

“OMG! How did it know that?”

Reactions to Highly-Personalized Ads

Aleksandar Matic
Telefonica Alpha
Placa Ernest Lluch i Marti 5
Barcelona, Spain
aleksandar.matic@telefonica.com

Martin Pielot
Telefonica Research
Placa Ernest Lluch i Marti 5
Barcelona, Spain
martin.pielot@telefonica.com

Nuria Oliver
Vodafone Research
Paddington Central, 1
London, UK
nuria@alum.mit.edu

ABSTRACT

In this paper, we explore the question “would people be willing to share their personal data in exchange for highly-personalized online ads?” through a Wizard-of-Oz deception study. Our volunteers were exposed via a web browser to three different highly-personalized ads, designed by people who knew them well. They were made believe that the ads had been generated automatically by an Artificial Intelligence engine on the basis of their browsing & location history and/or personal traits. The participants’ reactions were surprisingly favorable: in more than 50% of the cases, the ads triggered spontaneous positive emotional reactions; almost 90% of participants would share at least two of the three data sources with advertisers; and about 50% would share all data sources. Our results provide evidence that highly-personalized ads may offset the concerns that people have about sharing their personal data. Thus further efforts in building increasingly personalized online ads would represent a worthwhile endeavor.

CCS CONCEPTS

•Security and privacy→Social aspects of security and privacy;

KEYWORDS

Personalization; advertising; privacy; highly-personalized ads

1 INTRODUCTION

As technology becomes pervasive in our daily lives, computing devices and services are increasingly logging our activities, such as *e.g.* our searches, the content we browse and our whereabouts and movements. One of the primary purposes of this data collection is targeted advertisement: in exchange for free services, we agree to receive personalized ads that have been selected on the basis of our interests, which are inferred automatically from our data traces. Though using human

behavioral data to personalize ads promises to make them less bothersome and more effective [19], ads have been still perceived as annoying [1, 4, 17]. At the same time, only a small fraction of users would pay for an ad-free experience [8, 11].

Personalization has raised concerns: some users find personalized ads to be *scary or creepy* [11, 17], whereas others worry about the personal data that has been used to create the personalized ads [1, 3]. In previous work, when given a choice, many people expressed hesitation to share information: their concerns of sharing personal data—such as browsing and location history—outweigh the perceived usefulness of personalized ads [3, 11, 13]. In this context, we formulate our main research question as:

RQ: *What are people’s reactions and insights towards highly personalized ads given their potential privacy implications?*

Since no real-world system has yet delivered what we will refer to as *highly-personalized ads*, *i.e.* ads that are *extremely-well targeted to their viewers’ current needs, wishes, and intents*, we are unable to determine if in such a scenario the perceived usefulness would outweigh the privacy concerns.

To answer the question above, we conducted a Wizard-of-Oz study where we captured the reactions of 17 participants to 3 highly-personalized ads each. Through elaborate deception, we made the participants believe that the ads had been created by an Artificial Intelligence-based research prototype which had used their browsing logs, location histories, and their inferred personal traits. In reality, the ads had been created individually for the participants by people who knew them well (*e.g.*, a friend or a partner). To the best of our knowledge, no study to date has attempted to investigate the users’ attitudes towards ads that have been personalized beyond what today’s personalization systems can achieve. The purpose of our study is to capture reactions and insights, with a particular focus on qualitative feedback, regarding highly-personalized ads and the use of personal data to deliver such a personalized experience.

In contrast to previous work, where concerns outweigh the perceived usefulness, our study revealed surprising insights:

- Participants found over 90% of the ads to be currently relevant to them;
- participants spontaneously showed positive emotional responses (such as excitement, positive surprise, and enthusiasm) for more than half of the presented ads. This kind of feedback has not typically been reported in previous work; and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UMAP’17 Adjunct, July 09–12, 2017, Bratislava, Slovakia

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5067-9/17/07...\$15.00

<http://dx.doi.org/10.1145/3099023.3101411>

- more than half of the participants would be willing to share all of the three tested personal data sources in order to receive ads of similar quality and relevance, namely: 71%, 65% and 59% for browsing history, location history, and their personal traits respectively.

These findings provide evidence that striving towards computational models to create highly-personalized ads would be a difficult yet worthwhile endeavor.

2 RELATED WORK

Over the past fifteen years, numerous studies have focused on understanding user attitudes towards online advertisements. One of the first studies from 2003 [14] reported that only 13% of the participants (students) enjoyed Internet ads and only 20% found them useful or informative. Negative attitudes towards ads have been consistently reported until today: consumers have found ads to be annoying [17], and it has been demonstrated that online ads can lead to lower return rates to websites [10].

Various strategies have been suggested to make ads more attractive. A good match of ad and topic of the page improved advertised brand recall [5]. However, not all approaches were successful. The use of personal media, such as photos, names, and holiday destinations resulted into higher attention, but caused issues related to acceptability and privacy [9]. Ads which are similar to the content of a website lower the site's perceived credibility, even when said ads were rated highly in isolation [4].

Personalized ads have been shown to decrease annoyance [15] and increase the ads' effectiveness and click-through rates [6]. This commonly applied approach is referred to as online behavioral advertising (OBA). It relies on tracking the users' online activities and behavior to infer their interests and preferences, and to use this information to select relevant advertisements. Even though personalization can make ads more relevant to the consumer, it has also raised mixed emotions and attitudes related to the use of sensitive personal data. Privacy concerns were found to remain stable across time, but that the level of concern with respect to the data used to generate personalized advertisements had increased with time [2].

Several studies have focused on the willingness to share personal data versus the associated privacy concerns. Turow *et al.* [16] reported that 68% of the participants "definitely would not" and 19% "probably would not" allow advertisers to track them online, not even anonymously. Another survey found that 68% of participants were opposed to targeted advertising because they did not want their online behavior to be tracked and analyzed [13]. Users show little desire in targeted ads: 20% of participants expressed interest in personalized advertising but the majority of them (64%) found the idea to be too invasive [11]. A report found that 57% of surveyed individuals were not comfortable with the use of Web browsing history for behavioral advertising (even when anonymized¹). Highly personalized ads might one day use

the consumers' personal traits (such as their personality) as input. As seen before, personality inference made participants feel uncomfortable and they opposed sharing the results in many situations [18]. These studies unanimously indicate that concerns about data sharing clearly outweigh the perceived benefit of personalized ads.

Yet, to date, previous work has only explored the perceptions towards personalized ads that can be created with today's technology. However, in the future, advancements in user modeling and Artificial Intelligence (AI) may allow to create ads that are *highly personalized*, *i.e.* selecting products or services that match the interests of a consumer as well as a product or service that would be chosen by a person who knows that consumer very well. What is missing are insights into people's attitudes – perceived usefulness, associated concerns, and willingness to share the required data– towards such highly-personalized ads. The study described in this paper aims at shedding light on this topic.

3 RESULTS AND DISCUSSION

To answer our research question, we study people's reactions towards the use of three data sources already available today: Web browsing history, location history, and personal traits, *e.g.* personality. We exposed 17 volunteers to a mock-up website that presented highly-personalized ads to them. Participants were made believe that those advertisements had been created by an AI system based on those three data sources. In a semi-structured interview, the experimenters recorded the participants feedback to the individual ads and to the data sources allegedly used to create those ads.

3.1 Creation of Highly-Personalized Ads

To achieve the highest level of personalization and relevance possible, the content of the ads presented during the study was created individually for each participant by a person who knew the participant very well (*e.g.* a friend, a relative or a spouse/partner). The ad content was selected on the basis of information that would be available from **web browsing history, location history, and inferred personal traits**. Note that the creation of ads did not rely on dedicated observation of the participants' data: the ad creators leveraged longitudinal knowledge of the participants' interests, whereabouts and traits. To achieve consistent ads, we gave specific guidelines (described below) to the *ad creators*.

We defined **browsing history** as the set of web pages visited via web browsers from any of the participants' devices. The ad creators were instructed to come up with an ad that could be inferred from browsing history or the information that can be derived from it. An example for using web browsing history is an advertisement for specific shoes that the participant had been looking for in online stores the week before the study took place.

¹ TRUSTe. 2011. Privacy and Online Behavioral Advertising. <https://www.eff.org/files/truste-2011-consumer-behavioral-advertising-survey-results.pdf>. (July 2011)

We defined **location history** as the history of places on a neighborhood level that the person had visited in the past, including the timestamp and duration of the visit. The ad creators were instructed to define an ad that could be inferred directly from the participant's whereabouts. An example for using location history is an advertisement about a holiday package in a region that the participant had recurrently visited in the past.

Inferred **personal traits** were loosely defined as any form of individual characteristics or behaviors that could be relevant to determine ads related to products or services for a particular individual. We intentionally left this aspect open interpretation, so that the ad creators could suggest truly personalized ads based on any information from their long-run history of knowing a participant. An example for using personal traits is an advertisement that captured the fact that the participant had very irregular routines and had been frequently changing hobbies.

The ads themselves were designed by the authors, such that all the ads had the same *look and feel* as shown in Figure 1. As input for the process we took the products and services selected by the ad creators. We then searched for a matching, actually-existing product or service online in order to set the price, select the image, and create a brief description of the item. At the end of this process, there were three highly personalized ads per participant: (1) one based on hypothetical web browsing history; (2) one using browsing- and location history; and (3) a third one using the above two data sources plus personal traits/characteristics of the participant.

The 51 ads (17 participants x 3 ads) covered a wide range of common products and services: 39% advertised products (such as shoes, books, games, furniture, and so on); 22% advertised experiences (including trips, spa centers, gyms, and similar); 26% were about services (for instance, courses, classes, catering, real-state, etc.); and finally 13% were coupons and offers for shopping malls, clothing brands, restaurants, and similar.




Figure 1: Example of one of the highly-personalized ads

3.2 The Highly-Personal Advertising System

The ads were presented through a website that we created for the purpose of this study. Figure 1 displays one of the ads as it was shown to one of the participants. As seen in the Figure, the ad is placed on a page that contains a box with an image and a short

description of the ad. The page has a matching background image for aesthetic purposes. In order to avoid sequence effects, *e.g.* to avoid biases from gradually increasing the number of used data sources, the ads were presented in random order.

Once the three ads had been shown to the participants, they saw a summary web page, which displayed the three ads alongside the data sources that the system had allegedly used to create them.

3.3 Deception Strategy

To obtain ecologically-valid reactions from the participants, we needed to make people believe that the ads had been created by an AI algorithm on the basis of their actual data traces. Thus, we used several strategies to make the study as credible as possible. First, in the informed consent, we emphasized that the prototype that we were testing in this study had been developed in a collaboration between teams of several well-known research- and industrial organizations, and that we were only assigned to help this consortium to evaluate the prototype in a user study. Second, we added a non-disclosure agreement clause to the informed consent, and asked participants not to talk about the study with others because of the novelty and potential value of this technology. Third, after participants introduced their mobile phone number on the study's website, we introduced a fake waiting time of more than 10 seconds that would make the participants believe that the system was performing complex calculations in order to generate their personalized ads. The experimenters would then make an excuse about overload in the servers to justify the delay.

Prior to the study, we tested the deception strategy with four of our colleagues who, given their technical background, could have easily realized that such a system was unlikely to be real. All of them fell for the deception, demonstrating that our strategy was effective. Participants were debriefed individually. We revealed that the system, in fact, did not exist, and that the personalized ads had been manually created for them.

3.4 Data Collection

The data collection was done via semi-structured interview. Each session with each participant was audio-recorded (with explicit consent from each participant) and lasted for approximately 15 minutes. The interviews were structured along the following two questions for each of the presented ads:

“Q1: Right now, how relevant do you find this ad?” and “Q2: What is your reaction to the fact that the system presents you this ad?”.

Similarly, when being shown the data sources that had allegedly been used to create the ads, the experimenters structured the discussion along the following two questions:

“Q3: What is your reaction to learning what information the system has used to generate the ad?”, and “Q4: Would you consider sharing such data to get similarly personalized ads?”

Often, the participants spontaneously elaborated their perceptions and reactions beyond those four questions, in which case we followed up on the elicited topics. As mentioned above,

we distanced ourselves from the prototype by explaining that we had just been “drafted” to test a system, and that we had no further involvement in its development. The purpose was to facilitate the collection of the participants’ honest feedback.

3.5 Participants

In total, 17 volunteers, 8 female, 9 male, took part in the study. Their ages ranged from 26 to 71 years ($M = 39.6$, $SD = 13.6$). The participants came from various European countries, the most frequent being Serbia ($n = 7$), Germany ($n = 5$), and Spain ($n = 2$). They had diverse occupations, including jobs in medical fields ($n = 4$), architecture ($n = 3$), engineering ($n = 3$), and designers ($n = 3$). All participants went online frequently (at least once per day) and were smartphone users. No monetary compensation was given for their participation. For each participant, we also needed to find a person (ad creator) who would know him/her well in order to be able to create highly-personalized ads. Therefore, we used a *snowball sampling* recruitment technique. While snowball sampling is prone to have issues with the representativeness of the sample, our participants had a wide range of ages, nationalities and occupations, thus mitigating potential biases in the data collection, such as common biases from student samples.¹ our sample is gender-balanced.

3.6 Procedure

The study was carried out via Skype or in person, in setting where participants had access to their own computer. At the beginning, the experimenter walked the participant through the informed consent and explained that the study. Once the participant gave informed consent, the experimenter asked him/her to open the study’s website on the computer and to enter his/her mobile phone number. At this point, the rest of the study was recorded using a voice recorder.

After some artificial waiting time, the first ad was displayed. For each ad, the experimenter would then ask the participant about his/her reactions to seeing this ad, following questions Q1 and Q2. The experimenter clarified that *relevant* meant that the content of the ad was aligned to his/her interests, *i.e.* it would be for an item that (s)he would use/consume. Once the participant had seen all three ads, the system showed them the web page with the explanation of the data sources that had been used to create the seen ads. For each of the data sources: browsing history, location history, and inferred personal traits, the experimenter first clarified the meaning if necessary, and then asked Q3 and Q4.

Finally, the experimenter debriefed the participant by explaining that the system and the ads had been manually created, and by disclosing the true purpose of the study. with sharing sensitive personal data traces for this purpose.

4 RESULTS

All of the 17 participants completed the study without showing signs of suspicion that the system was a deception. From the interview transcripts, we identified reactions related to the *current relevance* of the ads and the *emotions* they evoked, both positive,

such as surprise or excitement, or negative, such as anger or worry. Interestingly some participants expressed negative and positive emotions towards the same ad.

In brief, 90.2% of the 51 ads were considered to be relevant, 54.9% evoked positive emotional reactions, and 39.2% evoked negative emotional reactions. Participants expressed both positive and negative emotions towards 19.6% of the presented ads.

Contrasting previous work, the participants showed a surprisingly high willingness to share their personal data in exchange for receiving similar, personalized ads:

- 9 (52.9%) of the 17 participants would share all three kinds of data sources with advertisers in exchange for ads of similar quality and relevance;
- 15 (88.2%) of 17 participants would share two (but not all three) of the data sources in exchange for ads of similar quality and relevance; and
- only 2 (11.8%) participants would *not* share any personal data with advertisers.

4.1 Reactions towards the Personalized Ads

In the following, we report the reactions of the participants to the highly-personalized ads. Table 1 breaks down the reactions by data sources.

Table 1: Reactions towards the three classes of ads

Data Sources	Considered Relevant	Positive Reaction	Negative Reaction
Browsing History	88,2%	58,8%	26,5%
Browsing & Location History	88,2%	35,3%	35,3%
Browsing, Location, Personal Traits	94,1%	70,6%	52,9%

1. Web Browsing (WB) History: Regarding the 17 ads that were hypothetically generated from web browsing history, 15 (88.2%) ads were found to be **currently relevant**. The participants expressed different reasons for finding the ads relevant: “[The presented ad about a game] is very relevant, I googled it recently!” (P0037); “This is the author that I happen to be reading right now.” (P4134)

Positive reactions were recorded in 10 (58.8%) of the cases, *e.g.*: “This is great, I need this, this is perfect for me, how does [the website] know that?” (P7501); “The system really knows me well. I would not have selected these things so well myself! I am amazed!” (P2485)

Negative reactions were recorded in 5 (29.4%) of the cases, *e.g.*: “It annoys me! Because this means that there is a system that follows me constantly...” (P6561); “[Swearing]. How do they know it? I am angry, very angry!” (P6644)

2. Web Browsing (WB) & Location (L) History: Regarding the 17 ads that were hypothetically generated from location history combined with web browsing history, 15 (88.2%) ads were found to be **currently relevant**. Examples include: “This is very relevant, because we just visited the Brother Grimm Museum in [city]!” (P6795); “These are my last days [at the current location] and this ad is really relevant to the move.” (P7001) Participants **reacted positively** in 6 (35.3%) of the cases, *e.g.*: “This is really cool because I like kite surfing!” (P7335); “I would immediately take that! I was looking for holiday

apartments when we last came visiting [frequently visited city]" (P6795)

Participants **reacted negatively** in 6 (35.3%) of the cases, e.g.: "That's pretty personal.. so it's creepy that it knows [my location]!" (P5134) When P6644 realized that the ad could not have been inferred without knowing his location history, he riled that the ad was *creepy! I am angry!*

3. WB & L History, and Personal traits: Regarding the 17 ads that were hypothetically generated from personal traits, location and web browsing history, 16 (94.1%) ads were found to be **relevant**. P5134 found the ad to be "ridiculously personalized [laughing] this was the most useful in terms of showing me something that I hadn't thought about."

Participants **reacted positively** in 12 (70.6%) of the cases, e.g.: "They have captured my hobbies, etc... all the important things for me and it's already captured." (P1566); "I would go to this with pleasure! .. Now I am even more positive about the system." (P7110)

Participants **reacted negatively** in 9 (52.9%) of the cases, e.g.: "I feel like I am in a psychological experiment. I wouldn't like to receive this ad at all!" (P9220); "I am shocked what they can know and what they can calculate about me." (P6644)

4.2 Reactions Towards the Used Data Sources

Table 2 breaks down the participants' willingness to share their data by data source. The participants' reactions are described next.

Table 2: Fraction of participants who required clarification about the data sources and who would be willing to share such data for ads of similar quality

Data Sources	Asked for Explanation	Would Share Data
Browsing History	23,5%	70,6%
Location History	5,9%	64,7%
Personal Traits	41,2%	58,8%

1. Web Browsing History After seeing the explanations, 12 (70.6%) of the 17 participants expressed their willingness to share their browsing history with advertisers in order to receive ads of similar relevance. Clarifications about the data that the browsing history included were required in 4 (23.5%) cases. Examples of positive and negative reactions are: "It's okay [to share browsing history], it is no secret data" (P6795); "I do not really care. I'm not a terrorist. I don't care that people know what I am looking at because it is harmless" (P7501); or "It is a bit uncomfortable that I don't know what exactly they look at [in my browsing history]" (P6101)

2. Location History: 11 (64.7%) of the 17 participants reported their willingness to share their location history with advertisers in order to receive ads of similar quality. Clarifications about the data that the location history included were required in 1 (5.9%) case. Examples for positive and negative reactions are: "Sharing browsing history is ok, location history is not a problem either. I am fine with that" (P0056); "I am totally opposed to this. This is the first step into a surveillance state!" (P0377)

3. Personal Traits: 10 (58.8%) of the 17 participants would allow advertisers to infer their personal traits in order to receive similar ads. Clarifications about what the personal traits meant were requested in 7 (41.2%) cases. Examples for positive and negative reactions are: "If they can model my personality, go for it! why not?" (P7501); "This can save your time [...] I just do not want to receive thousands of emails with different offers" (P5067); "Personality! Oh my God! ... how they know modeling personality is a little creepy... How can the computer know how I am?" (P9220).

5 DISCUSSION AND IMPLICATIONS

5.1 High Relevance brings Positive Emotions

According to the responses of the participants, over 90% of the presented ads were relevant. This fraction is considerably higher when compared to the state-of-the-art relevance of computational ads which ranges between 40-50% [3, 12, 4]. The amount of positive reactions that we recorded was unexpectedly high: more than 50% of all the ads evoked positive reactions. In contrast, while in previous work, participants recognized the usefulness of personalized ads, positive emotions are rarely reported [3, 12, 4, 17].

The fraction number of positive reactions occurred when combining all three sources of personal data. In this case, given the availability of different sources of information, the created ads were very suited to the person. Yet, as discussed above, the ads created based solely on web browsing history were almost equally successful in evoking positive reactions, while being possibly more feasible to be deployed in practice by an algorithm.

The fraction of negative emotions was comparably low. In previous work annoyance was a common reaction to the irrelevance of the ads. e.g. as reported in [17], the second most common reaction was that the online ads are annoying, whereas in [1] more than 70% voiced that the targeted ads are repetitive and annoying.

5.2 High Willingness to Share Personal Data

A surprisingly large fraction of the participants was willing to share their personal data with advertisers if they were to receive ads of similar relevance as the ones shown in our study: 70.6% of the participants would share their browsing history, 64.7% their location history, and 58.8% their personal traits. These numbers are notably higher than those reported in previous work, where 50 to 80% of participants rejected the use of their personal data for advertisement [11, 3, 13, 18]. We hypothesize that the key difference lies in that fact that this study presented the participants actual highly-personalized ads, so that they could evaluate the potential benefits in a more direct way.

Interestingly, negative emotions did not automatically translate into not wanting to share data. In 6 out of 20 cases where people expressed negative emotions, participants were still willing to share their personal data. In 4 of the 10 cases where people had expressed both negative and positive emotions towards the ads, they would still share their personal data with advertisers. Despite

having concerns, a non-negligible percentage of the participants perceived enough value in highly-personalized ads to share their personal data.

5.3 Factors for User Acceptance

About one third of the participants commented that their willingness to share data depends on the availability of data control mechanisms and the possibility to opt-in or opt-out. This confirms findings by [12, 4] and [8] where *being in control* also emerged as an important factor. It has also been suggested that granular control over data sharing may result in increased users' willingness to share their data [7]. However, providing control over personal data for this purpose is by no means a trivial task. It has been shown that most users do not understand the purpose of privacy tools [8].

More than half of the participants in our experiments expressed that they believed privacy to be a *lost battle*, such that their personal information would be used in any case, with or without their consent. These findings corroborate previous work and show that there is a need and an opportunity to regain people's trust and goodwill by designing systems that (1) enable more control over their personal data and (2) provide increased transparency on how the data is being used. In this respect, projects like the *Data Transparency Lab* aim to stimulate research on personal data transparency.

6 LIMITATIONS

When interpreting the findings from this study, a number of limitations have to be considered. First, since the ads had to be highly personalized, they differed in terms of content and price across participants. As previously explained, we made an effort to ensure a consistency across ads, which was achieved to the extent that was possible. Second, we resorted to snowball sampling to be able to recruit both participants and ad creators. We made an explicit effort to mitigate potential sampling biases by including participants of a wide range of ages, countries and professions. Yet, snowball sampling is known to yield biased samples. In our case, all our participants had a college degree. Third, our sample size is small ($N=17$), though within the range of the most related work that relied on semi-structured interviews. Thus, the percentages that we provide should be interpreted as tendencies. Forth, our participants had presumably never seen such highly-personalized ads before. The strong emotional responses, positive and negative, may be explained by the novelty of the experience. Finally, the ads we exposed our participants to were harmless. None touched potentially sensitive content, such as medicines for a particular illness. Such factors have to be considered when striving towards highly-personalized ads.

6 CONCLUSIONS

In this paper, we describe the results of a deception study in which 17 participants were exposed to highly-personalized ads. Through semi-structured interviews we captured the participants' attitudes towards the ads, and investigated their willingness to

share their personal data –namely, browsing- and location-history, and inferred personal traits– in order to receive such personalized ads.

In contrast to previous work, more than half of the displayed ads elicited positive reactions, such as surprise or enthusiasm. Almost 60% of the participants would share all the three kinds of personal data and almost 90% of the participants would trade two of the three data sources in exchange for ads of similar relevance.

This work provides evidence that increased personalization and perceived relevance of online ads, combined with improved control and transparency on how the data is being used, may boost acceptance of online ads and enhance customer experience. Highly-personalized ads would bring value to users and be worth pursuing for publishers and online service providers.

We believe that these results will motivate future work in the domains of user modeling, HCI, and computational advertisement that aim at enhancing the personalization of ads while providing more transparency and control to their end users.

REFERENCES

- [1] Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. 2013. Do Not Embarrass: Re-examining User Concerns for Online Tracking and Advertising. In *Proc SOUPS '13*. ACM
- [2] Annie I. Anton and Julia B. Earp Jessica D. Youn. 2009. *How Internet Users' Privacy Concerns Have Evolved Since 2002*. Technical Report.
- [3] Chanchary and Sonia Chiasson. 2015. User Perceptions of Sharing, Advertising, and Tracking. In *Proc SOUPS '15*.
- [4] Henriette Cramer. 2015. Effects of Ad Quality: Content-Relevance on Perceived Content Quality. In *Proc CHI '15*. ACM, 2231–2234.
- [5] Marco de Sa, Vidhya Navalpakkam, and Elizabeth F. Churchill. 2013. Mobile Advertising: Evaluating the Effects of Animation, User and Content Relevance. In *Proc CHI '13*. ACM, 2487–2496.
- [6] Jun Yan, Ning Liu, Gang Wang, et al. 2009. How Much Can Behavioral Targeting Help Online Advertising?. In *Proc WWW '09*. ACM, 261–270.
- [7] J. Gomez, T. Pinnick, and Soltani. A. 2009. KnowPrivacy. http://knowprivacy.org/full_report.html. (2009)
- [8] Pedro Giovanni Leon, Blase Ur, Yang Wang, et al. 2013. What Matters to Users? Factors that Affect Users: Willingness to Share Information with Online Advertisers. In *Proc SOUPS '13*.
- [9] Miguel Malheiros, Charlene Jennett, Snehal Patel, et al. 2012. Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-media Personalized Advertising. In *Proc CHI '12*. ACM, 579–588.
- [10] Scott McCoy, Andrea Everard, Peter Polak, and Dennis F. Galletta. 2007. The Effects of Online Advertising. *Commun. ACM* 50, 3 (March 2007), 84–88
- [11] Aleecia McDonald and Lorrie Faith Cranor. 2010. Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising. In *Proc TPRC '10*.
- [12] Katie O'Donnell and Henriette Cramer. 2015. People's Perceptions of Personalized Ads. In *Proc WWW '15*. International World Wide Web Conferences Steering Committee, 1293–1298
- [13] K. Purcell, J. Brenner, and L. Rainie. 2012. Search engine use.
- [14] R. Rettie, H. Robinson, and B. Jenner. 2003. Does Internet advertising alienate users? *Occasional Paper Series* 52 (2003)
- [15] Christian Rohrer and John Boyd. 2004. The Rise of Intrusive Online Advertising and the Response of User Experience Research at Yahoo!. In *Proc CHI EA '04*. ACM, 1085–1086
- [16] Joseph Turow, Jennifer King, Chris Jay Hoofnagle, et al. 2009. Americans Reject Tailored Advertising and Three Activities that Enable It.
- [17] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, et al. 2012. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. In *Proc SOUPS '12*. ACM, Article 4, 15 pages.
- [18] Jeffrey Warshaw, Tara Matthews, Steve Whittaker, et al. 2015. Can an Algorithm Know the "Real You"? Understanding People's Reactions to Hyper-personal Analytics Systems. In *Proc CHI '15*. ACM, 797–806.