The WEDF Student Initiative Submission

Towards a Normative Evaluation of Personalization

Travis Greene
Institute of Service Science
National Tsing Hua University
Hsinchu, Taiwan
travis.greene@iss.nthu.edu.tw

1. Introduction

Digital reality is no longer clearly separable from reality. Philosophers, marketing scholars, and media theorists now speak of a *posthuman* world of *capta*, *data shadows*, *dividuals*, *assemblages*, and *inforgs* (Floridi, 2011; Kitchin and Dodge, 2011). Recognizing this digital shift, governments have enacted new and far-reaching data privacy and protection regulations. Meanwhile, professional societies (e.g., the IEEE) and international organizations have churned out a sea of AI (artificial intelligence) and ML (machine learning) ethics frameworks, principles, and guidelines (Mittelstadt, 2019). Likewise, human subjects research ethics and practices are being re-evaluated as the lines between *data* and *people*, *computer science* and *social science*, increasingly blur (Shmueli, 2017).

ML drives *personalization*, and behavioral big data (BBD) provide the material from which predictive models are built. These predictive models have immense business value. Netflix reports its personalized recommendations have contributed to over $1B in revenues (Gomez-Uribe and Hunt, 2015). Yet despite the lip service paid to personalization, very little consensus exists about what it is and what it should be.

According to philosopher of information Luciano Floridi, personalized recommendations, as a form of information technology, shape our online and offline experience by influencing 1) our self-conceptions; 2) our mutual interactions; 3) our conception of reality; and 4) our interactions with reality (Floridi, 2013). Floridi asks us to consider how technology impacts “who we are, who we think we are, who we might become, and who we think we might become” (Floridi, 2011).

As such, personalization is more than a business or design strategy—it raises deep ethical and metaphysical questions with important social and political ramifications (Knijnenburg et al., 2016; Milano et al., 2019; Paraschakis, 2017; Rossi and Mattei, 2019; Stray et al., 2020). But only now have these issues trickled into personalization research. For example, Zanker et al. (2019) urgently note the need to examine the deeper societal and scientific “pitfalls” deriving from personalization in a post-General Data Protection Regulation (GDPR) world.

We thus offer a new, humanistic framework for evaluating the degree of personalization, relying on the existing ethical and legal norms of the GDPR. This preliminary conceptual framework can guide regulators, professional groups, and practitioners in assessing current levels of personalization and tailoring future regulation to new algorithms and data collection methods.

2. The GDPR and the Subjective Turn in Data Representation

The European Union’s 2018 GDPR can be seen as a political and legal counterbalance to the increasing datafication of society (Greene et al., 2019). Two crucial legal notions underlie the GDPR’s attempt to preserve human dignity in the digital age: *informational self-determination* and the *right to the free development of one’s personality*. Informational self-determination is “an individual’s control over the data and information produced about him,” and is necessary for human self-determination (Rouvroy and Poullet, 2009). We note this view is shared by the IEEE’s Ethically Aligned Design (EAD), which empowers “individuals to curate their identities and manage the ethical implications of their data” (IEEE Standards Association, 2019). Self-determination is a precondition for “a free democratic society based on its citizens’ capacity to act and to cooperate” (Rouvroy and Poullet, 2009). Finally, the right to personality is based on the German Constitution and protects persons’ freedom to develop their personalities as long as they do not violate the rights of others, the constitution, or the moral code (Coors, 2010).

The GDPR expresses various themes from the European (Continental) philosophical tradition. The right to informational self-determination allows individuals to *subjectively* create their digital representations and express themselves. Further, data subjects’ rights to access, delete, and modify their personal data grant them the ability to both more deeply understand themselves and subjectively narrate their personal identities over time. The right to modify one's personal data to fit one's self-narrative provides an expression of *agency* to modify the description under which one's recorded behaviors are understood (see Figure 1 below). Further, rights to access and download one's personal data in portable, machine-readable formats permit the conscious *reflective endorsement* (Bratman, 2000) of in-app and on-device behaviors. Requirements for explicit consent to processing also ensure reflective endorsement.

Lastly, the *right to be forgotten* mirrors the natural workings of autobiographical memory, in which the act of forgetting is essential in constructing a self-narrative over time. Persons actively take part in constructing self narratives to understand themselves, their behavior, and their roles in society (McAdams, 1996). Our narrative identities emerge through a continual cognitive process by which human experience is shaped into "temporally meaningful episodes" (Polkinghorne, 1988).

3. The Evolution and Current State of Personalization

In context of the adaptive web in the late 1990s and early 2000s, three main personalization research streams emerged (Brusilovsky and Maybury, 2002). The first considered various taxonomies and recommender system (RS) designs, some of which considered business strategies and privacy trade-offs (Awad and Krishnan, 2006; Burke, 2002; Fan and Poole, 2006). The second aimed to differentiate personalization from customization and its potential for identifying "long-tail" customers (Sundar and Marathe, 2010). And the third considered various kinds of data—implicit or explicit—as input to personalized RS (Nichols, 1998).

But as computation became cheaper and new sources of data became available, implicit data collection became the focus (Zuboff, 2019). Not only is BBD cheaper and easier to collect, especially by major platforms such as Google and Facebook—who have access to the online behaviors of millions of people—but many in industry believe it is *more* predictive and revealing of underlying preferences than explicit data (Stephens-Davidowitz, 2017). Deep learning and reinforcement learning approaches to personalization are now commonplace. Today the discourse of personalization centers around *preference*, *interest*, and *taste* elicitation, prediction and inference (Pu et al., 2011).

What is Personalization? Despite its economic value and large literature on applications, personalization remains conceptually ambiguous. In the paper, "What Is Personalization?"

Fan and Poole (2006) classify various approaches to personalization, but overlook a simple, but fundamental point: the user is a *person*. Further, their analysis must be updated to remain relevant in today's BBD-era. Some researchers see the difference between customization and personalization as one of *who* initiates the selection of content: the system (system-induced personalization) or the user herself (user-induced customization) (Sundar and Marathe, 2010).

Those developing personalized RS also struggle to define it consistently: Cremonesi et al. (2010) say what personalization is *not*: giving "any user a pre-defined, fixed list of items, regardless of his/her preferences." Jannach et al. (2010, pg.1) say it is assigning "different" lists to users based on their "tastes." In some cases, personalization appears to be equivalent to *profiling*: the assignment of a particular group-derived (often behavioral) profile to each individual in a dataset (Tondello et al., 2017).

3.1. Pitfalls of Personalization

The literature on personalization pays little to no attention to the *normative* or *ethical* aspects of personalization beyond privacy (see, e.g., Adomavicius and Tuzhilin (2005)). Few have tried to align personalization with GDPR and EAD normative principles. Below we discuss where current approaches in personalization fall short of the ideals of the GDPR and EAD.

The person as a feature vector: Figure 1 illustrates the process of narrowing the person into a feature vector. "Important" attributes are only those that contribute to the algorithm's predictive accuracy—predictive goals dictate that certain aspects of human behavior are valued more than others. For example, only the behaviors most amenable to measurement are recorded (clicks, likes, swipes, etc.). This leads to selection bias in the representation of the person. Unpredictable behavior is deemed noise, and measured behaviors which do not add to predictive accuracy are redundant. In contrast, the GDPR and EAD grant data subjects the power and agency to determine "important" factors for personalization based on their subjective meaning, not predictive power.

Nonconscious ideomotor behaviors: Ideomotor actions are "movements or behaviours that are unconsciously initiated, usually without an accompanying sense of conscious control" (Gauchou et al., 2012). They are typically triggered by environmental cues expressing information not consciously accessed. We claim many of the "prefer ences," "interests," and "tastes" personalization claims to predict are *unreflective* in nature. Doing justice to the

GDPR and EAD requires a shift from representing persons using an implicit, behavior-based paradigm (Ekstrand and Willemsen, 2016) dominated by non-conscious “organismic” interests, to one centered on conscious, reflective, explicit feedback, meaning and intentionality.

Aggregating across persons: Personalization often uses data not only from the focal user, but from other users. For instance, personalization based on social network data relies on direct and indirect connections to other users and groups to infer a focal user’s interests. This raises several questions. Does the “community” of nearest neighbors reflect one’s *chosen* social identity? Should personalization include the use of *aggregate* user data to make predictions, or should it only be based on the focal user’s data? The answer will affect the appropriate level of explanation given (nomothetic or idiographic). Further, personalization using aggregate data across persons rests on a “pseudoempirical" assumption (Smedslund, 1991): an unjustified metaphysical claim that abstract groups of persons (populations), and not specific, concrete individuals are scientific objects of interest.

3.2. Personalization That Takes the *Person* Seriously

We suggest the machine learning concept of *personalization* take its etymology seriously. A person could be a “human animal,” “moral agent,” “rational, self-conscious subject,” “possessor of particular rights,” or “being with a defined personality or character" (Schechtman, 2018). *Person* comes from the Latin *persona*, which derives from the Ancient Greek word for a type of mask worn by dramatic actors. Thus the concept of a person is inherently connected to one’s social roles. One’s persona is a specific kind of self-identity that is *public, socially-defined*, and varies depending on context. Western thinkers, following Kant, expanded the idea of the person into two main parts: an outer personnage of public roles and an inner conscience and consciousness, complete with a moral identity and phenomenology off-limits to outside observers (Fukuyama, 2018).

Yet current personalization research fails to recognize the representational duality of the self-conscious person and thus fails to properly understand the concept of context in human experience. Persons exist and act in two epistemically asymmetric inner and outer domains—roughly corresponding to the *emic* and *etic* perspectives used in ethnography—making them different from the objects studied in the natural sciences (Blumer, 1986; MacIntyre, 1985;

Taylor, 1980). These ontological differences necessitate methodological differences (Dilthey, 1989; Habermas, 2015). Epistemological differences further relate to explanations: by failing to distinguish between stimulus-response behavior and intentional action, for instance, a *blink* and a *wink* are indistinguishable and would be (inappropriately) explained in identical ways.

Even highly-cited papers about “context aware personalization,” (e.g., Adomavicius and Tuzhilin (2011)), misunder stand the problem of context. Epistemic asymmetry introduces an ethical problem for personalization which relies solely on *other-defined* BBD at the expense of subjective, interpretive input from persons themselves. A person’s digital behaviors are measured and collected, but their *meaning* and *identity* derive one-sidedly from data controllers. There is no single “correct” representation or encoding of BBD—our shared “forms of life” ultimately ground such questions of meaning and identity of observed regularities (Wittgenstein, 2009). Instead, there are simply different representations under different interpretations about what counts as what.

Kelly and Teevan (2003) conclude more personalization research should focus on “what observable behaviors *mean* and how they change with respect to *contextual factors*." We claim the GDPR fills this interpretive vacuum by viewing informational self-determination as necessary to protect human dignity. Absent their original intentional context, behavioral data lose their subjective meaning, essentially “making numbers that appear to be identical actually different from each other” (boyd and Crawford, 2012). Seaver (2015) laments how “context can be simultaneously missing from data science and central to it.” We see a similar pattern: the *person* has been paradoxically missing from *person*alization.

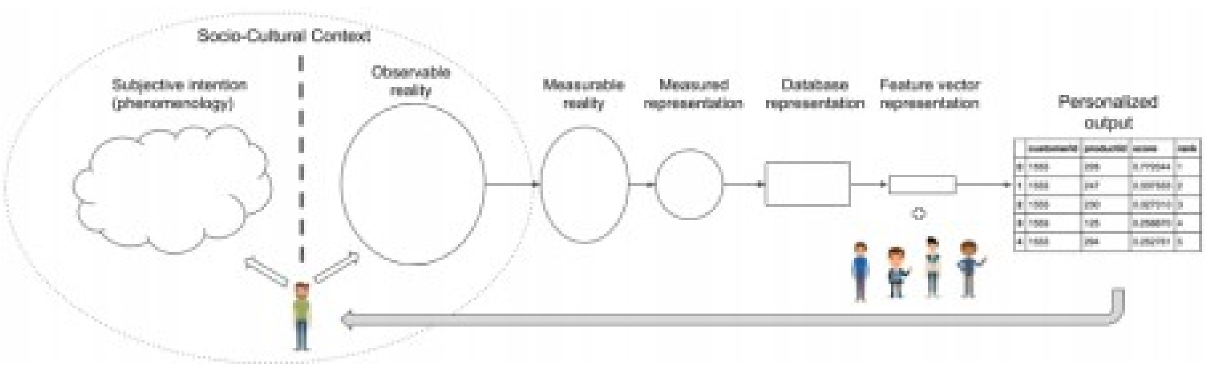

Figure 1: The personalization feedback loop. Personalization narrows a person embedded in subjective, social, and cultural space to an objective feature vector embedded in feature space. Intentionality is ignored in favor of easily measurable behaviors.

4. Towards a Normative Evaluation of Personalization

We hope to move towards a normative understanding of personalization founded on the human rights and values expressed in GDPR and EAD principles. This is a major undertaking—much beyond the scope of a short article— and requires a full-fledged "theory of personalization." Below is a sketch of a preliminary conceptual framework that could help regulators, professional groups, and practitioners in evaluating personalization methods and crafting new legislation to preserve these humanistic values.

First, personalization exists on a continuum. Second, the degree of personalization could be evaluated in terms of input data, model, and output from two independent perspectives reflecting the human condition: subjective and objective (Levine, 1983). Objective measures are those nearly any reasonable observer would agree on and are relatively quantifiable. Subjective measures, however, are those that only the target user can judge. Below we propose three examples of GDPR-based constructs.

Self-determination Objective: How many options exist for a user to interact with and control her data representation? Subjective: Does your data representation fit your narrative identity?

Self-expression Objective: How many degrees of freedom exist for logged interactions with the RS? Subjective: How constrained do you feel when interacting with the RS?

Moral salience Objective: Do the recommendations have "significant effects" according to the GDPR's Article 22? Subjective: How important are these recommendations to your moral or narrative identity (Atkins, 2010)?

Ideally subjective and objective measures would be highly-correlated, a state we refer to as "narrative accuracy," while discrepancies between the two measures can reveal areas deserving more research and legislative attention.

References

Adomavicius, G. and Tuzhilin, A. (2005). Personalization technologies: a process-oriented perspective. *Communi cations of the ACM*, 48(10):83–90.

Adomavicius, G. and Tuzhilin, A. (2011). Context-aware recommender systems. In *Recommender systems handbook*, pages 217–253. Springer.

Atkins, K. (2010). *Narrative identity and moral identity*. Taylor & Francis.

Awad, N. F. and Krishnan, M. S. (2006). The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS quarterly*, pages 13–28. Blumer, H. (1986). *Symbolic interactionism: Perspective and method*. Univ of California Press. boyd, D. and Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society*, 15(5):662–679.

Bratman, M. E. (2000). Reflection, planning, and temporally extended agency. *The Philosophical Review*, 109(1):35–61.

Brusilovsky, P. and Maybury, M. T. (2002). From adaptive hypermedia to the adaptive web. *Communications of the ACM*, 45(5):30–33.

Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction*, 12(4):331–370.

Coors, C. (2010). Headwind from europe: The new position of the german courts on personality rights after the judgment of the european court of human rights. *German Law Journal*, 11(5):527–537. Cremonesi, P., Koren, Y., and Turrin, R. (2010). Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM.

Dilthey, W. (1989). *Introduction to the human sciences*, volume 1. Princeton University Press. Ekstrand, M. D. and Willemsen, M. C. (2016). Behaviorism is not enough: better recommendations through listening to users. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pages 221–224. Fan, H. and Poole, M. S. (2006). What is personalization? perspectives on the design and implementation of personalization in information systems. *Journal of Org. Comp. and Elec. Comm.*, 16(3-4):179–202. Floridi, L. (2011). The informational nature of personal identity. *Minds and machines*, 21(4):549. Floridi, L. (2013). *The ethics of information*. Oxford University Press.

Fukuyama, F. (2018). *Identity: Contemporary identity politics and the struggle for recognition*. Profile books. Gauchou, H. L., Rensink, R. A., and Fels, S. (2012). Expression of nonconscious knowledge via ideomotor actions. *Consciousness and cognition*, 21(2):976–982.

Gomez-Uribe, C. A. and Hunt, N. (2015). The netflix recommender system: Algorithms. *Business Value*, 6. Greene, T., Shmueli, G., Ray, S., and Fell, J. (2019). Adjusting to the gdpr: The impact on data scientists and behavioral researchers. *Big data*, 7(3):140–162.

Habermas, J. (2015). *Knowledge and human interests*. John Wiley & Sons.

IEEE Standards Association (2019). The IEEE global initiative on ethics of autonomous and intelligent systems. ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems.

Jannach, D., Zanker, M., Felfernig, A., and Friedrich, G. (2010). *Recommender systems: an introduction*. Cambridge University Press.

Kelly, D. and Teevan, J. (2003). Implicit feedback for inferring user preference: a bibliography. In *Acm Sigir Forum*, volume 37, pages 18–28. ACM New York, NY, USA.

Kitchin, R. and Dodge, M. (2011). *Code/space: Software and everyday life*. Mit Press. Knijnenburg, B. P., Sivakumar, S., and Wilkinson, D. (2016). Recommender systems for self-actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pages 11–14.

Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific philosophical quarterly*, 64(4):354–361. MacIntyre, A. (1985). How psychology makes itself true-or false. In S. Koch, D. L., editor, *A Century of Psychology As Science*, pages 897–903.

McAdams, D. P. (1996). Personality, modernity, and the storied self: A contemporary framework for studying persons. *Psychological inquiry*, 7(4):295–321.

Milano, S., Taddeo, M., and Floridi, L. (2019). Recommender systems and their ethical challenges. *Available at SSRN 3378581*.

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical ai. *Nature Machine Intelligence*, pages 1–7. Nichols, D. (1998). Implicit rating and filtering. In *Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering*, pages 31–36. ERCIM.

Paraschakis, D. (2017). Towards an ethical recommendation framework. In *2017 11th International Conference on Research Challenges in Information Science (RCIS)*, pages 211–220. IEEE.

Polkinghorne, D. E. (1988). *Narrative knowing and the human sciences*. Suny Press.

Pu, P., Chen, L., and Hu, R. (2011). A user-centric evaluation framework for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pages 157–164.

Rossi, F. and Mattei, N. (2019). Building ethically bounded ai. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9785–9789.

Rouvroy, A. and Poullet, Y. (2009). The right to informational self-determination and the value of self-development: Reassessing the importance of privacy for democracy. In *Reinventing data protection?*, pages 45–76. Springer. Schechtman, M. (2018). *The constitution of selves*. Cornell university press.

Seaver, N. (2015). The nice thing about context is that everyone has it. *Media, Culture &*

Society, 37(7):1101–1109. Shmueli, G. (2017). Research dilemmas with behavioral big data. *Big data*, 5(2):98–119. Smedslund, J. (1991). The pseudoempirical in psychology and the case for psychologic. *Psychological Inquiry*, 2(4):325–338.

Stephens-Davidowitz, S. (2017). *Everybody lies*. Harper Collins.

Stray, J., Adler, S., and Hadfield-Menell, D. (2020). What are you optimizing for? aligning recommender systems with human values. In *ICML 2020 Workshop on Participatory Approaches to Machine Learning*. Sundar, S. S. and Marathe, S. S. (2010). Personalization versus customization: The importance of agency, privacy, and power usage. *Human Communication Research*, 36(3):298–322.

Taylor, C. (1980). Understanding in human science. *The Review of Metaphysics*, 34(1):25–38. Tondello, G. F., Orji, R., and Nacke, L. E. (2017). Recommender systems for personalized gamification. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 425–430. Wittgenstein, L. (2009). *Philosophical investigations*. John Wiley & Sons.

Zanker, M., Rook, L., and Jannach, D. (2019). Measuring the impact of online personalisation: Past, present and future. *International Journal of Human-Computer Studies*, 131:160–168.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books.