. Algorithms themselves may need to determine these opportunistic moments*. In summary, the **design implications** we identified are the following:

- Intertwine profile composition with primary activities for anticipation, control, and trust
- Have algorithms detect and seize opportunistic moments for user participation in profiling
- Provide prominent and specific explanations for unexpected or potentially sensitive inferences
- Promote choice and efficacy of control as opposed to "transparency theater"
- Give users choice on profiling specificity as a compromise

6.1 Engaging Users in Profiling

Enable anticipation, regulation, and optimization. Not all users are expected to accept invitations to engage with their profile or participate in its composition. However, a few factors might motivate participation. For instance, participation might be motivated by the opportunity to set boundaries on what can or cannot be turned into ads in the first place, or to monitor for sensitive items one might want removed from ad targeting. The most engaged users might even want to correct inaccuracies or manually add items. This finding is in line with Seymour et al.'s [43] probe findings for smart, Internet-connected devices in the home. Their participants found and speculated about new control mechanisms, suggesting new approaches to address their own concerns. Allowing users to achieve these goals may lead to more trust in ad targeting and increased ad engagement.

How to do it? Ad explanations and ad-targeting profiles should come to the forefront of user interactions. Furthermore, control mechanisms should be more tightly integrated both with users' primary activities and the ads they see. In contrast, on the Google Personalized Ads dashboard, users cannot see the specific originating activity when looking up different items on their profile. When choosing to investigate further, they are taken to a page that includes all of their activities on Google services, without any connection to the profile items whatsoever. Therefore, more integration is necessary between people's identities, originating actions, and the ads that result. Intertwining profiling with primary activities may allow users to take preventive measures when doing an online search or visiting a physical location to define their own boundaries and prevent undesired outcomes like embarrassment, unintended disclosures about past behaviors (e.g., on shared devices), and irrelevant ads resulting from conflated identities. Engagement might vary nonetheless, with some users only acting when they are "bothered enough" by a specific ad. Other users might take a proactive role in managing ad outcomes. Some ways to give users control are in-the-moment toggles or "modes," such as a "do not track this activity" mode, "transaction" mode, "work" mode, or "not doing this for myself" mode. Users who are not motivated to engage with such anticipatory mechanisms might want to be invited biweekly or monthly to review new items on their profile, or to be given prominent explanations and controls upon seeing an ad that might trigger them to act.

"Online" and "offline" expectations: meeting users where they stand. In the diary notes, most accounts of unexpected tracking or surprises were due to inferences made from offline activities captured on the participant's phone. We believe this to be due to people having an "online-first" mental model when it comes to tracking and ad targeting, which led them to have stronger reactions to, and deeper interest in, profile items originating from offline contexts. For these items that originated offline, participants often had a strong reaction and wanted to regulate boundaries more often, in addition to giving the algorithm more authority than otherwise. For instance, in online contexts (e.g., website visits), participants could more readily speculate about mistakes and did not demonstrate much concern about items inferred because those inferences were somewhat expected. However,

in offline contexts, participants had more triggering reactions, often expressing that they did not know such inferences could be made or wished the inference were not made due to finding such profiling invasive. This finding further expands the knowledge of users' preferences with regard to the level of detail in explanations. For instance, a prior user study revealed users prefer "medium-level" explanations [15], yet the findings of our study suggest that the answer might actually depend on the context. Participants also were more likely to try to justify mistakes by thinking the phone could have significantly greater sensing capabilities than what might be true. We found these differences to be interesting because they surfaced with increased transparency, despite targeting users with ads based on app usage and location being an increasingly common practice.

How to do it? One could provide more specific explanations when users likely may not recall what may have triggered a certain ad (e.g., a one-off search), or when the originating action may not be a part of how users expect to be tracked and profiled. It could be that users would want very specific explanations for unexpected outcomes of profiling and ad targeting. Moreover, the fact that explanations are often a couple of clicks or taps away triggers unintended reactions from users when they see a highly accurate ad they think may have been triggered from a potentially private activity, such as a conversation with a friend. In such cases, participants are likely to use heuristics to explain the ad instead of clicking through to see the explanation offered. In some diary notes, participants wondered whether the explanation was hiding something (thinking the algorithm had more information), or whether their profile purposefully chose to hide certain items from them. This might be a common reaction in real circumstances as well, which can make satisfactory explanations unattainable, particularly if users come from a place of distrust and resignation regarding apps and services that run ads.

Intrusiveness as a driver of engagement. With ads becoming increasingly intrusive and pervasive, users' interest in understanding them and having them be accurate might increase. Ads are also increasingly blended into feeds consumed by users on Facebook, Instagram, Twitter, and other social networking apps. Free-to-play games rely on showing video ads as a source of revenue. This means ads are becoming harder to escape, which also pushes users to opt into paying for a premium version of services, such as Spotify or YouTube Premium. Our findings suggest that even users who want to escape ads would prefer ads to be accurate if they have to sit through them when using online services. Essentially, if ads are to become ever more prominent in user interactions, it may make sense both for service providers and users to implement feedback, transparency, and control mechanisms that make the experience more worthwhile.

Inaccurate *and* sensitive. The most expressive reactions participants had – triggering them to act – were when they considered an inference as both sensitive and inaccurate. Participants reported removing items they thought were inaccurate and sensitive, such as addiction, smoking, or sleep patterns. For the most part, participants' emphasis and reactions to inaccurate profile items revolved around sensitive activities. When an inference was potentially sensitive yet accurate, participants' often wanted to investigate the inference, such as how it was made and what data was involved. Either way, potentially sensitive inferences seem to motivate engagement because users would like either to remove them (if they are inaccurate) or to investigate them (if they are accurate).

6.2 The Challenges Ahead

Feasibility of specific explanations. If ad networks are to show specific and complete explanations, it may mean that they have to store such data for accountability purposes. One question that emerges relates to feasibility: given the highly distributed places in which tracking is performed, is it even possible for companies to determine which specific actions have triggered a certain targeted ad, let alone enforce user controls? One paradox that might exist is that for the most accurate predictions of user intentions, machine learning models that made the underlying inference may not be interpretable at all, yet these might be the moments in which users demand satisfactory and complete explanations. Although such contextual and specific information could be collected for

accountability purposes, its possession also carries risks, including data breaches, privacy violations through unintended disclosure, and secondary uses of the data. For example, Facebook used phone numbers collected for two-factor authentication to target users with ads [36]. Just as litigation followed that revelation, policymakers might need to intervene regularly to enforce processes preventing secondary use of information collected for the accountability of automated processes. With highly decentralized tracking and profiling being done by multiple companies, it may be unmanageable for users to exercise control, or even to see that their controls work. This may also be a problem relevant to the field of Explainable Artificial Intelligence (XAI), where highly accurate models are known to be less interpretable. The distributed nature of profiling may compound this problem. However, even limited explanations should be attempted if users are able to provide feedback. In a recent user study of explanations for decisions made by machine learning models, the opportunity to provide feedback to explanations was found to be equally (or even more) important than the original explanation provided [45]. An additional question that emerges for highly specific explanations is one of authentication; ad networks must ensure that only the user who originally performed an action is shown a specific explanation mentioning that past action. This problem of authentication is particularly thorny for shared devices.

Impressions of "transparency theater." In comparing the probe's design with ad-interest dashboards and ad explanations offered by Google and Facebook, participants were dissatisfied with the vagueness and generality of the information these existing tools presented to them. This finding aligns with prior work identifying that users often consider a vague and incomplete explanation a non-explanation [4]. It may be argued that with the state of tracking and profiling, it is in the best economic interests of companies to keep it that way, but mounting evidence suggests that the current approach is not helping users build trust in, nor feel comfortable with, ad targeting. Participants in our study thought that while these tools exist, they give only the illusion of control and make it seem like service providers do not really want users to engage with their ad-targeting profiles. This attitude might reinforce users' distrust of Internet companies and new technologies, such as Internet-connected devices for the home [47]. Users may need to start seeing their choices enforced. For example, some participants mentioned that if they spent time and effort to control their ad-targeting profile, they would want to see that their actions would actually have some effect. Often, doubts about the efficacy of control were introduced by the very control affordances provided by Facebook and Google. In its user-facing documentation, Facebook says, *"Removing yourself from an interest category… doesn't affect the number of ads you see overall. We may still show you ads related to these categories if we think these ads may be relevant to you."* Trying to turn off certain interests on Google's dashboard results in a message saying, *"Advertisers you turn off will stay off for at least 90 days,"* rather than forever. These types of communications cast further doubt about the efficacy of actions users take to control their ad-targeting profiles. Addressing this friction is important because, after all, OBA makes free Internet services and apps possible. Users may indeed benefit more from relevant and personalized ads, but tensions are running high around the extent and the means of tracking to sustain such practices. The lack of transparency around these practices may result in measures aimed at completely preventing profiling, with one example being Apple's recent iOS update.[5] This update prompts users to consider whether or not to allow cross-app tracking. It is possible that continuing to track and profile users in obscure ways will not only increase distrust, but also make ad experiences less worthwhile for users if profiling were no longer possible.

6.3 Limitations

Our probe's inferencing mechanism is a local substitution model aimed at simulating, to some extent, what inferences can be made from user behavior. Participants' actual ad-targeting profiles might differ from the probe's. However, when asked to compare their actual profile data from companies' dashboards with our probe's inferred interests, participants did not find the information to be vastly different. However, their actual profiles had more

[5] https://developer.apple.com/app-store/user-privacy-and-data-use/

items, which participants mentioned was expected given decades of Internet use. Their actual profiles were also more specific. Given that the probe is only a substitution model, participants' level of control and actions (e.g., removing items) might differ if the study were conducted with their actual profiles. For example, some participants reported in the exit interview that they did not remove any items from their profile because they wanted the researchers to have an accurate picture of what the probe had captured. This is a limitation of the methodology, yet it was mitigated in the exit interview via probing participants to situate themselves within the context of their actual ad-targeting profiles.

7 CONCLUSION

Profiling processes in OBA are increasingly pervasive and sophisticated. The picture of people's behaviors, interests, likes, and needs is progressively enhanced by newly introduced devices and increased inferencing capabilities. While such processes have become a part of people's everyday online experiences, they are also still very opaque to users. Transparency can offer benefits to users, but the reality of OBA is that user acceptance and utility can be mixed. It can also be said that many users are unmotivated to control their ad-targeting experiences because doing so would be a secondary, and often intrusive, part of their online activities. With this in mind, we designed and deployed a technology probe with 25 Internet users based in the US. The probe provided incremental and in-the-moment visibility into inferences that could be a part of participants' ad-targeting profiles as they went about their lives for two weeks. The probe also provided details and controls identified as desirable in prior work. We found that participants wished to be invited to review, and at certain moments also to control, pieces of their ad-targeting profiles. Most participants were not aware of their inferred identities as created by Google or Facebook, and they wished such identities had been more prominent and integrated with their primary activities. While more transparency was desired, participants also wished to have choices and the ability to anticipate whether certain aspects of their activities could be used for ad targeting. They especially wanted control over potentially sensitive activities/topics and unexpected profiling, particularly in offline contexts. These results illuminate directions that will inform the future design of transparency and control mechanisms within increasingly pervasive and ubiquitous tracking and profiling environments.

ACKNOWLEDGMENTS

This research was funded in part by a Facebook Privacy Preserving Technologies research award. In addition, this material is based upon work supported by the National Science Foundation under Grant No. CNS-2047827. We thank Michael Twidale and Sauvik Das for their feedback on this work, as well as our anonymous reviewers for their thoughtful comments and suggestions. We thank the participants in our field study for their contributions and insightful conversations, as well as our pilot participants for their time and effort spent testing earlier versions of the technology probe and providing meaningful feedback.

REFERENCES

[1] Gunes Acar, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. The web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 674–689.

[2] Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. 2013. Do not embarrass: Re-examining user concerns for online tracking and advertising. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*. 1–13.

[3] Athanasios Andreou, Márcio Silva, Fabrício Benevenuto, Oana Goga, Patrick Loiseau, and Alan Mislove. 2019. Measuring the Facebook advertising ecosystem. In *Proceedings of the Network and Distributed System Security Symposium*.

[4] Athanasios Andreou, Giridhari Venkatadri, Oana Goga, Krishna Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Investigating ad transparency mechanisms in social media: A case study of Facebook's explanations. In *Proceedings of the Network and Distributed System Security Symposium*.

[5] Anja Bachmann, Christoph Klebsattel, Matthias Budde, Till Riedel, Michael Beigl, Markus Reichert, Philip Santangelo, and Ulrich Ebner-Priemer. 2015. How to use smartphones for less obtrusive ambulatory mood assessment and mood recognition. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*. 693–702.

[6] Muhammad Ahmad Bashir, Umar Farooq, Maryam Shahid, Muhammad Fareed Zaffar, and Christo Wilson. 2019. Quantity vs. quality: Evaluating user interest profiles using ad preference managers. In *Proceedings of the Network and Distributed System Security Symposium*.

[7] BBC News. 2017. Is your phone listening in? BBC. https://www.bbc.com/news/technology-41802282

[8] Victoria Bellotti and Abigail Sellen. 1993. Design for privacy in ubiquitous computing environments. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work*. 77–92.

[9] Alastair R. Beresford and Frank Stajano. 2003. Location privacy in pervasive computing. *IEEE Pervasive Computing* 2, 1 (2003), 46–55.

[10] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. I always feel like somebody's watching me: Measuring online behavioural advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*. 13:1–13:13.

[11] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. Universal sentence encoder for English. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 169–174.

[12] Julie E. Cohen. 2000. Privacy, ideology, and technology: A response to Jeffrey Rosen. *Georgetown Law Journal* 89 (2000), 2029–2045.

[13] Lorrie Faith Cranor. 2004. I didn't buy it for myself. In *Designing Personalized User Experiences in eCommerce*. 57–73.

[14] Leyla Dogruel. 2019. Too much information!? Examining the impact of different levels of transparency on consumers' evaluations of targeted advertising. *Communication Research Reports* 36, 5 (2019), 383–392.

[15] Claire Dolin, Ben Weinshel, Shawn Shan, Chang Min Hahn, Euirim Choi, Michelle L. Mazurek, and Blase Ur. 2018. Unpacking perceptions of data-driven inferences underlying online targeting and personalization. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 493:1–493:12.

[16] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I "like" it, then I hide it: Folk theories of social feeds. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2371–2382.

[17] Motahhare Eslami, Sneha R. Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating algorithmic process in online behavioral advertising. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 432:1–432:13.

[18] Facebook. 2021. Facebook ad preferences. https://www.facebook.com/adpreferences/ad_settings.

[19] Facebook. 2021. Facebook advertising targeting options. https://www.facebook.com/business/ads/ad-targeting.

[20] Denzil Ferreira, Vassilis Kostakos, and Anind K. Dey. 2015. AWARE: Mobile context instrumentation framework. *Frontiers in ICT* 2 (2015), 6.

[21] ftramer. 2020. ad-versarial. https://github.com/ftramer/ad-versarial.

[22] Google. 2021. Google ad settings. https://adssettings.google.com/authenticated.

[23] Google. 2021. Targeting your ads. https://support.google.com/google-ads/answer/1704368?hl=en.

[24] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking data out of context to hyper-personalize ads: Crowdworkers' privacy perceptions and decisions to disclose private information. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[25] Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic understanding of image and video advertisements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1705–1715.

[26] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: Inspiring design for and with families. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*. 17–24.

[27] Xiaodong Jiang, Jason I. Hong, and James A Landay. 2002. Approximate information flows: Socially-based modeling of privacy in ubiquitous computing. In *Proceedings of the 4th International Conference on Ubiquitous Computing*. 176–193.

[28] Yucheng Jin, Karsten Seipp, Erik Duval, and Katrien Verbert. 2016. Go with the flow: Effects of transparency and user control on targeted advertising using flow charts. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*. 68–75.

[29] Harmanpreet Kaur, Harsha Nori, Samuel Jenkins, Rich Caruana, Hanna Wallach, and Jennifer Wortman Vaughan. 2020. Interpreting interpretability: Understanding data scientists' use of interpretability tools for machine learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[30] Matthew Kay, Eun Kyoung Choe, Jesse Shepherd, Benjamin Greenstein, Nathaniel Watson, Sunny Consolvo, and Julie A. Kientz. 2012. Lullaby: A capture & access system for understanding the sleep environment. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*.

[31] Byoungjip Kim, Jin-Young Ha, SangJeong Lee, Seungwoo Kang, Youngki Lee, Yunseok Rhee, Lama Nachman, and Junehwa Song. 2011. Adnext: A visit-pattern-aware mobile advertising system for urban commercial complexes. In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*. 7–12.
[32] John Krumm. 2010. Ubiquitous advertising: The killer application for the 21st century. *IEEE Pervasive Computing* 10, 1 (2010), 66–73.
[33] Marc Langheinrich. 2001. Privacy by design — Principles of privacy-aware ubiquitous systems. In *Proceedings of the International Conference on Ubiquitous Computing*. 273–291.
[34] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–31.
[35] Scott Lederer, Anind K. Dey, and Jennifer Mankoff. 2002. *A conceptual model and a metaphor of everyday privacy in ubiquitous computing environments.* Technical Report. Computer Science Division, University of California.
[36] Natasha Lomas. 2018. Yes Facebook is using your 2FA phone number to target you with ads. Tech Crunch. https://techcrunch.com/2018/09/27/yes-facebook-is-using-your-2fa-phone-number-to-target-you-with-ads
[37] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research* 15, 4 (2004), 336–355.
[38] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *Proceedings of the Telecommunications Policy Research Conference*.
[39] Suman Nath, Felix Xiaozhu Lin, Lenin Ravindranath, and Jitendra Padhye. 2013. SmartAds: Bringing contextual ads to mobile apps. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*. 111–124.
[40] James Pierce. 2019. Smart home security cameras and shifting lines of creepiness: A design-led inquiry. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 45:1–45:14.
[41] Emilee Rader and Rebecca Gray. 2015. Understanding user beliefs about algorithmic curation in the Facebook news feed. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 173–182.
[42] Emilee Rader, Samantha Hautea, and Anjali Munasinghe. 2020. "I have a narrow thought process": Constraints on explanations connecting inferences and self-perceptions. In *Proceedings of the Sixteenth Symposium on Usable Privacy and Security*. 457–488.
[43] William Seymour, Martin J. Kraemer, Reuben Binns, and Max Van Kleek. 2020. Informing the design of privacy-empowering tools for the connected home. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
[44] Vidhi Kirit Shah. 2020. *User acceptance of online tracking if 'forgetting' was an option.* Ph.D. Dissertation. Carleton University.
[45] Alison Smith-Renner, Ron Fan, Melissa Birchfield, Tongshuang Wu, Jordan Boyd-Graber, Daniel S. Weld, and Leah Findlater. 2020. No explainability without accountability: An empirical study of explanations and feedback in interactive ML. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
[46] James Stables. 2021. Facebook is working on a smartwatch – and it could be Google powered. Wareable. https://www.wareable.com/smartwatches/facebook-smartwatch-2022-release-8304
[47] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2014. Intruders versus intrusiveness: Teens' and parents' perspectives on home-entryway surveillance. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 129–139.
[48] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*. 1–15.
[49] Max Van Kleek, Reuben Binns, Jun Zhao, Adam Slack, Sauyon Lee, Dean Ottewell, and Nigel Shadbolt. 2018. X-ray refine: Supporting the exploration and refinement of information exposure resulting from smartphone apps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 393:1–393:13.
[50] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinshel, Michelle L. Mazurek, and Blase Ur. 2020. What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users' own Twitter data. In *Proceedings of the 29th USENIX Security Symposium*.
[51] Ben Weinshel, Miranda Wei, Mainack Mondal, Euirim Choi, Shawn Shan, Claire Dolin, Michelle L. Mazurek, and Blase Ur. 2019. Oh, the places you've been! User reactions to longitudinal transparency about third-party web tracking and inferencing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 149–166.
[52] Mark Weiser. 1993. Some computer science issues in ubiquitous computing. *Commun. ACM* 36, 7 (1993), 75–84.
[53] Mark Weiser. 1999. The computer for the 21st century. *ACM SIGMOBILE Mobile Computing and Communications Review* 3, 3 (1999), 3–11.
[54] Yoram Wurmser. 2019. US time spent with mobile 2019. eMarketer. https://www.emarketer.com/content/us-time-spent-with-mobile-2019
[55] Yaxing Yao, Davide Lo Re, and Yang Wang. 2017. Folk models of online behavioral advertising. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 1957–1969.
[56] Shoshana Zuboff. 2015. Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology* 30, 1 (2015), 75–89.

A PROBE TRACKING AND INFERENCING

Table 2. Summary of sources of inferences and mechanisms used in the probe. Inferences are considered speculative if it is not clear from our research that Google or Facebook provides a way for advertisers to target users based on them. Open-source algorithms were used for semantic similarity with Google AdWords, sentiment analysis, and language detection. Details on heuristics can be found on Appendix B. * indicates an Android-only feature.

Source	Inferences about	Mechanism	Speculative?
PC browsing	Interests	Google AdWords	
	Browser and OS used	Native APIs	
	ZIP code	IPStack API	
Phone location (AWARE)	Interests	Google, Yelp, Overpass	
	City, state, county	Overpass API	
	Income	Yelp price levels	
	Travel activity	Heuristic	
	Local weather	OpenWeather API (AWARE)	✓
Phone screen (AWARE)	Device usage	Native APIs	
Phone motion (AWARE)	Physical activities	Native APIs	✓
	Commute mode	Native APIs	
	Interests	Heuristics	
	State of mind	Heuristics	✓
Phone hardware	Device info	Native APIs	
Phone charging (AWARE)	Sleep patterns	Heuristic	✓
Phone apps (AWARE)*	Products and services	Native APIs	
	Interests	Play Store app category	
Phone keystrokes (AWARE)*	Mood	Sentiment analysis	✓
	Languages spoken	Language detection	
	Multicultural affinity	Heuristic	
	Destination	Place Autocomplete API	✓

Activities tracked: On the computer browser, webpages visited by users are tracked in order to make inferences about interests users may have. Specifically, each page's title and a summary of its contents are used for inferences. Upon first using the probe, the IP address is also used to determine the user's ZIP code, state, and city. The probe also collects the browser and operating system used on the computer.

The following are the activities tracked on the mobile app:

- **Location (via AWARE):** The probe obtains the latitude and longitude values every 2 minutes and 30 seconds and records when the user has made significant movement (e.g., more than about 100 meters or about 328 feet).
- **Device usage (via AWARE):** The probe obtains device usage from the duration of user sessions, between timestamps of when the screen is turned on and off.
- **Physical motion (via AWARE):** Motion data is obtained from native APIs that detect human activity (e.g., stationary, in a vehicle, walking, cycling, running). Physical activity is tracked every two minutes.
- **Device information:** The probe records the brand of the device's manufacturer and the phone's OS.
- **Local weather information (via AWARE):** The probe queries OpenWeather using the phone's location coordinates. Information includes minimum and maximum temperatures, as well as weather conditions like cloudy, clear, rainy, and snowy.

- **Device charging (via AWARE):** The probe obtains time measurements of the duration of device charging sessions in the evening.
- **Apps used (via AWARE, Android-only):** The probe records information about other apps used by users, such as their name and distribution package or ID.
- **Keystrokes (via AWARE, Android-only):** Keystrokes are captured from instant messaging apps (Messages, WhatsApp, Facebook Messenger, and WeChat) in order to locally infer the user's mood. Keystrokes typed into Google Maps are also used to infer possible user destinations.

The mobile app developed for the probe had the AWARE Android [6] and iOS [7] libraries as dependencies, using sensor event broadcasts to obtain sensor data. The app then registered for sensor broadcasts from the AWARE sensor layer. Upon receiving sensor data from the AWARE library, the app made inferences. For example, upon receiving the latitude and longitude values from the AWARE location sensor broadcast, the app used the Google Maps Web API for reverse geocoding and the Google Places, Yelp, and Overpass Web APIs for POI details.

Inferences made. Most inferences that can be added to a participant's ad-targeting profile originate either from website visits on the computer or from phone sensors. Interests added from browsing activity on the computer each map to an exact interest in the Google AdWords categorization. They are therefore representative of "online" behavior. Interests drawn from location sensing are mapped from Google Maps, Yelp, and Overpass APIs via reverse-geocoding and radius search of points of interest. A smaller number of profile items are mapped based on heuristics, such as the hours a device was charged in the evening hinting at the number of hours of sleep, the sentiment of words entered into messaging apps hinting at a person's mood, and searches made on Google Maps hinting at a person's immediate intentions.

Profile items traced from phone activity are representative of "offline" behavior. Some of the profile items are speculative in that they are (likely) not considered in ad-targeting today, but they are technically feasible to infer and could be used for targeting in the future. These were mainly used to elicit participants' responses on unexpected and potentially sensitive activities informing their ad-targeting. For example, the probe attempts to infer a user's mood from performing on-device sentiment analysis on recent chat messages. It attempts to infer a user's lifestyle from physical movement, such as whether a person is physically active or their mode of commute (e.g., cycling, vehicle). Studying these inferences also enabled the unearthing of a kind of "playbook" in the ever-expanding boundaries of profiling. That is, these factors capture potential user reactions to unexpected activities being tracked and profiled. Table 2 summarizes inferences and whether or not they are speculative.

From the activities tracked by the browser extension and mobile app, inferences are made and mapped to three different sections of the profile: "you may see ads about" (containing potential user interests), "you may be ad-targeted as" (containing potential user demographics and behavior), and "hints about your intents today" (containing speculative targeting attributes). The wording for these sections was framed in a way that alludes to the utility of profiling rather than labeling sections in a descriptive manner, as suggested by prior works on how people perceive OBA inferences [42]. The mapping to different profile items was inspired by the different attributes Facebook and Google provide to advertisers when targeting users. For example, inferences were made about pages visited on the user's computer and measured for similarity against Google's AdWords list of interest categories. From the list of custom targeting based on demographics and behaviors on Facebook, inferences were made about the user's location, device usage, and apps used on the phone.

Sensitive items. Some profile items are marked sensitive by default. For example, if the comfort value (as measured by Weinshel et al. [51]) of the selected AdWords category is negative, the item is marked as sensitive by default upon insertion. If either the top-level category or any of the lower-level items have a negative value for comfort, a notification is shown to the user on the browser about the item being added to their profile. Only

[6] https://github.com/denzilferreira/aware-client
[7] https://github.com/tetujin/AWAREFramework-iOS

notifications involving potentially sensitive items are shown on the browser in order to avoid overwhelming users with notifications at every page visited. These AdWords interests are added under the "you may see ads about" section. A list of items originating from location data could also be marked as sensitive by default. Whether to mark such data as sensitive was decided by comparing the tags from Yelp Fusion and Google Places APIs via semantic similarity with sensitive AdWords categories, manually inspecting potential matches. From manual inspection, a threshold was picked to decide whether the item should or should not be marked as sensitive. For example, "doctors" in Google AdWords and "doctors" on Yelp are a perfect match, but "pharmacies" and "drug stores" are not a perfect match, yet highly related. Perfect matches were automatically tagged as sensitive, while the highest scores were ordered and manually inspected until the relationships became weaker, and a threshold was chosen: 0.7. If the cosine similarity between the closest sensitive AdWords interest and the tag was greater than 0.7, the tag was marked as sensitive. This process resulted in 149 Google and Yelp tags being marked as sensitive by default based on their relationships with the sensitive Google AdWords categories.

The ultimate goal of these inferencing processes is to identify ad-targeting possibilities and show users what those might be. In some ways, the probe simulates the individual user's subset of items in targeting lists provided by Google or Facebook to advertisers [19, 23]. Importantly, these inferences are at best a local substitution model for the ground truth, with efforts made to mimic the underlying process as best as possible while keeping room for exploration (e.g., in the case of speculative attributes). More details about the inferencing processes are presented in the next section.

B TECHNICAL DETAILS - PROBE INFERENCING

This section explains in detail how the different inferences are made by our probe system.

From the pages visited on the computer browser, an inference is made about potential interests by identifying words in the Google AdWords list[8] that are most closely related to the title and content of the page visited by the user. This relationship is determined by a calculation of the AdWords category with the highest cosine similarity from sentence encoding embeddings between AdWords categories and the title and a summary of the content of the page (i.e., a Word2Vec approach). Upon selecting the category with the highest similarity, items in its hierarchy are also added. For example, if the highest similarity category is "thriller, crime & mystery films" and its hierarchy is "arts & entertainment > movies > thriller, crime & mystery films," each item in the hierarchy leading up to the selected item is also added as a profile item. On the profile visualization, "arts & entertainment" becomes a major group, which can be expanded to find both "movies" and "thriller, crime, mystery films" inside. This means that if "movies" is instead the most closely related item, "arts & entertainment" and "movies" would be added. The similarity calculation is performed server-side, where the sentence embeddings for the AdWords interests are stored and pre-processed for fast computation. An inference is triggered only when the cosine similarity of the most closely related item is greater than 0.5. If the comfort value (as measured by Weinshel et al. [51]) of the selected AdWords category is negative, the item is marked as sensitive by default upon insertion. If either the top-level category or any of the lower-level items have a negative value for comfort, a notification is shown to the user on the browser about the item being added to their profile.

From physical locations visited, a reverse geocoding request is sent to the Google Maps API, from which a point of interest is captured. The point of interest is further looked up on the Google Places API in order to collect its types, such as "art gallery," "bakery," or "doctor." Once the point of interest is identified, the Yelp Fusion API is also used to collect more specific tags from the point of interest, such as "vegetarian food" or "coffee & tea," in addition to the price point for each point of interest, used as an inference associated with a person's income or purchasing power. The Overpass API is also used to detect any underlying cuisines, diets, and franchise operators for a given point of interest. All of these are checked for duplicates and added as profile items under the "you

[8] https://developers.google.com/adwords/api/docs/appendix/verticals

may see ads about" profile section. In order for an inference about a visit to a point of interest to be made, at least two requests to the reverse geocoding API have to return the same point of interest, which means that two requests separated by two and a half minutes have to map to the same point of interest. This process tries to prevent detecting points of interest that the user just happens to drive or walk by. Locations were also used to make inferences about a user's travel habits, adding "traveler" under "you may be ad-targeted as" if the location sensor detects the user in a new city. When traveling, the item "visiting [city name]" is also added.

From a running sum of the length of time users use the phone on a given day, the inference "heavy smartphone user" can be added to a profile under "you may be targeted as" once they use their phone for more than three hours on a single day, which is the average daily phone usage for US-based users [54].

From physical motion, inferences are made about users' mobility preferences. The following items can be added to the "you may see ads about" section upon a related activity detection: "vehicle commuter," "runner," "physically active" (once the user walks for a number of minutes), and "cyclist." In addition, once the "runner," "cyclist," or "vehicle commuter" inferences are made, "sports apparel," "cycling apparel" and "auto insurance" are respectively also added under "you may see ads about." Finally, "exercised today" is added under "hints about your intents today." With activity tracking happening every two minutes, the inferences above were only made after detecting a streak of three readings, such as cycling being detected for three consecutive readings. This again was a heuristic implemented to mitigate items that could be added because of sensing errors or from short-natured physical effort, such as a quick run to catch the bus.

Upon app installation, the phone manufacturer's brand and the phone's operating system are added to "you may see ads about" and "you may be ad-targeted as," respectively.

From local weather information, as provided by the AWARE framework's integration with the OpenWeather API, maximum and minimum temperatures and weather conditions are recorded. When detecting maximum temperatures over 30C (86F) or minimum temperatures under 8C (46.4F), "hot weather where you are" and "cold weather where you are" are respectively added to the "hints about your intents today" section. The OpenWeatherAPI also provides a succinct description of weather conditions, such as "clear" or "heavy rain," which was added as "[condition] where you are" under "hints about your intents today."

From tracking device charging activity, an inference is made about the user's sleep activity based on the following heuristics. If the charging session starts between 9pm and 2am, and it also ends between 5:01am and 12:59pm, and if the charging session lasted over four hours, the length of the session was calculated. If the length equals to or exceeds seven hours, the item "had enough sleep" is added to the "hints about your intents today" section. If it is under seven hours, "had little sleep" is added.

From apps used on Android, the app's ID is sent to an unofficial Google Play API[9] to gather the app's name and category. This category is added under "you may see ads about." For example, using the Chase Banking app leads to "finance" being added under "you may see ads about," and "Chase Banking user" under "you may be ad-targeted as."

From keystrokes on Android, keyboard input is captured to (1) predict the user's language or languages spoken, (2) predict a user's mood from sentiment analysis of the words entered in messaging apps, and (3) predict a user's potential destination from terms entered into Google Maps. The inferences made from the messaging apps are all performed locally on the mobile device using open source libraries for language detection[10] and sentiment analysis[11]. A Place Autocomplete API search is performed with the terms entered into Google Maps in order to detect the nearby search results for the term and make inferences like "looking for auto parts store" under "hints about your intents today."

[9] https://github.com/facundoolano/google-play-scraper
[10] https://github.com/wooorm/franc
[11] https://github.com/thisandagain/sentiment

C ADDITIONAL PROBE FEATURES

Item details: Upon clicking or tapping a profile item, users can see a description of the related actions that led to the addition of the profile item. For example, they will see that "health > pharmacy" was added because they visited a page on `cvs.com`, where the title of the page they visited is also visible, in addition to the date on which the item was added to the profile. If the item came from a location visit, users will see the name of the point of interest, along with its physical address. On the details pop-up, users will also see ad examples related to the item they are seeing, which are explained in more detail below.

Ad examples: Pilot testers thought that, besides seeing the items on their profile, it would be nice to see a visual representation of their profiles in the form of ads that the profile may help inform. Accordingly, this feature was implemented based on a manually labeled open data set of ads [25]. For each item, ad examples were gathered based on the semantic similarity of the item's title with a category of an ad, and also based on some heuristics for speculative items. For example, when detecting that the user may have had little sleep the night before, the item "had little sleep" is added under the section "hints about your intents today." In this case, ad examples showing coffee and energy drinks are shown. Similarly, if the probe detects hot weather at the user's location, the ad examples showed cold drinks and ice cream.

Ad tags: Because the probe's profile is a local substitution model (i.e., it does not capture users' actual ad-targeting profiles), ad tags were implemented to attempt to show the relationship between items on a user's profile and actual ads users saw online when visiting different websites. For each ad detected, the probe showed the item on the user's profile most closely related with the ad based on a semantic similarity calculation considering the profile items and the text recognized on the ad (via image OCR). This was both a way to integrate the probe with the profiling outcome – ads– and to gauge the similarity of items on the profile with the actual ads seen by users. The latter served as a loose proxy for profile accuracy. Ads were detected via a combination of open-source pre-trained models from the Ad-Versarial [21] project and models created by Hussain et al. [25].

Profile controls: Users can remove any item from their profile. Doing so means the item is removed from the profile, yet could be added again if any future activity leads to the same inference. Users can also mark any item as sensitive, which means that the user will start receiving notifications when the item is inferred again in the future. Some items are marked as sensitive by default, based on a crowdsourced list of people's comfort with Google AdWords categories reported by Weinshel et al. [51]. Users can remove the "sensitive" marker from any item, including those items marked as sensitive by default.

D PARTICIPANT DEMOGRAPHICS

Table 3. General participant demographics and the number of profile views throughout the two weeks of study.

#	Gender Identity	Age	Education	Phone	Computer Browser	Probe Profile Views
P1	Male	25-34	Bachelor's	Android	Chrome	68
P2	Male	45-54	Associate's	Android	Chrome	50
P3	Female	18-24	Bachelor's	iPhone	Firefox	69
P4	Female	45-54	Master's	iPhone	Chrome	103
P5	Male	25-34	Bachelor's	Android	Chrome	174
P6	Male	25-34	Master's	Android	Chrome	113
P7	Male	25-34	Bachelor's	Android	Chrome	109
P8	Male	25-34	Doctoral	iPhone	Chrome	115
P9	Male	35-44	High school	Android	Chrome	20
P10	Male	25-34	Master's	iPhone	Chrome	166
P11	Male	18-24	Some college	Android	Firefox	101
P12	Female	35-44	Master's	iPhone	Firefox	151
P13	Female	25-34	Bachelor's	iPhone	Chrome	76
P14	Male	35-44	Associate's	Android	Firefox	94
P15	Male	18-24	Some college	Android	Chrome	62
P16	Female	25-34	Some college	Android	Chrome	67
P17	Female	25-34	Bachelor's	Android	Firefox	113
P18	Female	35-44	Master's	iPhone	Chrome	94
P19	Female	55-64	Associate's	iPhone	Chrome	53
P20	Female	25-34	Bachelor's	Android	Chrome	103
P21	Female	45-54	Bachelor's	Android	Chrome	101
P22	Female	25-34	Some college	Android	Edge	43
P23	Female	35-44	Bachelor's	iPhone	Chrome	72
P24	Female	18-24	Bachelor's	iPhone	Chrome	50
P25	Female	25-34	High school	iPhone	Chrome	66

E EXIT INTERVIEW PROTOCOL

(1) What are your general impressions of WHO AM I?

(2) Which features do you like or dislike the most?

(3) *Ask one or two individual questions about diary notes and logs to get context from the participant, such as "why did you or did you not remove items?"*

(4) *If participant could not make changes in first week:* Why did you or didn't you change your profile when you could?

(5) Over the two weeks, have you noticed any change in your own behavior? For instance, did you find yourself doing something to prevent or ensure WHO AM I noticed? Please give examples.

(6) *Remind participants about the main point they made regarding their experiences with OBA from the entry interview, then ask:* Reflecting on your responses in the entry interview, how has your experience over these last two weeks affected your understanding and/or attitudes about Online Behavioral Advertising, if at all?

(7) What transparency and/or control features do you desire about inferences made in Online Behavioral Advertising?

(8) *Ask participants to look at one ad explanation from scrolling on their social media feed "why am I seeing this ad?" explanations and an ad interests dashboard, from Google or Facebook (click Interest categories and Other Categories) (even-odd counterbalancing). Ask participants if they were aware of them and used them before. Then ask:* Were you aware of this? How does the information there compare with WHO AM I?

(9) How would you compare WHO AM I with these existing solutions? For example, what are the pros and cons of each? Are there things from WHO AM I you wish were available in the existing solutions?

(10) *You have probably noticed over the two weeks that some inferences about you may represent long-term (e.g., learning a new skill) or short-term (e.g., looking for a restaurant) interests you may have. Others items may stem from a one-off activity such as a search or interest (e.g., when was this movie released?), or recurring and reinforced interests (e.g., investing).* How would you want these differences to be accounted for on your ad-targeting profile, if at all? For example, would you want them to be displayed differently? How and when would you prefer to manually set these differences, if at all?

(11) *You have also probably noticed that there are mistakes on your profile. For example, inaccuracies and/or conflated identities, such as when doing something for work or sharing a device with others.* Would you be interested in correcting those? Why or why not? If so, how might you want to do it?

(12) *Some profile items may be about a concrete interest (e.g., autos and vehicles) while others may hint about who you are or your state at the moment (e.g., socioeconomic status, state of mind).* How would you want your profile to distinguish these, if at all? For example, would you want priority in one over the other?

(13) Did you have something in mind that you wish WHO AM I picked up so you could get ads about it? If so, what was it? How would you want these instances to be accounted for? For example, one indirect way of doing this is searching for it in order to start seeing ads about it. Can you think of a more direct way to do it?

(14) Have you experienced during the two weeks any inference that you thought was inappropriate, sensitive, or creepy? For example, something very accurate and specific, or something negatively surprising. If so, how might you want to be told about and control these items?

(15) I will ask for a moment that you imagine a designer from Google or Facebook being in this call with us. How would you tell them about your preferences to see and control your ad-targeting profile? For example, with what frequency would you like to review pieces of your profile? Every day? At the time they are inferred? Every week or month?

F RESULTS OF MIDTERM AND EXIT QUESTIONNAIRES

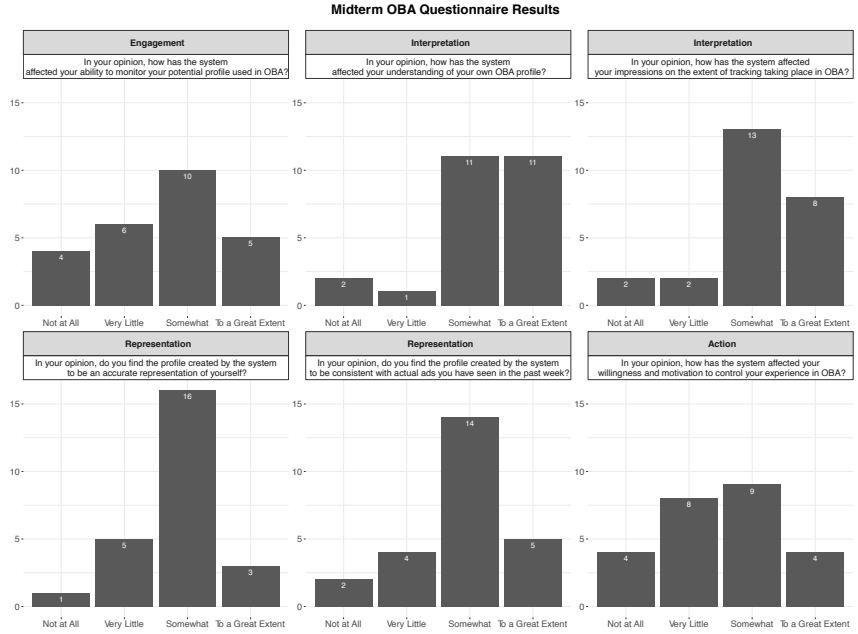

Fig. 4. Results of the Midterm Questionnaire. Responses were collected at the end of the first week.

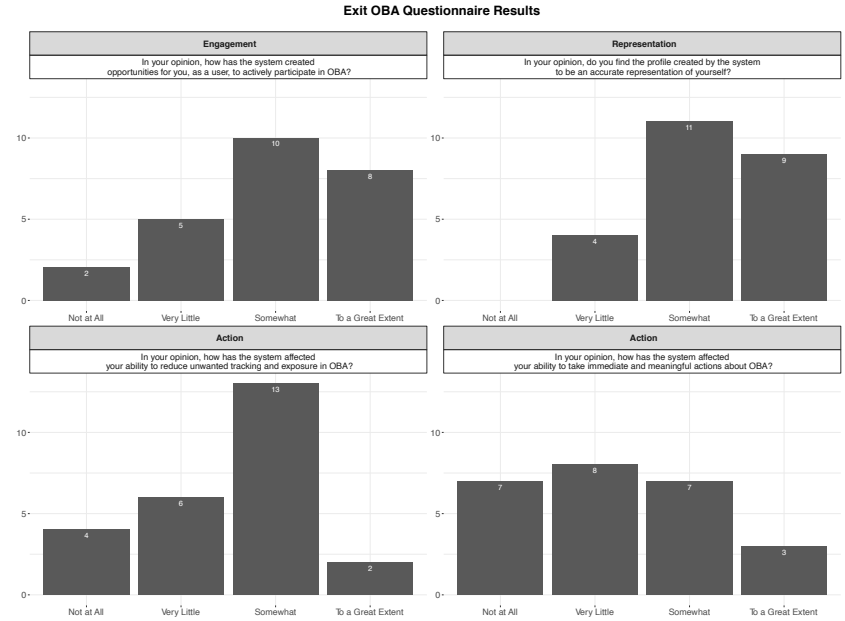

Fig. 5. Results of the Exit Questionnaire. Responses were collected at the end of the second week, before the exit interview.

G QUALITATIVE ANALYSIS CODES

Table 4. Codebook for the qualitative analyses.

Entry Interview

Concept of privacy		**OBA experience**		**Perceived control**		**Data considered**	
Ability to control	11	Ads follow me	14	Resignation	12	Mostly browsing	11
Contextual integrity	9	Mic eavesdropping	9	Use or don't	9	Concrete examples	9
Security	5	Likes ads	8	Specific strategies	9	Blanket "everything"	4
Left alone	4	Mixed about ads	5	Not in control	7		
		Consent unclear	5	Doubts of efficacy	7		
				Control is in behavior	5		
				Not motivated	4		
				Doesn't know how to	3		

Tools used	
AdBlocker	8
Too much work	6
Incognito	6
No tools	4
It breaks things	3
Location off	2
Standard	3

Exit Interview Questions

Liked features		**Behavior explanations**		**Changes in behavior**		**Changes in attitude**	
Behavior traceback	13	"Nothing sensitive"	12	More mindful of connections	6	Motivated to control	10
In-the-moment notifications	8	Removed sensitive	6	Identity performance	2	Unchanged	10
Structure	7	Growth rate of profile	5			More aware	9
Timeline	6	Identity performance	5			Empowered	6
Visibility of data collected	4	OK with mistargets	3			Underestimated tracking	5
Mark sensitive	4	General data	2			Reaffirmed	4
Device filter	4	Monitor for sensitive data	2			Less worried	4
Auto-sensitive	3						
Prominence of items	2						

Desired features		**Comparison**	
In-the-moment controls	13	Behavior traceback	10
Control specificity	6	Mark sensitive	10
More profile visibility	4	Notifications	8
Mark off-limits	3	Timeline	6
Precise explanations	3	Structure	5
Prominence	2	Ad examples	5
Mark sensitive	2	Remove	3

Co-Design Questions

Temporality		**Inaccuracies**		**Interest v. Identity**		**Missing pieces**	
"Figure it out"	14	Fix for more relevant ads	13	No manipulation	3	Identity performance	6
Choice on one-offs	6	Only act if bothered	10			Optimize outcome	5
Prominence	2	Obfuscation is good	5			Questionnaire	4
Direct control	4	Through notifications	4			No incentive	4
		No motivation	2			New interests	2

Sensitive		**Open-ended designer**	
Notification	14	More visibility	13
Warn on ad	2	Invitation to review	8
		Behavior traceback	5
		More like conversation	5
		Perceived control	4

Diary notes

Identity reflection	19	Identity performance	14	Inaccurate and sensitive	10	Utility of inferences	6
Unexpected tracking	17	Boundary regulation	12	Curiosity on mechanics	8	System hiding something?	5
Disillusionment	17	Elucidation	12	Ubiquitous tracking	7		
Omniscience	15	Recalling actions	11	General profile	7		