DSSV 2023 Data Science, Statistics & Visualisation
ECDA 2023 European Conference on Data Analysis

The trade-offs of obscuring your digital footprints

S. Goethals[1], S. Matz[2], F. Provost[3], Y. Ramon[1], and D. Martens[1]

[1]University of Antwerp, Department of Engineering Management, Antwerp 2020, Belgium
[2]Columbia Business School, Management Department, New York, United States
[3]NYU Stern School of Business, TOPS Department, New York, United States

Keywords— Targeted advertising, Privacy, Machine Learning

A growing part of human life is happening online, including shopping, entertainment, social interactions, and are mediated by digital platforms [1]. This digitalization led to an explosion of online traces we leave behind, think about your Spotify playlist, Google search history of Instagram feed, and creates an extensive picture of our personal habits and preferences. Content that used to be intensely private, and worthy of legal protection, is now freely available to one's network of friends

People may wish to keep certain aspects of their lives private, such as their sexual or political orientation, and yet this information may be revealed from the digital traces that they leave behind. It has been shown that properties such ethnicity, sexual or political orientation, personality traits, mental health and religious views can be accurately predicted from a set of someone's Facebook likes [1]. Revealing an individual's private traits without their consent or even their knowledge could have very consequential implications: for example, in countries where homosexuality is illegal, governments could obtain the identity of people that are more likely to be homosexual; neo-nazi organizations could identify people in certain regions that are likely to be Jewish; or health assurance companies could attempt to identify people with unhealthy habits (interested in smoking, drugs, fast food, etc.) or specific health problems, and increase their health insurance cost (or even not accept them at all).

But how can we protect the privacy of individuals? Cloaking your digital traces has been suggested as a potential solution in the past. Chen et al (2017) propose a cloaking device that points users to the online traces without which the model would not have made the inference and can thus guide users to better decide which data they feel comfortable sharing. This cloaking device has been shown [2] to be effective; however, it is not clear how effective this mechanism would be over time. There can be a wealth of redundant information in online behavioral data, as it is sparse and extremely high-dimensional, and if the cloaking does not also cover closely related features, one might end up being targeted again in the future [2]. On the other hand, we do not know what the impact will be on other prediction tasks, for which the user still wants to receive personalized content. In this study, we explore the effect of cloaking the metafeatures instead of the fine-grained features and answer the following questions:

- Can cloaking metafeatures instead of fine-grained features avoid future inferences to a further extent?
- What is the impact on desirable inferences when users cloak a larger part of their data?

In this study, we use data from the MyPersonality project, which contains the liked Facebook pages of 58,000 volunteers in the United States, along with their scores on the Big 5 personality traits and some personal characteristics such as gender, age, and sexual / political orientation [1]. We investigate the effect of cloaking gender, political orientation, and sexual orientation. *Cloaking* your data means changing the data of a user so that it was as if the user did not exhibit this behavior (for sparse, behavioral data, this means setting the feature value to zero).

This research highlights the potential danger of using digital traces to make private inferences about individuals. Many people are unaware of the extent to which online activity can reveal personal information about them, and do not understand how concealing part of their data can have an impact on the performance of other prediction tasks. It should be up to the user to decide whether he is willing to give up some personalization to gain privacy.

[1] M. Kosinski, D. Stillwell, and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," *Proceedings of the national academy of sciences*, vol. 110, no. 15, pp. 5802–5805, 2013.
[2] D. Chen, S. P. Fraiberger, R. Moakler, and F. Provost, "Enhancing transparency and control when drawing data-driven inferences about individuals," *Big data*, vol. 5, no. 3, pp. 197–212, 2017.