# Position Paper for W3C DNT Workshop

**Rebecca Balebako, Pedro G. Leon, and Blase Ur**
Carnegie Mellon University CUPS Lab

The current multi-stakeholder process to define the Do Not Track standard has considered inputs from many parties, including browser vendors, network advertisers, privacy researchers, and the government. However, we feel that average users, the millions of people who will be directly affected by a Do Not Track standard, have been underrepresented. These are the users whose information is collected, and also those users who benefit from having more customized and relevant advertising.

Over the past year, our research group has published a series of studies analyzing online behavioral advertising from the perspective of end-users. These studies have added the voice of these users to the debate about DNT and online behavioral advertising. At the W3C Workshop on Do Not Track, we hope to continue to give voice to users as part of the debate on the future of online privacy.

The W3C Tracking Protection Working Group is currently working on the definition of the DNT standard, aimed to improve user privacy and user control.[1] Of course, DNT is not the first attempt at letting users express online privacy choices. In a study published at CHI 2012, our group evaluated the usability of nine popular privacy tools provided by the advertising industry, third-party developers, and browser vendors. Our study found substantial usability flaws in all of the tools we tested, ranging from inappropriate default settings to inappropriate feedback to confusing configuration options [Leon et al. 2012a].

While analyzing the usability of existing privacy tools is essential in understanding the current situation for end-users, a detailed look at what users understand and think about online behavioral advertising is also essential to moving consumer privacy forward. We have provided such an analysis through in-depth interviews with 48 average users, as described in a paper we titled “Smart, Useful, Scary, Creepy” [Ur et al. 2012]. The duality of this title reflects the conflicting nature of end-users’ opinions of OBA. While the idea of targeting advertising based on past behaviors on the web seemed smart and useful to many study participants, the tracking that provided this data without consumers’ knowledge struck participants as scary and creepy.

In particular, the results of these interviews provide directions for improving current transparency and control mechanisms to support non-experts’ online privacy preferences. This task is challenging for a number of reasons. The online behavioral advertising ecosystem is complex, and users often have difficulty

[1] http://www.w3.org/2011/tracking-protection/

understanding the different players in this ecosystem.  In particular, how their information flows is confusing, even for privacy experts. Users would need to reason about the privacy practices of the many companies in the advertising ecosystem to make decisions about online tracking.

Among the lessons from these interviews is that participants had difficulty making decisions about different companies involved in online behavioral advertising. On one hand, this result lends support to current Do Not Track implementations in which users make blanket decisions about tracking, in contrast to many existing privacy tools that requires users to make filtering decisions on a per-company basis.

However , the idea that Do Not Track enables users either to send or not to send a blanket Do Not Track signal conflicted with the context-sensitive nature of our participants' opinions about online behavioral advertising. In particular, we learned that context of browsing matters. Both utility and privacy fuel users' attitudes towards allowing or not allowing advertising to be targeted based on their browsing. As a result, an all or nothing DNT approach does not seem to support user preferences fully. Furthermore, the solution is not as simple as specifying globally in which browsing circumstances DNT would apply. We observed that some browsing scenarios that were innocuous to certain users were perceived as privacy violations by other users, and vice versa. These results suggest that users would benefit from a Do Not Track variant enabling them to communicate their preferences dynamically.

Through these interviews, we also found that users strongly fear the collection of personal information for the purpose of tailored advertising. Of course, the form that users' collected information takes in databases is often a mystery to users. While most network advertisers claim that no personal information is collected, research has found instances in which popular websites leak personal information to third parties [Krishnamurthy et al. 2011]. Furthermore, the distinction between first-party and third-party advertisers blurs for companies like Google, Yahoo, Adobe, and Microsoft, who in certain cases are a first party, while in other cases are a third party. This line is further blurred by "social plugins," such as Facebook "Like" buttons, Google "+1" buttons, and Twitter widgets. We believe that there is tremendous opportunity for added transparency about the detailed types of data collected about users.

While we advocate increased transparency concerning the data that has been collected about individual users, we have also investigated transparency about the practice of online behavioral advertising in general. In particular, we investigated the icons and taglines (e.g., "AdChoices") that are displayed across the Internet to determine what message average users glean from these privacy disclosures [Leon et al. 2012b].  We found that the icons and taglines overall did not communicate clearly about behavioral advertising.  Our study participants were more likely to think these icons would let them purchase their own advertisements than to understand that they could make choices about their privacy related to targeted advertising. In contrast, we found that including verbs and other action words could communicate more successfully about online behavioral advertising and eliminate misconceptions.

In order to develop effective privacy options for users, it is also necessary to understand the current privacy situation. Over the past year, we have developed and tested a novel method for measuring the effectiveness of privacy tools that claim to limit online behavioral advertising [Balebako et al. 2012]. This method enables systematic measurements of OBA that is based on past browsing. We automated the collection of advertisements in a way that controlled for the time of visit, IP address, machine setup, and Flash LSOs. In a case study of this method, we found that third-party browser plugins and cookie-based tools we tested were effective in reducing behavioral advertising. However, in our case study, Do Not Track headers were ineffective in reducing behavioral advertising. Our method is particularly relevant in analyzing the efficacy of DNT as its standardization progresses.

Our current work has extended our measurement study across a wider range of advertising agencies to capture more completely the effectiveness of existing privacy tools. We will continue to measure the effectiveness of OBA tools through a longitudinal study, examining the effectiveness of DNT, cookie blocking, and browser plugins over the next two years. Our experience measuring tools has provided insight into the actual result of browsing information being used to target ads, supplementing users' perspectives with a snapshot of current targeting practices.

Moving forward, we believe that the Do Not Track standard should consider providing users with meaningful information that they can use to make informed decisions about tracking. For instance, we believe that sharing, retention and secondary uses of information, as defined by the Privacy Rulesets project,[2] are important aspects to consider. In addition, we believe that further investigating non-expert users' mental models could shed light about other important elements of the Do Not Track standard. Does it make sense to include basic or advanced options in DNT to accommodate the needs of users for whom a simple on/off switch doesn't provide support? What kind of feedback is needed for users to understand the implications of their filtering decisions? What level of access to their collected data would be appropriate for users? When should users be required to make a decision?

We would like to supplement the Do Not Track conversation with our belief that providing users control over data that has been collected about them can better align privacy options with users' wishes. Currently, a number of network advertising companies provide a dashboard that allows user to both learn what information is being collected and request the removal of certain data. We believe that the DNT standard or subsequent privacy tools should provide the same level of

[2] http://dev.w3.org/2009/dap/privacy-rulesets/

access and control without requiring users to identify themselves. This approach could help protect users' privacy and allow ad networks access to richer information. In the same way cookies are used to identify unique users, they can be used to provide access and control capabilities.

Taken as a whole, our work over the past year is particularly valuable to the current debate about online privacy by providing a voice to average users. Many users have little or no knowledge about the mechanisms that enable online tracking, and all humans are subject to a number of cognitive limitations and biases. For instance, to use behavioral economics terms, humans have bounded rationality, employ hyperbolic discounting, and display both overconfidence and a tendency to stick with default options.

Our research has demonstrated that current privacy tools and mechanisms for providing users privacy notice and choice present many opportunities for improvement. Our deep investigations of users' attitudes and abilities to use existing privacy tool have provided insight into designing privacy mechanisms that better align with users' expectations and mental models. While the needs of network advertisers, privacy advocates, and other stakeholders are fundamental to the design of the Do Not Track standard, we believe that the voice of users is also essential. We hope to provide that voice at the workshop.

## References

[Balebako et al. 2012] Rebecca Balebako, Pedro Leon, Richard Shay, Blase Ur, Lorrie Faith Cranor, "Measuring the Effectiveness of Privacy Tools for Limiting Behavioral Advertising," In Web 2.0 Security and Privacy Workshop (W2SP), San Francisco, California, May 2012.

[Krishnamurthy et al. 2011] Balachander Krishnamurthy, Konstantin Naryshkin, and Craig E. Wills. Privacy leakage vs. protection measures: The growing disconnect. In Web 2.0 Security and Privacy Workshop (W2SP), Oakland, California, May 2011.

[Leon et al. 2012a] Pedro G. Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, Yang Wang, "Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising," In Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI), Austin, Texas, May 2012.

[Leon et al. 2012b] Pedro G. Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, Guzi Xu, "What Do OBA Privacy Disclosures Communicate to Users," In Workshop on Privacy in the Electronic Society (WPES) 2012.

[Ur et al. 2012] Blase Ur, Pedro G. Leon, Lorrie Faith Cranor, Richard Shay, Yang Wang, "Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising,"

In Proc. of the Symposium On Usable Privacy and Security (SOUPS). Washington, D.C., July 2012.