

The Personalization Paradox: the Conflict between Accurate User Models and Personalized Adaptive Systems

Santiago Ontañón*

Drexel University
Philadelphia, PA, USA
santi.ontanon@gmail.com

Jichen Zhu†

Drexel University
Philadelphia, PA, USA
jichen.zhu@gmail.com

ABSTRACT

Personalized adaptation technology has been adopted in a wide range of digital applications such as health, training and education, e-commerce and entertainment. Personalization systems typically build a *user model*, aiming to characterize the user at hand, and then use this model to *personalize* the interaction. Personalization and user modeling, however, are often intrinsically at odds with each other (a fact some times referred to as the *personalization paradox*). In this paper, we take a closer look at this personalization paradox, and identify two ways in which it might manifest: *feedback loops* and *moving targets*. To illustrate these issues, we report results in the domain of personalized exergames (videogames for physical exercise), and describe our early steps to address some of the issues arisen by the personalization paradox.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing systems and tools**;

KEYWORDS

personalization, user modeling, adaptive systems

ACM Reference Format:

Santiago Ontañón and Jichen Zhu. 2021. The Personalization Paradox: the Conflict between Accurate User Models and Personalized Adaptive Systems. In *26th International Conference on Intelligent User Interfaces (IUI '21 Companion)*, April 14–17, 2021, College Station, TX, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3397482.3450734>

1 INTRODUCTION

Personalized adaptation is the idea of designing systems that automatically adapt the user experience to the current user. This idea has been adopted in a wide range of digital applications such as health, training and education, e-commerce, and entertainment [30]. By using artificial intelligence (AI) to tailor themselves to individual users' needs and preferences, adaptive digital applications have

shown to improve learnability [8], usability [15] or user enjoyment [28]. Moreover, personalization technology has also risen a number of open challenges (such as interpretability or controllability [16]). This paper focuses on an inherent problem of personalized systems which we refer to as the *personalization paradox*¹. The term was first used by Jarno Koponen, who, in a 2015 *TechCrunch* article², described it as:

“there lies a more general paradox at the very heart of personalization. Personalization promises to modify your digital experience based on your personal interests and preferences. Simultaneously, personalization is used to shape you, to influence you and guide your everyday choices and actions. Inaccessible and incomprehensible algorithms make autonomous decisions on your behalf. They reduce the amount of visible choices, thus restricting your personal agency.”

While this is an insightful observation, the discussion around the personalization paradox has so far remained abstract. Reflecting on existing work on personalization, we argue that the personalization paradox is a result of the fundamental **conflict between user modeling and personalized adaptation**. We use this angle to analyze the paradox deeper and identify two main ways in which the paradox can manifest when integrating user modeling and personalization, namely: *feedback loops* and *moving targets*.

In the remainder of this paper, we first briefly introduce some background on user modeling and personalization, then present our analysis of the personalization paradox, and conclude by illustrating the issues we identified within the context of a personalized system in the domain of exergames [5], aiming at behavioral change to promote physical activity. For more details on our work on personalized exergames, and how the personalization challenge arose in our experiments, the reader is referred to [29].

1.1 Background

A significant amount of research exists in personalized adaptation, especially in Internet applications. This ranges from just offering customization options that users themselves configure to complex automatic customization using machine-learned models of user preferences, needs and behavior [20]. These systems are typically composed of two subsystems [30]: *user modeling* and *adaptation*.

*Currently at Google.

†Currently at the IT University of Copenhagen, Denmark.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

IUI '21 Companion, April 14–17, 2021, College Station, TX, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8018-8/21/04.

<https://doi.org/10.1145/3397482.3450734>

¹The term “personalization paradox” is some times also used to refer to the “privacy-personalization paradox” [1] (personalization creates users' sense of vulnerability and lower adoption rates), which is different from the use in this paper.

²Retrieved from <https://techcrunch.com/2015/06/25/the-future-of-algorithmic-personalization/>, Oct 1, 2020.




Figure 1: Personalization might create (1) Self-reinforcing loops that push user behavior to become what the system predicts, and (2) Make the user a moving target for user modeling by changing the digital environment context.

User modeling approaches range from the collection of user profiles to the use of machine learning algorithms, and has been applied to domains such as e-commerce [9], training/education [19], social media (e.g. Facebook, Twitter), health [6], and video games [28]. Depending on the application domain, user models have been used to capture a range of user characteristics. For example, in the domain of video games, work exists on capturing player behavior [7, 13, 27], player types [4, 14], user preferences [25, 26], skill level [17, 21], knowledge [18] and goals [12].

User models are then used to adapt the user experience given a target *personalization goal*. This personalization goal depends on the target application. For example, in a personalized exergame, the goal is to personalize the experience to maximally motivate the user to perform physical activity [29]. In training/education, the goal is to maximize learning outcomes [18, 19]. In summary, a standard personalization system works by building a user model, and then adapting the experience based on that model, attempting to affect user behavior in order to achieve the personalization goals.

2 THE PERSONALIZATION PARADOX

We argue that the personalization paradox is a result of the fundamental conflict between user modeling and personalized adaptation. The key underlying problem is that while user modeling tries to acquire a model of some aspects of interest of the user (such as their preferences), personalized adaptation changes the context the user interacts with. There is a large body of literature in psychology that consider human behavior and preferences to be contextual, hence pointing out that that personalization influences the user herself. This “influencing the subject that we are trying to model” issue can manifest in several ways. Through an analysis of existing work, we identified two main forms this paradox can take (not mutually exclusive), listed below and illustrated in Figure 1:

Self-Reinforcing Loops: This problem happens when a personalized system “forces” a user into what its user model categorizes, regardless whether the model is accurate. This self-reinforcing nature of personalized technology has been documented by other researchers [22–24]. For instance, imagine the *Netflix*’s user model inaccurately predicts a specific user’s preference to be only Sci-fi based on her viewing history and thus only recommends Sci-fi content. Since the user can only express her preference through the digital environment controlled by the personalization algorithm, she is more likely to further display Sci-fi preference due to the lack of other choices. In this way, the adaptation reinforces its user

model without a chance to adjust the latter. Another recent example was reported in a recent study by Arnold et al. [3], who showed that adaptive predictive text interfaces affect the text being written by users. As the authors put it “predictive text encourages predictable writing”.

Moving Target: Even when no self-reinforcing loops occur, the fact that personalization affects the user’s behavior is still problematic. Continuing the above example, let us assume that the user model correctly categorized a user’s preference to Sci-fi when she is in the “neutral” viewing environment with a wide range of different genres. However, when adaptation personalizes this context into a Sci-fi-heavy one, the user’s preference may change to, say, documentary. This case is the closest to Kopenen’s description above – by changing the users’ digital environment through adaptation, users’ preferences and behavior become a moving target for modeling. A special case of this problem occurs when the goal of the personalization system is to induce behavior change. In this situation, the system’s explicit goal is to push the user’s preferences or behavior in a particular direction. As a result, user modeling might reflect the user at the start, rather than what she has become.

2.1 Adaptive Exergames Case Study

In our recent work on personalization in the context of physical activity (PA) [10], we designed an adaptive system based on *social comparison* [2] to motivate users to perform PA. Specifically, we designed a web-based platform in which users can compare themselves, including their daily steps with other users’ PA-related profiles. The user’s steps are captured by *Fitbit* and synced automatically with our platform. We use the AI technique of multi-armed bandits to model individual users’ social comparison preferences and adapt the comparison targets shown to them [10, 11]. For example, if the user model predicts that a particular user tends to prefer *upward comparisons*, the system will show more profiles with a larger number of daily steps.

Concerning feedback loops, in our work, we hypothesize that when the users become more physically active, they will perform less downward comparison (when they compare themselves with other users that performed less physical activity). What is captured in the user model is often connected to the desired actions by design. In our case, our underlying assumption is that whose activity level a user compare herself to (upwards or downwards) will impact the user’ PA and motivation. Thus, if were to add other mechanics to our system that further foster these changes (e.g., competition), this will likely influence players towards upward comparison behaviors, creating a feedback loop, and thus affecting user modeling. In order to avoid this self-reinforcing loop, we avoided the inclusion of such mechanics, and also provided different profiles the users could choose to compare themselves to. This was designed this way to allow them to express a comparison preference different from what the user model predicted at the time.

Our attempt to minimize the *moving target problem*, however, was met with mixed success. In our domain, the personalization component affects the type of other user profiles the user can choose for comparison. Similarly as the *Netflix* example above, as the personalization component changes the distribution of profiles, it is unclear if the user behavior will remain constant. For instance,

when a user increases her PA, will she compare upwards more often? Currently, psychology literature does not have a definitive answer. If so, the moving target problem of the changing user in our case may be exacerbated by the targeted behavior change, in addition to the adapted digital environment.

To mitigate this Moving Target problem, we used two strategies. First, our user modeling does not try to model user behavior based on their comparison choice, but tries to model the expected amount of physical activity a user will be expected to do in each of the contexts our personalization system can choose from. Second, we designed the app to be as neutral as possible — we intentionally left out the design elements (e.g., competition) which are known to increase engagement and PA but risk pushing users towards upward comparison and penalizing downward comparison. However, this may have significantly reduced the app's ability to motivate changes in PA. Conversely, if we had decided to incorporate features known to motivate PA effectively, we risk skewing users' preferences and thus jeopardizing the accuracy of the user model.

Our project shows that the tension created by the personalization paradox is especially prominent when combining personalization with behavioral change. To accurately capture and model users' innate social comparison preferences and reactions, we designed the app to be as neutral as possible. By removing design elements such as competition and goals, we aimed to avoid incentivizing upward comparison and penalize downward comparison, and thus to allow users to express their “innate” preference for social comparison. However, this may have significantly reduced the app's ability to motivate changes in PA. Conversely, if we had decided to incorporate features known to motivate PA effectively, we risk skewing users' preferences and thus jeopardizing the foundation of personalization. For more details on these experiments, the reader is referred to [29].

Another approach that is worth studying is to further separate the user modeling and adaptation stages of personalization. For example, our app could first, during a period of time, stay neutral as it is in our study to collect accurate user data on their social comparison tendencies. Once it has a robust model, then the app could unlock other design features to explicitly motivate PA. However, more research is needed to see if users' social comparison tendencies modeled in a neutral environment can be transferred into a different context with new features that may reward upward comparison. As this research matures, we may want to strike a different balance to have stronger behavioral outcomes. As personalization becomes increasingly used in areas associated with behavior change further research is needed to balance personalization model accuracy and behavioral change effectiveness.

ACKNOWLEDGMENTS

This work is partially supported by the National Science Foundation (NSF) under Grant Number IIS-1816470.

REFERENCES

- [1] Elizabeth Aguirre, Dominik Mahr, Dhruv Grewal, Ko de Ruyter, and Martin Wetzels. 2015. Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of retailing* 91, 1 (2015), 34–49.
- [2] Danielle Arigo, Megan M Brown, Kristen Pasko, and Jerry Suls. 2020. Social Comparison Features in Physical Activity Promotion Apps: Scoping Meta-Review. *Journal of Medical Internet Research* 22, 3 (2020), e15642.
- [3] Kenneth C Arnold, Krysta Chauncey, and Krzysztof Z Gajos. 2020. Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 128–138.
- [4] Richard Bartle. 1996. Hearts, clubs, diamonds, spades: Players who suit MUDs. *Journal of MUD research* 1, 1 (1996), 19.
- [5] Karina Caro, Yuanyuan Feng, Timothy Day, Evan Freed, Boyd Fox, and Jichen Zhu. 2018. Understanding the effect of existing positive relationships on a social motion-based game for health. In *Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare*. 77–87.
- [6] Nitesh V Chawla and Darcy A Davis. 2013. Bringing big data to personalized healthcare: a patient-centered framework. *Journal of general internal medicine* 28, 3 (2013), 660–665.
- [7] Anders Drachen, Alessandro Canossa, and Georgios N Yannakakis. 2009. Player modeling using self-organization in Tomb Raider: Underworld. In *2009 IEEE symposium on computational intelligence and games*. IEEE, 1–8.
- [8] Anushay Furqan, Chelsea Myers, and Jichen Zhu. 2017. Learnability through adaptive discovery tools in voice user interfaces. In *Proceedings of the 2017 CHI Conference*. 1617–1623.
- [9] Anna Goy, Liliana Ardisson, and Giovanna Petrone. 2007. Personalization in e-commerce applications. In *The adaptive web*. Springer, 485–520.
- [10] R C Gray, J Zhu, D Arigo, E Forman, and S Ontañón. 2020. Player Modeling via Multi-armed Bandits. In *in Proceedings of the 15th International Conference on the Foundations of Digital Games*.
- [11] R C Gray and S Zhu, Ontañón. 2020. Regression Oracles and Exploration Strategies for Short-Horizon Multi-Armed Bandits. In *in Proceedings of the 2020 IEEE Conference on Games (CoG 2020)*.
- [12] Eun Young Ha, Jonathan P Rowe, Bradford W Mott, and James C Lester. 2011. Goal recognition with Markov logic networks for player-adaptive games. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [13] Brent Harrison and David L Roberts. 2011. Using sequential observations to model and predict player behavior. In *Proceedings of the 6th International Conference on Foundations of Digital Games*. ACM, 91–98.
- [14] Carrie Heeter, Yu-Hao Lee, Ben Medler, and Brian Magerko. 2011. Beyond player types: gaming achievement goal. In *ACM SIGGRAPH 2011 Game Papers*. ACM, 7.
- [15] Kristina Höök. 1998. Evaluating the utility and usability of an adaptive hypermedia system. *Knowledge-Based Systems* 10, 5 (1998), 311–319.
- [16] Anthony Jameson. 2007. Adaptive interfaces and agents. In *The human-computer interaction handbook*. CRC Press, 459–484.
- [17] Martin Jennings-Teats, Gillian Smith, and Noah Wardrip-Fruin. 2010. Polymorph: dynamic difficulty adjustment through level generation. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*. ACM, 11.
- [18] Pavan Kantharaju, Katelyn Alderfer, Jichen Zhu, Bruce Char, Brian Smith, and Santiago Ontañón. 2019. Tracing player knowledge in a parallel programming educational game. In *Proceedings of the 15th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'19)*.
- [19] Robert Kass. 1989. Student modeling in intelligent tutoring systems—implications for user modeling. In *User models in dialog systems*. Springer, 386–410.
- [20] Joseph Kramer, Sunil Noronha, and John Vergo. 2000. A User-Centered Design Approach to Personalization. *Commun. ACM* 43, 8 (2000), 44–48.
- [21] Olana Missura and Thomas Gärtner. 2009. Player modeling for intelligent difficulty adjustment. In *International Conference on Discovery Science*. Springer, 197–211.
- [22] Safiya Umoja Noble. 2018. *Algorithms of oppression: How search engines reinforce racism*. nyu Press.
- [23] Cathy O'neil. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.
- [24] Eli Pariser. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
- [25] Manu Sharma, Santiago Ontañón, Manish Mehta, and Ashwin Ram. 2010. Drama management and player modeling for interactive fiction games. *Computational Intelligence* 26, 2 (2010), 183–211.
- [26] David Thue, Vadim Bulitko, and Marcia Spetch. 2008. PaSSAGE: A Demonstration of Player Modeling in Interactive Storytelling. In *AIIDE*.
- [27] Josep Valls-Vargas, Santiago Ontañón, and Jichen Zhu. 2015. Exploring player trace segmentation for dynamic play style prediction. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [28] Georgios N Yannakakis, Pieter Spronck, Daniele Loiacono, and Elisabeth André. 2013. Player modeling. (2013).
- [29] Jichen Zhu, Diane H. Dallal, Robert C. Gray, Jennifer Villarealle, Santiago Ontañón, Evan M. Forman, and Danielle Arigo. 2021. Personalization Paradox in Behavior Change Apps: Lessons from a Social Comparison-Based Personalized App for Physical Activity. In *2021 ACM Conference on Computer Supported Cooperative Work and Social Computing*.
- [30] Jichen Zhu and Santiago Ontañón. 2020. Player-Centered AI for Automatic Game Personalization: Open Problems. In *Proceedings of the Fifteenth International Conference on the Foundations of Digital Games*.