Towards Societally Beneficial Personalized Realities: A Conceptual Foundation for Responsible Ubiquitous Personalization Systems

Jannis Strecker
University of St. Gallen
St. Gallen, Switzerland
jannisrene.strecker@unisg.ch

Simon Mayer
University of St. Gallen
St.Gallen, Switzerland
simon.mayer@unisg.ch

Kenan Bektaş
University of St. Gallen
HOCH Health Ostschweiz
St.Gallen, Switzerland
kenan.bektas@unisg.ch

Abstract

Personalization of online realities is today ubiquitous to support decision making or reduce information overload. Recently, through the expanding capabilities and pervasiveness of Mixed Reality and Ubiquitous Computing technologies, we observe increasing personalization also of physical reality. This might yield more convenient, efficient and inclusive everyday interactions. However, it may readily lead to serious societal consequences such as the loss of shared worlds and the emergence of perceptual filter bubbles. To mitigate such harms while retaining the benefits of personalization, it is important to understand how ubiquitous personalization systems may operate responsibly. Responding to this need, we propose a conceptual model that overcomes the limitations of established personalization models and expands their applicable scope to physical, virtual, and hybrid environments. We validated our model in relation to existing literature and show how it provides a conceptual foundation for the analysis and study of responsible personalization systems that create individually and societally beneficial Personalized Realities.

CCS Concepts

• **Information systems → Personalization**; • **Human-centered computing → Mixed / augmented reality**; **Ubiquitous and mobile computing systems and tools**.

Keywords

personalization, mixed reality, societal implications, mediated reality, conceptual model, responsible computing

ACM Reference Format:
Jannis Strecker, Simon Mayer, and Kenan Bektaş. 2025. Towards Societally Beneficial Personalized Realities: A Conceptual Foundation for Responsible Ubiquitous Personalization Systems. In *Designing Interactive Systems Conference (DIS '25), July 5–9, 2025, Funchal, Portugal.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3715336.3735709

1 Introduction

The primary focus of personalization research and practice is mostly on content that users access virtually through Web browsers or mobile applications, such as social media feeds [80] or shopping experiences [6]. Already today, however, personalization impacts people's lives beyond virtual environments [207], e.g., by influencing what services they use, what things they buy, or what places they visit. Additionally, passive personalization technologies, such as sunglasses and prescription glasses, modify people's perception of physical reality in ways that are—nowadays—culturally and societally so well accepted that they are not considered as reality mediators. Such personalization is increasingly joined by emerging—and socially not yet as accepted [68]—*active* mediation of physical reality, e.g., through Mixed Reality (MR) head-mounted displays, that enable the dynamic personalization of individual perceptions of physical reality in real-time. Additionally, Ubiquitous Computing (UbiComp) technologies are extending the application space of personalization by allowing individuals to interact with smart environments that dynamically adapt to their needs. This already includes, e.g., home automation and lighting systems that adapt to match a user's task or mood [145] and may expand to systems that seem more alien today: utilizing hardware such as ChainFORM [126], door handles might change their shape to fit an individual's hand's grip, or they might actively draw attention when the system wants to motivate a user to leave for their next meeting—or to pick up their children from kindergarten.

The dynamic personalization of people's reality—directly or indirectly—may address diverse sensory modalities (cf. [65, 146, 179, 213, 215]) and it is hence conceivable that systems may personalize literally anything that can be perceived by humans through their sensory organs. The increasing availability and capabilities of MR and UbiComp technologies [66, 136, 159] provide the technical means to integrate such personalization more and more directly with *physical* realities. Previously, we introduced the term *Personalized Reality (PR)* [185] to describe the perceivable output of such *ubiquitous personalization (UP)* systems. On the one hand, UP allows to bring known benefits of personalization such as equal access for users with diverse abilities [57, 103], reduction of information overload [24], and better preference matching [200] to a wider range of applications in physical reality. Possible applications of such ubiquitous personalization include the personalization of navigational cues [171], expertise-adapted instructions for industrial workflows [56], allergy- or health-oriented product recommendations in supermarkets [187], personalized learning and coaching (cf. [135, 197], worker support through personalized movements of industrial robots (cf. [81]), and personalized assistance to help people overcome social barriers (cf. [114]).

DIS '25, Funchal, Portugal

ACM ISBN 979-8-4007-1485-6/2025/07
https://doi.org/10.1145/3715336.3735709

On the other hand, all-encompassing personalization could evoke harmful effects on individuals and society such as undermining the social co-construction of a common reality [14] and the loss of *shared worlds* [24], through the creation of "perceptual filter bubbles" that situate users in isolated and fragmented personalized perceptions of reality where elements in the physical or virtual context that contradict their (e.g., political) beliefs are selectively suppressed (cf. [1, 163]). Investigations and visions of *pervasive MR* (e.g., by wearing MR glasses that overlay virtual content on physical reality continuously) [11, 66, 87, 128, 132, 143, 159, 209] anticipate similar issues and harms such as a loss of control [87], possibilities for manipulation and deception [47, 118], or isolated perceptions of reality [1, 163, 186]. However, these concepts mention personalization only as one of many examples (e.g., personalized AR advertising [5], or personalized knowledge display [132]) without considering the concrete implications of widespread UP. As physical reality with its shared worlds and physically-grounded experiences constitutes the foundation of how people understand and communicate about reality, and as researchers and practitioners anticipate that the integration of personalization into physical reality will even accelerate in the present and future [46, 83, 107, 110, 115, 124, 185], e.g., through AI-enabled methods [79], there is an urgency to implement effective countermeasures that will permit realizing positive outcomes of UP while avoiding or mitigating negative consequences.

To address potential harmful outcomes of personalization systems, existing personalization research has proposed frameworks to design personalization algorithms based on ethical principles [64, 141] and has investigated how different types of recommender systems affect individuals and society [138, 194]. However, this focus on the algorithm misses the notion that personalized content is consumed by persons and potential bystanders in a specific context with a specific device. Thus, recent research has requested a more holistic perspective that includes these aspects and considers the wider implications of personalization systems on individuals and society (cf. [97, 113]), i.e. *responsible personalization.* Furthermore, to study the potential implications of personalization systems, researchers have created conceptual models that capture the relevant mechanisms and relationships between the involved components (cf. [22, 48, 59, 88, 111, 201]). However, these models focus on traditional (Web) personalization and do not adequately cover personalization systems that are deployed using current and emerging technologies; they are hence not anymore fit for a structured analysis of personalization phenomena in this greatly extended space. Instead, to understand the phenomenon of UP and its implications on individuals and society—and to analyze concrete personalization systems in this scope—a new conceptual model is needed that can describe UP systems in physical, hybrid, and virtual environments.

Responding to this need, in this work, we propose a conceptual model for responsible UP systems (RUPS) that accounts for the described expansion of the scope of personalization systems and incorporates the notion of Personalized Reality [185] as the output of such systems. Our model is based on existing personalization research and extends widely-accepted personalization models to comprehensively describe personalization systems. We validate the RUPS model with respect to its coverage as well as its predictive and explanatory capabilities by mapping existing personalization systems to its components. We further explain how we intend the model to be used by researchers, practitioners, and regulators. This work thus provides a conceptual foundation for the study of any kind of personalization system.

2 Related Work

The term *personalization* is commonly used when referring to the processing of personal data by a system as input, the adaptation of the system's functioning in response to personal data, and the personalized content that such a system outputs (cf. [24, 52, 166, 196, 200]). The use of *personal data* is hence central, and we follow the "future-proof" [185] definition of the European Union's General Data Protection Regulation (GDPR) of personal data as "[...] any information relating to an identified or identifiable natural person [...] who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person" [50, p.33].

The terms *personalization* and *customization* are often used interchangeably and many applications feature a combination of both; *customization* typically denotes a user-initiated or -controlled adaptation [52, 166, 189] (i.e. a system is *adaptable* [134]) while *personalization*, in contrast, is a system-initiated or -controlled adaptation [52, 166, 189] (i.e. a system is *adaptive* [134]), e.g., when a social media platform adapts the selection and order of posts in a user's timeline. The term *hyper-personalization* is sometimes used to refer to instances of personalization that use big data and machine-learning methods, e.g., to target individuals with tailored ads [76, 118]. On another dimension, personalization is often used synonymously with *Web Personalization* [166]. On the Web, personalization is widespread and applied, e.g., to search results [182], news [210], advertisements [45], music [23] or video [7] recommendations, social media feeds [88], or shopping experiences [6]. In this work, we focus on personalization as system-initiated adaptations in adaptive systems that process personal data, as described above. This subsumes Web personalization as well as hyper-personalization since we do not prescribe the delivery medium nor the amount of personal data that is needed or the methods for processing the data.

2.1 Personalization and Society

Research highlights the benefits of personalization, such as improved experiences [196, 200], or support for diverse user needs [103], but also raises concerns about users' growing difficulty in distinguishing personalized from non-personalized content [24] often because the fact that a certain content is personalized is not disclosed to users [89]. While some claim that personalization may lead to increased user autonomy [21], others warn that its consequences include a perceived or actual loss of control for users [196], privacy risks [200], and manipulation possibilities [214].

On a societal level, the potential of personalization technologies to amplify polarization through the formation of echo chambers and filter bubbles is debated for (traditional) personalization systems such as search engines or social media platforms [27, 67, 193]. However, researchers expect that *perceptual* filter bubbles and more isolated views on reality will likely emerge when personalization is increasingly present in people's physical realities [1, 83, 163]. This relates to the concern of eroding *shared worlds* [24] and negative effects on *intersubjectivity* (i.e. "the common-sense, shared meanings constructed by people in their interactions with each other and used as an everyday resource to interpret the meaning of elements of social and cultural life" [170, p.1126]) through personalization. Recent studies, e.g., show that already the personalization of online news content leads to a loss of shared news rituals, and thus undermines shared worldviews [108]. As personalization is "fundamentally concerned with individual rather than collective or communitarian life" [96, p.43], it may hence foster increasingly unconnected *Fragmented Realities* [74]. While Fragmented Realities per se describe the (neutral) philosophical realist notion that different people's realities might be based on different subjective truths for one objective truth, the considered proliferation of personalization might reinforce the belief in one's own fragment as the objective truth without regarding others' as equally valid and without supporting the creation of common or connected truths and realities [14]. Thus, as personalization impacts how people perceive and communicate about their realities, it may be considered ubiquitous in a *social* sense already today.

2.2 Towards the Personalization of Physical Reality

In a *technical* sense, the increasing ubiquity of personalization is enabled through MR and Ubicomp technologies. Milgram et al. defined MR on the *Reality-Virtuality (RV) continuum* as the (visual) segment in-between Real and Virtual Environments [120, 121] where "real world and virtual world objects are presented together within a single display" [121, p.283]. Their formulation implies that virtual environments are not 'real' in the same sense as physical environments. While the 'realness' of virtual realities and objects with respect to physical reality and objects is debated by philosophers [28, 116, 175], we use the term *physical reality* when referring to physical objects and environments (i.e., the left end on the RV continuum), in-line with Chalmers [28]. Skarbez et al. extend Milgram et al.'s notion of MR to include all technology-mediated realities, including, e.g., eXtended Reality (XR) and Virtual Reality (VR) [173, 174]. While other definitions of MR exist (e.g., [157, 177]), we use the term MR when referring to any kind of mediated reality that contains physical and virtual objects and stimuli in a single percept, following Skarbez et al. [174]. In addition to the personalization of physical objects in our environment and the (well-known) personalization of virtual content (such as Websites), personalization in MR may make use of Augmented Reality (AR), but also seamless interactions between physical and virtual realities [18, 35, 49, 162], and Augmented Virtuality (AV), in which, e.g., cues and objects from physical reality are included in VR environments [12, 117, 192, 204], or virtual reconstructions of physical reality compose (part of) a VR experience [106, 172]. While MR is often only associated with hand-held or worn (visual) displays, such as HMDs or smartphones, it may stimulate all human senses [174], thus, UP may be delivered through any sensory modality that is compatible with the intended users, e.g., through vestibular [179] and electromuscular stimulation [20] as well as auditory [213], tactile [151, 215], tastable [156], and olfactory interfaces [99].

In addition, connected devices that are embedded in the environment (i.e., in UbiComp as envisioned by Weiser [206]) may personalize a users' physical reality *directly*, if these devices have properties that can be dynamically adapted by a system, without relying on mediators such as MR devices. Recent examples of such *direct* personalization include personalized interactions with smart coffee machines [125], interactions with educational robots as personalized quizmasters [149], humanoid robots that prepare personalized drinks based on interactions with a user [86], self-actuating furniture that automatically improves its ergonomics for a specific user [211], and self-balancing bicycles that optimize their users' experience of gravity [208].

The direct or mediated personalization of a user's reality might require information about their context and personalization systems are hence typically *context-aware*: they use contextual information as input and adapt their interface or behavior accordingly [42, 43] to deliver the right services at the right time [94]. In mobile settings, such systems can dynamically adapt their behavior and adjust the form and amount of information they provide with respect to the physical environment (e.g., conditions, infrastructure, location) and human factors [169] such as user or bystander activity [17], cognitive load and task [105], or attention and interest [15, 146, 147]. Similar to MR that may stimulate all human senses, also context-aware systems have been becoming more multimodal—e.g., equipped with optical, audio, motion, and biological sensors)—to adapt more specifically to a user's sensed context. However, in contrast to context-aware systems, UP systems necessarily need personal data, as we explained above, and thus context-aware computing frameworks such as the *Context Toolkit* [165] may not be readily applicable.

Summarizing, we observe that traditional personalization systems are joined by approaches to personalize a user's physical reality as well. This may affect all human sensory modalities, and may be implemented through *direct* or *mediated* personalization. Combined, this development provides the basis to *personalize a user's entire experience of their physical, hybrid, and virtual environments*. We refer to this perceivable product of UP systems as *Personalized Reality* (PR): "Personalized Reality describes a physical, virtual, or mixed reality that has been modified in response to personal user data and may be perceived by one or multiple users through any sensory modality." [185, p.2] Yet, currently this definition lacks contextualization in the wider literature, as well as a detailed description of the possible implications stemming from PR.

2.3 Responsible Ubiquitous Personalization

Such an all-encompassing PR places creators of UP systems in a powerful position: They can personalize a space that is not confined to a smartphone or computer screen and thus essentially have the possibility to "determine how users experience the world, how they conceive of themselves, and how they regard others" [112, p.99], and hence essentially influence people's perceptual and conceptual worldview [183]. Therefore, a *responsible* approach to UP is needed to mitigate potential harmful implications for individuals and society. Towards a basic understanding of what the term *responsible* might entail, it its beneficial to look at related disciplines and fields.

Responsible Computing describes the perspective of taking ethical, social, and societal aspects of computing into account during the full life cycle of a technological artifact. This includes considering the implications of a technology on individuals and society [98], creating inclusive and accessible systems [33], and respecting users' rights (e.g., towards data privacy cf.[40, 95]). *Responsible Research and Innovation* [139] denotes an approach to research and innovation that considers societal needs and moral values [44] with an orientation towards social or environmental benefits [190] and the inclusion of all relevant stakeholders [203]. Frameworks for *Responsible Artificial Intelligence (AI)* highlight the need to consider human well-being and autonomy [144], social responsibility [30], and ethical impacts [198] of AI systems. Furthermore, also computing and design associations highlight responsible behavior in their Codes of Ethics, such as the ACM[1], IEEE[2], or UXPA[3]. Existing research on responsibility in personalization systems often only concerns the personalization algorithm, e.g., regarding fairness [104], ethics [141] or biases [22]. A more holistic view on personalization systems that include the algorithm, devices, users, and the environment they are situated in (including possible bystanders and objects), would help to ensure that the design and analysis of responsible personalization systems considers all relevant components and stakeholders.

Responsible ubiquitous personalization in our context thus combines these notions and considers systems' societal and ethical impact, and addresses the concerns of all involved stakeholders to "actively promote human flourishing and autonomy" [64, p.1545] through personalization systems.

3 The Need for a new Personalization Model: From the Web to Physical Reality

In this section, we show why existing methods for studying personalization and its implications are not adequately applicable when also physical reality is personalized. On this basis, we derive the required features for our conceptual model for *responsible* UP systems.

3.1 Personalized Reality May Become a Nightmare Without Responsible UP Systems

Towards estimating the possible implications of PR on individuals and society, concepts that also pervasively mediate physical reality, such as *pervasive MR*, are helpful. We use *pervasive MR* as an umbrella term to refer to concepts such as *Pervasive Augmented Reality* [66], *Ubiquitous MR* [143], *Societal XR* [73], and related proposals (cf. [11, 87, 128, 132, 209]). These approaches are united by a common aim to study the benefits and harms of pervasive MR to ensure a responsible implementation. Especially the expected harms of pervasive MR overlap considerably with those of personalization, as discussed above: Researchers warn that pervasive MR interfaces may nourish the creation of perceptual filter bubbles [1, 163] and may provide ground for manipulation possibilities, e.g. through dark patterns or deception [47, 63, 75, 93, 118, 175]. Thus, possible harms of pervasive MR for foreseeable [1, 68] and envisioned futures [58, 87, 136, 140] are likely to be *amplified* when personalization and pervasive MR are combined in PR. PR might hence create increasingly fragmented and isolated perceptions of reality, if its development is not steered in a responsible direction. Existing speculative visions of ubiquitous PRs [107, 110, 115] give a first glance on how such amplified implications may look like, while also research is beginning to imagine personalization and its impact beyond the Web [38, 46, 83, 124, 132]. However, there is currently no systematic description of responsible UP systems that mitigate dystopian outcomes, e.g., such as depicted in the television series Black Mirror[4]. While researchers have, e.g., envisioned the sharing of content across users to counter potential isolated realities [1, 87, 186], these works do not describe how this idea is situated within a personalization system and what its wider implications may be. A more systematic and holistic view on UP would be needed to describe responsible UP systems and PR, and guide their future development.

3.2 Previous Personalization Models Were Created for the Web

To derive such a systematic description, we consider previous research in *Web personalization* where researchers created conceptual models that capture the relevant mechanisms and relationships between the components involved in a personalization system [22, 48, 59, 88, 111, 201]. These models provide a high-level abstraction of complex systems, guide design and architecture decisions, enable efficient problem detection and solving in the modeled domain, and help to facilitate interdisciplinary communication [61, 77, 130]—they hence should possess explanatory and predictive power regarding specific instances of systems in the modeled domain. We discuss two well-accepted and representative models of Web personalization in detail: Eirinaki and Vazirgiannis [48] capture personalization in the early Web, and Bozdag [22] extends their model to integrate the widespread commercial adoption of Web personalization in the 2010s.

In the year 2003, Eirinaki and Vazirgiannis define Web personalization as "the process of customizing the content and structure of a Web site to the specific and individual needs of each user taking advantage of the user's navigational behavior" [48, p.3]. In their model (see Figure 5 in Appendix A), they consider four types of data, according to the classification of Web data in Srivastava et al. [180], as inputs for Web personalization: *content* (e.g., text or images), intra-page and inter-page *content structure* (i.e., HTML markup and hyperlinks), *usage patterns* (e.g., the date and time of accesses), and *user profiles* (e.g., a user's demographic information). As the Web became an ever-more important business factor over the first two decades of the 21st century, Web applications have increasingly realized the personalization potential inherent in their possession of large amounts of usage data and the large value of adapting their specific content (e.g., for product suggestions [176]). Along with the emergence of the participatory Web and the transformation of users into prosumers [19], this required revisiting Eirinaki and Vazirgiannis' focus. It became necessary to also consider factors

[1]"To act responsibly, [computing professionals] should reflect upon the wider impacts of their work, consistently supporting the public good.", from https://www.acm.org/diversity-inclusion/code-of-ethics. Last accessed January 9, 2025.
[2]https://www.ieee.org/about/corporate/governance/p7-8.html. Last accessed January 9, 2025.
[3]https://uxpa.org/uxpa-code-of-professional-conduct/. Last accessed January 9, 2025.
[4]https://en.wikipedia.org/wiki/Black_Mirror. Last accessed January 10, 2025.

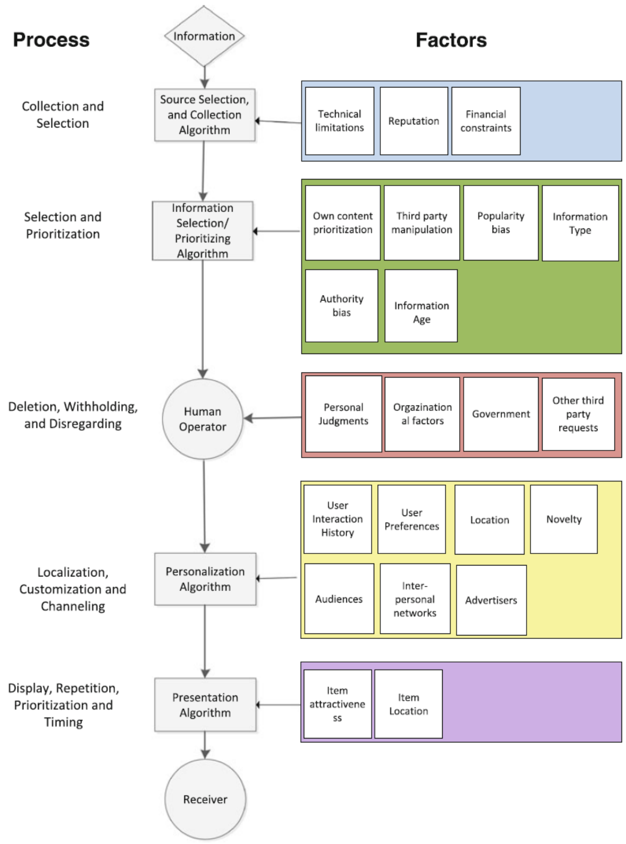

Figure 1: Bozdag's Web personalization model: "A model of filtering for online web services including personalization"(from [22, p.215]). While this model includes the human factor in personalization, it lacks its situatedness in the user's context, the inclusion of different delivery devices, and the circular nature of personalization .

such as biases, regulation, new technological possibilities (e.g., dynamic Web content), and third-party requests. Responding to this need, a more detailed "model of filtering for online web services including personalization" was proposed by Bozdag [22]. Their model (see Figure 1) adds societal (e.g., third-party manipulation), and economical (e.g., own content prioritization) aspects while making explicit the collection, selection, and prioritization of information that is used for personalizing content. The model furthermore considers these processes to be influenced by several additional factors such as popularity and authority biases, and technical limitations. Information that has been selected is then edited by a "Human Operator", who is in turn influenced by personal, organizational, and governmental factors and may be subjected to requests by other third parties—information may consequently be deleted, withheld, or discarded. The next step in Bozdag's modeled personalization process is the algorithmic personalization of the information that is based on factors like user characteristics, target audience, or advertisers. Finally, a presentation algorithm displays the resulting personalized information to the receiver, i.e., the user of the personalized service.

3.3 Shortcomings of Previous Personalization Models Concerning the Personalization of Physical Reality

We argue that both models [22, 48], as well as other models of personalization that we have surveyed [59, 88, 111, 201], are not anymore fit to represent current and emerging personalization systems. Not only do these models require updates and scope extensions that derive from the concept of UP itself, but they also need to integrate the refined notion of *responsibility* that has emerged in research (in HCI and elsewhere, see Section 2.3) over the past decade. To this end, we summarize six main shortcomings of past personalization models; from each shortcoming, we derive a requirement for an new model.

3.3.1 Stakeholder Conflation. The "Human Operator" is a well-circumscribed and distinct entity in Bozdag's model and in other representative personalization models; this perspective has evolved over the past decade and it is today more accurate to represent human decisions as influencing *each* step of a personalization system—including the personalization and presentation algorithms [181], which may contain organizational and cultural biases as well [92]. Current models of personalization furthermore show all processes—from information collection to information presentation—as being carried out by the same entity. This conflation of stakeholders and the modeling of a distinct human operator are understandable given the state of the Web when Bozdag's model was proposed. However, in the time since, the complexity of personalization systems has increased in terms of the number and heterogeneity of stakeholders as well as the systems' depth and scope. Personalization is today "one of the key applications in machine learning with widespread usage across e-commerce, entertainment, production, healthcare and many other industries" [60, p.1] and—not least known through the 2018 Facebook-Cambridge Analytica data scandal—is also used in the context of the microtargeting of political advertisements. Models that do not account for this complexity lose their explanatory and predictive power regarding (societal) phenomena in personalization systems, e.g., when considering cross-stakeholder interactions where the party that delivers the personalization is different from the party that delivers a system's core content. Today, such third-party personalization (e.g., in the case of online advertisements) has however become the default practice even on the Web itself. A contemporary personalization model hence needs to consider that personalization involves a large number of direct and indirect stakeholders.

3.3.2 Virtual-Only Focus. No current conceptual model of personalization considers the possibility that physical and mixed reality are personalized. However, the environment in which the user is situated in physical reality is fundamentally different compared to when they are using the Web. In physical reality, the user's environment serves an objective, unfiltered ground truth (given uniform sensory abilities) that is implicitly shared with others, even when personalized for one user. This implies that the user's context is more important for personalization in physical reality than on the Web where an objective, unfiltered ground truth is typically not available for all users who access a particular Web resource.[5] This implicit sharing in physical environments may provide opportunities for the *sharing* of personalized content *across users* to counter a loss of shared worlds, and for addressing multi-user scenarios. These aspects are not included in current personalization models because (personalized) Web content had been implicitly considered to be consumed individually. We argue that an explicit inclusion of such sharing for multi-user scenarios in a novel personalization model will serve to inform existing traditional as well as emerging personalization systems about sharing opportunities that have been missed, or would be missed, without a new model.

3.3.3 Disregard of Delivery Medium / Assumption of Homogeneous Delivery Media. Because of the uniformity of interaction media in Web personalization (e.g., browsers, mobile phones), past personalization research overall pays little attention to the *delivery medium.* However, in physical reality, the space of possible delivery media increases greatly since not only mediators such as MR head-mounted displays (HMD) but any "smart device" (e.g., lighting systems, car dashboards, collaborative robots) may personalize aspects of a user's physical reality through a multitude of modalities [136]. The possible consequences of such real-world compartmentalization will qualitatively perhaps be similar to current phenomena in digital compartmentalization (e.g., filter bubbles, echo chambers, or microtargeting). In physical reality, however, these implications are likely amplified due to the vastly increased pervasiveness of the personalization space [83, 87, 132, 186]: Any thing in peoples' environment may possibly be modified, emphasized, or hidden—based on personal user data and on the capabilities of the delivery medium. While current personalization models do not explicitly address the delivery medium, this needs to receive more attention in models that are used to study UP.

3.3.4 Limitation to Virtual Data Sources. The possibility that a personalization system is situated in physical or mixed reality also has consequences on the collection of (personal) data that is used to personalize content: On the one hand, data collection in physical reality is more expensive because it typically requires additional sensors and hardware while powerful tracking on the Web (e.g., through Click Trails [91]) can be implemented with few lines of code and remote-deployed instantly. Data collection in physical reality is also less accepted by users since, as Acquisti et al. state, people "rely, in part, on sensorial cues to navigate privacy choices" [3, p.270]. These sensorial cues are present in physical reality (e.g., visible surveillance cameras) but are mostly absent when navigating the Web or using other means of virtual personalization. Therefore, Acquisti et al. warn that "the more we transition from physical to digital interactions, the less equipped we may be for informed digital privacy decisions" [3, p.270]. On the other hand, if additional sensors and hardware are available in physical reality, they facilitate the collection of vast amounts of information about the user's context. While currently few personalization systems such as recommender systems make use of external sensor information, researchers expect this to increase in the near future [85, 153]. MR HMDs, for instance, typically already carry cameras and other sensors that enable capturing context information on many levels of granularity, from the visual detection of objects in the device's field of view [13, 178] to the logging of texts through optical character recognition [184], and the real-time recording and analysis of user gaze data [16]. Beyond sensors on MR HMDs, the Internet of Things is driving the proliferation of ubiquitously available sensors. Through standardization (such as through the World Wide Web Consortium's Web of Things Thing Description[6] standard) these sensors are become readily usable as sources of data from physical reality where, reflecting Weiser [206], "it is desirable that these sensing devices disappear physically, as well as psychologically, motivating the use of thin, deploy-and-forget wireless sensors" [188, p.1] that boost the availability of contextual data sources. Due to this development, and since current personalization models merely consider information about the interaction behavior of the user with digital content [22], we argue that a new model of personalization needs to account for the vast possibilities of data collection in pervasive computing spaces.

3.3.5 Missing Feedback. In traditional personalization systems, personalization mostly happens transactionally. In the example of Web personalization, for instance, a user-induced state change (e.g., the user selecting a product) triggers further personalization of the interface or content. Understandably, traditional personalization models hence treat the phenomenon as a one-way process without feedback. However, algorithmic filtering and the personalization of reality typically work in real time and hence instead heavily rely on real-time user feedback (cf. [78]) that itself is used for future personalization by a system. A novel model of personalization needs to account for this feedback loop, which represents a significant structural requirement to describe contemporary and future personalization systems.

4 A Conceptual Model for Responsible Ubiquitous Personalization Systems

Summarizing, the applicability of accepted models of personalization has hence suffered as personalization systems have proliferated and increased in complexity. Specifically, traditional models are limited when considering that personalization may also target physical reality through MR and UbiComp technologies. In this section, we thus propose a new conceptual model for responsible UP systems (RUPS). Our model design is based on previous personalization models (see Section 3.2) where we critically evaluated applicable and relevant components for current and future UP systems (see Section 3.3). After additionally considering current literature on personalization (Section 2.1), pervasive MR (Section 2.2 and 3.1) and responsibleness (Section 2.3), we decided on the components to include in our new model. The RUPS model updates and extends existing personalization models to capture relevant aspects of current and future personalization in physical, virtual, and hybrid environments, and functions as a framework for the analysis of *responsible* UP systems. It describes the necessary components for systems that may personalize *any* aspect of a user's perceived

[5]From a Web Architecture perspective, the resource state at an origin server may be considered such ground truth, but only *representations* of this state are transmitted to clients (hence, *Representational* State Transfer; REST) [55].

[6]https://www.w3.org/TR/wot-thing-description11/. Last accessed April 4, 2025.

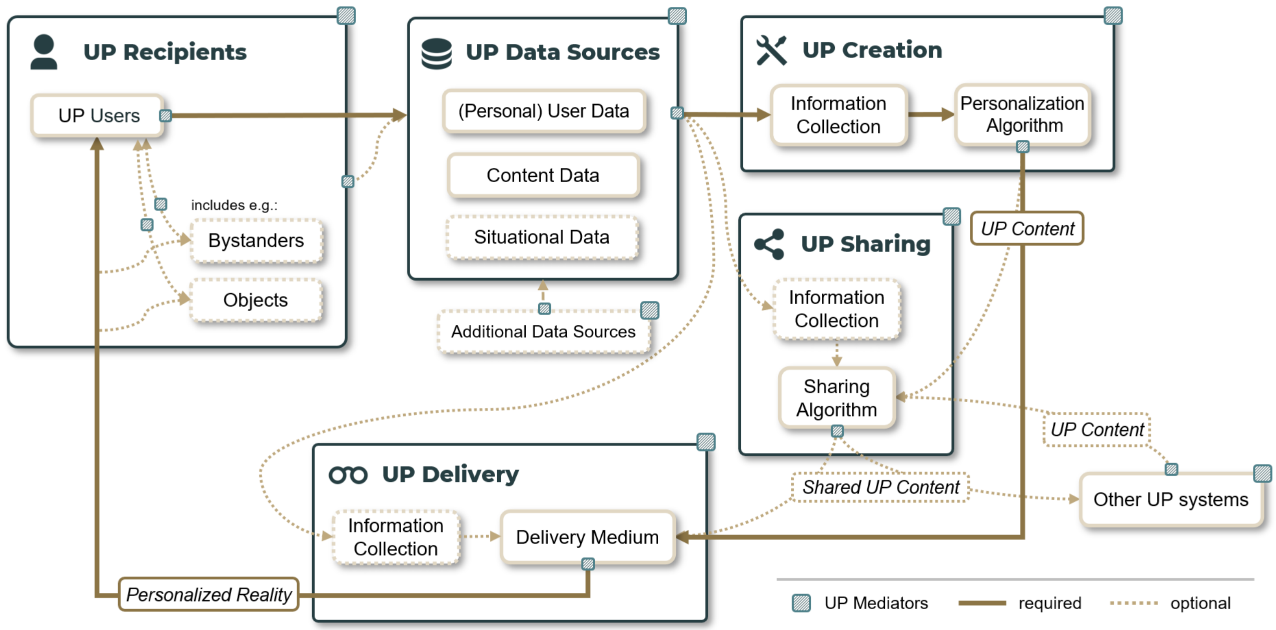

Figure 2: The RUPS model showing the components of a UP system. (Personal) User data from the *UP Users* and Content Data, along with potential Situational Data (e.g., from bystanders or the environment) and Additional Data Sources serve as an input to the *UP Creation*, and, possibly, the *UP Sharing*, and *UP Delivery* components. The *UP Creation's* (or, optionally, the *UP Sharing's*) output are then delivered to the UP User by the *UP Delivery* component who then perceive the PR. All components of a UP system may be influenced by UP Mediators. *Required* arrows between the components denote the minimum flow of data that is necessary for a minimal UP system, while *optional* ones describe the additional flow of data may be included if necessary.

reality—ranging from small and unimodal changes of a user's reality to the personalization of their full perception across modalities. To permit this significant scope extension, the RUPS model includes five main components: *UP Recipients*, *UP Data Sources*, *UP Creation*, *UP Sharing*, and *UP Delivery* (see Figure 2).

4.1 Ubiquitous Personalization Recipients

UP Users are the intended (one or multiple) recipients of PR, and may perceive PR through a single *UP Delivery Medium* or multiple *UP Delivery Media* depending on the concrete scenario. In Eirinaki and Vazirgiannis' [48] (see Figure 5) and in Bozdag's model [22] (see Figure 1), the user for whom the personalization is provided is not explicitly included in either model. This might be due to the assumption that personalization is a one-way process without feedback, as we detailed earlier (see Section 3). We argued that personalization systems should instead be modeled as closed feedback loops, which is specifically important when personalizing physical realty, and hence we include the *UP Users* separately in the RUPS model (see Figure 2). Including the users furthermore permits us to model users as being affected by *UP Mediators* (see Section 4.6), and it permits us to adopt a more user-centric viewpoint on personalization.

In addition to the *UP Users*, there might be *unintended* recipients of PR such as bystanders or objects in the environment, from whom data might be unintentionally collected as input for the *UP Creation*. Hence, we call the model's component *UP Recipients* to include all parties affected by PR. A PR in UbiComp environments is typically *implicitly* perceivable by others if personalization is applied directly to devices in the users' surroundings that are not invisibly embedded. Yet, this possibly infringes the privacy of the *UP Users*, and might, for instance, embarrass them in front of others—a phenomenon that is well-known for public displays (cf. [39]). Additionally, the personalized behavior of UbiComp devices may affect bystanders' perception of reality and thus, it should be guaranteed that the personalization of *UP Users'* physical reality does not result in negative or harmful implications for (non-PR-using) bystanders. For instance, the user of an MR HMD perceives their PR exclusively while individuals in the user's surroundings – even someone who might be shoulder-surfing – normally remain oblivious to the specific personalized content [32, 70, 133]. This creates an information asymmetry where *UP Users* might change their behavior based on the content of their PR without (perceivable) explanation for other people in their environment. This issue has been identified in pervasive MR settings [158], however, in PR settings, the information asymmetry may persist even when multiple users use the same UP system because the presented content has been personalized for each of them. While, e.g., an HMD or smart glasses could be exchanged between users, futuristic devices such as MR contact lenses may further complicate this issue [36].

To create a beneficial, transparent, and pleasurable experience for *UP Users* and other affected entities, *responsible* UP systems are transparent on what information is shared with a system [1] and on which parts of their reality are mediated or personalized [87, 159], give users agency over controlling these parts [129, 132] and over which parts of their UP content are shared with others [154, 186]. To make the interaction with the algorithmic parts of the *UP Creation* more explicit and understandable for users, creators of UP systems should follow recommendations from frameworks such as the *Algorithmic Experience* framework [9]. Additionally, responsible UP systems implement ways to moderate their impact on bystanders and the user's environment, e.g., regarding bystanders' awareness and consent [133], (perceptual) agency [160], and autonomy [132].

4.2 Ubiquitous Personalization Data Sources

According to our RUPS model, a UP system must minimally include *User Data* (from *UP Users*) that falls under the definition of personal data of the GDPR (see Section 2) and *Content Data* that is to be personalized in response to the *User Data*. UP systems might in addition include non-personal *User Data*, *Situational Data* from the users' environment and data from additional sources; yet, these are not strictly necessary. What specific and how much (personal) data a UP system needs to function, depends on the specific scenario. These aspects are not determined by the RUPS model. As depicted in our model (see Figure 2), these data sources inform the *UP Creation* and, potentially, the *UP Sharing* and *UP Delivery* components of a personalization system. Separating the data sources increases the RUPS model's modularity in contrast to existing Web personalization models ([48] and [22]) that do not differentiate between different data sources.

In our model, *User Data* refers to historic, current, and predicted data about the *UP Users*. This includes, for instance, a *UP User's* current age, preferences, interests, activities, and (predicted) goals; their medical history as well as instant physiological, cognitive and affective state; their individual skills and abilities; the level of their professional expertise; or information about their social and cultural background.

Content Data comprises data about the original (un-modified) content that is personalized during *UP Creation* and delivered in personalized form to the user via the *Delivery Medium* (see Section 4.4). We expect this data to typically overlap with *Situational Data*, since the personalized content often originates from the user's environment.

Situational Data refers to data about the environment the users are situated in, such as physical or virtual objects, or properties of the environment (e.g., temperature or brightness level). Further streams of data may be provided by interactions of users with their environment through diverse modalities with simple physical objects and devices (e.g., supermarket products or a light switch), complex systems (e.g., a robot, machine, or assembly line), as well as with artificial software agents. UP systems that include MR or UbiComp technology often use devices that can capture data from bystanders who may or may not have given consent. Preserving bystanders' privacy and agency is a challenge in MR and UbiComp environments (cf. [4, 29, 32, 41, 133, 148, 161]), and UP systems face similar challenges.

Additional Data Sources may include, for instance, in-house data that is only available to the *UP Creator* (e.g., usage profiles of other users), data from public Websites or APIs (e.g., Wikipedia, governmental open data portals, or OpenStreetMap API), or public datasets (e.g., MS COCO, or Stack Exchange Data Dump).

By differentiating between *Content*, *User*, *Situational*, and *Additional Data Sources*, the RUPS model makes the data collection more explicit and thereby enables a fine-grained mechanism for capturing users' consent to the data collection and processing, which may be implemented, e.g., by using the framework of *affirmative consent* [84]. Responsible UP systems furthermore include approaches to guarantee the security and privacy of users' personal data, e.g., by keeping the data and personalization flows transparent and actionable for *UP Users* (cf. [10, 40, 54]), while allowing them to experience the benefits of PR.

4.3 Ubiquitous Personalization Creation

UP Creation describes the creation of personalized content, and is comprised of the sub-components *Information Collection* and *Personalization Algorithm*. It corresponds to the *Information Acquisition & Searching, Content Management, Web Usage Mining, and Usage Patterns* modules in [48] (see Figure 5), and to all processes except for the *Presentation Algorithm* in Bozdag's model [22] (see Figure 1). In the process of *UP Creation*, the UP creator collects data about the users, the users' context, and from additional data sources (optional), using the means and sources that are available to them (see Figure 2). The *Personalization Algorithm* then takes this data as input, and outputs a version of the content (*UP Content*) that is personalized according to the data about the *UP Users* and their context. This output is either immediately forwarded to the *UP Delivery Medium*, so that the *UP Users* may perceive it individually, or to a *Sharing Algorithm* to provide multiple users with a shared UP experience. The possibilities of *UP Creation* are naturally constrained by the capabilities of the *UP Delivery Medium* and the available data sources.

UP Content might be *created* by *UP Users* themselves, a second party such as the manufacturer of an object in the user's physical environment or the manufacturer or owner of the *UP Delivery Medium*, or a third party, such as a provider of a specific UP application. UP creators may furthermore *collaborate* when creating PR experiences—specifically, *UP Users* might be enabled to further customize their PR (that was created by a second or third party). This could happen by modifying settings that are provided by the *UP Creator*, or through actually programming or creating virtual content that can be added to a PR. The three options are similar to how smartphone applications are created today but with more far-reaching consequences, as the creator of a UP system is in a much more powerful position: They can personalize a space that is not confined to a smartphone or computer screen and thus essentially have the possibility to determine how users experience their reality [112]. Thus, UP creators should use this power *responsibly* by, e.g., designing personalization algorithms transparently [167], or giving users (partial) control over the personalization through scrutable algorithms [89].

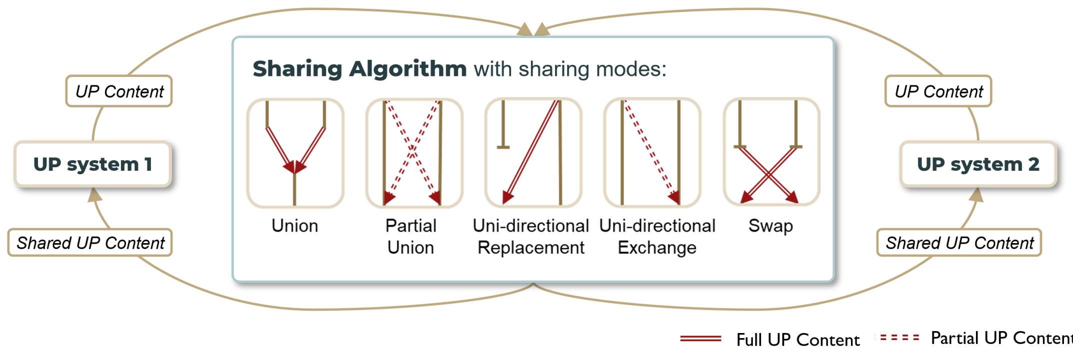

Figure 3: The Sharing Algorithm combines content from multiple UP systems based on sharing modes for multi-user *UP Sharing* (here: two users). All involved users agree on a mode before the sharing. The two systems and the included delivery media may be controlled by the same actor (e.g., the manufacturer of the delivery devices), or by multiple, interoperable actors (e.g., AR platform providers). (See the RUPS model in Figure 2 for the full details of each system.)

4.4 Ubiquitous Personalization Delivery

UP Delivery describes the component that ultimately lets *UP Recipients* experience their PR. This corresponds vaguely to the *Web Publishing* module (see Figure 5) in Eirinaki and Vazirgiannis' model [48] and to the *Presentation Algorithm* (see Figure 1) in Bozdag's model [22]. Both models assume that the delivery medium is largely standard for personalized (Web) content: on a desktop monitor or smartphone screen. However, PR features many more options for delivery media in MR and UbiComp environments (cf. [136]). Therefore, the RUPS model explicitly includes the *Delivery Medium* to account for this hardware dimension. The optional *Information Collection* may select data that is specifically necessary for delivering the PR, such as the geometry of the user's environment, location data, or background noises.

Through *UP Delivery*, various aspects of a user's reality may be personalized (e.g., through augmenting, diminishing, or mediating existing content) by addressing different sensory modalities of the user, dependent on the concrete scenario. To select a suitable *UP Delivery* medium, we propose the following requirements: Fundamentally, the medium must permit the real-time adaptation of the content it displays to the users, or even the adaptation of the interface itself, and the medium's output format must be perceivable by the *UP User*, i.e., the users' sensory abilities must match with the medium's delivery modalities. Additionally, it needs to provide a possibility for the UP creator to access and modify content (e.g., through an API). Examples of such suitable delivery media include, e.g., MR HMDs [53], robots [82], smart furniture [211], public displays [119], coffee machines [90], earphones [213], or smartwatches [25]. *Responsible* UP systems consider the social implications stemming from the delivery of personalized content with specific delivery devices, e.g., an MR HMD might be prone to creating information asymmetries (cf. [158]), while a public display might cause embarrassing situations for users (cf. [39]).

If multiple users simultaneously attempt to access their individual UP application while targeting the same physical device (e.g., a robotic arm, or a lighting system), the respective UP system's delivery medium needs to detect and resolve potential conflicts. This problem is already well-understood since it also occurs in non-personalized multi-user scenarios, and different conflict resolving algorithms (e.g., using preference integration, priority-based assignment, or time slicing) for context-aware applications have been proposed in the literature [195]. Thus, UP systems should incorporate these approaches if the delivery medium requires it.

4.5 Ubiquitous Personalization Sharing

The *Ubiquitous Personalization Sharing* (*UP Sharing*) is an optional component that aims to overcome isolated and fragmented realities, and to mitigate information asymmetries. It builds on our previous work for sharing personalized MR content across multiple users [186], and extends it to capture the sharing of any personalized content and embeds it into the RUPS model (see Figure 2). Previous personalization models do not consider shared experiences. *UP Sharing* comprises an *Information Collection* process that selects information used to decide on how to specifically share UP Content, and contains a *Sharing Algorithm* that takes a generated UP Content, modifies it, and shares it with another UP system. The perceivable PR is then displayed for each user with their respective *UP Delivery* based on the shared UP Content. *UP Sharing* may work synchronously or asynchronously, and the involved people may or may not be co-located in the same environment. It may be delivered by one or multiple delivery media by the same or different device types. We expect that *UP Sharing* will be more relevant in environments, where a person needs to *explicitly* share their PR with others so that they might perceive it as well, e.g., when the delivery medium is an MR HMD. While there are personalization systems that include input data from multiple people, e.g., group recommender systems [2, 8, 202] to create an output that is perceived by all involved users as a group, *UP Sharing* aims at sharing the outputs of such systems. This allows users to still perceive content that is personalized exclusively for them while also perceiving content that was personalized for others, depending on which sharing modes are employed.

We previously suggested five sharing modes for personalized MR that a sharing algorithm could make use of: *Union*, *Partial Union*, *Uni-directional Replacement*, *Uni-directional Exchange*, and *Swap* [186]. We adopt these for our model, as they provide a good

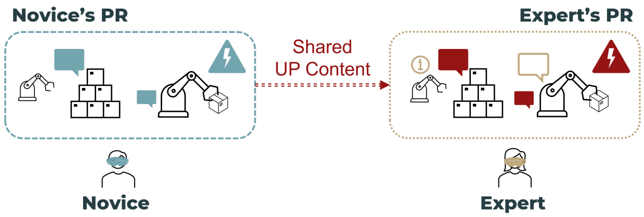

Figure 4: A *Uni-directional Exchange* of UP Content between the PRs of a Novice and an Expert robot operator who are both wearing an MR HMD. In this example, the Expert, who has already access to their own personalized content (in beige), is additionally exposed to the MR overlays (in dark red) that are initially personalized for the Novice (in blue). The shared content might help the Expert to better understand potential difficulties that the Novice might experience.

starting point to also show the potential for *UP Sharing* (see Figure 3). Depending on the specific application scenario with concrete devices, variations of these or new sharing modes might additionally be useful. For instance, two workers in a factory might prefer to share a *Union*, or *Partial Union*, of their UP Contents with each other to facilitate their collaboration. A *Uni-directional Replacement* or *Uni-directional Exchange* might be helpful for experts to support novices in learning, e.g., how to operate a certain machine (see Figure 4). Furthermore, the *Uni-directional Exchange* may also allow *UP Users* to share partial UP Content with bystanders who have suitable devices but not their own PR experience (cf. [69]). The (temporary) *swapping* of full UP Content could be beneficial in contexts like conflict resolution sessions or during group workshops that explore divisive societal issues, as it might build bridges between fragmented or isolated perceptions of reality. As democratic societies generally strive for cohesion [62], e.g., through public schools or public broadcasting, where different identities and ideologies are present in the same space, *UP Sharing* might provide a means to broaden these efforts to people's digital realities, and might even be a way to bring divided societies closer together.

When developing or analyzing UP systems, researcher and practitioners should consider whether these system induce information asymmetries or if these are prone to create isolated perceptions of reality, and investigate how *UP Sharing* might be able to mitigate such implications. Furthermore, *responsible* UP systems should provide users with an intuitive interface for administering the shared PR experiences, e.g., based on the delivery media (cf. [123]) or the type of social relationship the users have (cf. [127]), and regarding further privacy- and security-related settings (cf. [164]).

4.6 Ubiquitous Personalization Mediators

Catering to the diversity of entities and processes that are involved in the creation of a PR experience, the RUPS model considers the processes and interactions across all components to be mediated by additional factors that we refer to as *UP Mediators* (see Figure 2). This reflects Bozdag's and others' external factors such as *Government* or *Organizational Factors* (see Figure 1), but we significantly extend the coverage of this mediation to possibly influence *each* component and connection (see Section 3). Our RUPS model considers each component to be influenced, for instance, by human judgment and biases (e.g., popularity bias, authority bias, or novelty bias), company policies (e.g., business goals, or financial constraints), or third-party requests (e.g., court rulings, or advertisers). Also, technical limitations, such as the capabilities of the available sensors, or the processing power of the *Delivery Medium*, may constrain the components of a UP system further. Furthermore, cultural mediators might influence UP, such as implicitly shared personalized content that is considered embarrassing in front of others [39], or cultural differences in how willing people are to share their personal data with applications [152]. We furthermore emphasize the role of regulation on UP systems: Prominently, the *UP Data Sources* component as well as the *Information Collection* processes in the *UP Creation*, *UP Sharing*, and *UP Delivery* are constrained by legal regulation that limits what parties may collect, process, store, and share (personal) data, for what purpose, and on which legal grounds (e.g., consent, necessary performance of a contract, etc.). Examples of such regulation are the California Consumer Privacy Act (CCPA) [131], the Chinese Personal Information (PI) Security Specification [31], and the European Union's General Data Protection Regulation (GDPR) [50]. Other regulation directly affects the *Personalization Algorithm* and *Sharing Algorithm* processes as well as the *Delivery Medium* in the RUPS model. For instance, the European Union's *General Product Safety Regulation* [51]—and, with it, other consumer safety regulation around the planet—govern aspects from the manipulation of end users to deceptive practices and product hardware safety. Finally, *UP Recipients* themselves are (trivially) in-scope of such regulation, since they are citizens of a specific country, live and work in a certain region, and are therefore subject to the locally applicable regulations while using PR. Summarizing, *UP Mediators* may mediate PR on different levels and should hence be carefully considered in the development, design and deployment of UP systems.

5 Literature-based Validation and Discussion of the Model

In Section 4, we demonstrated that the RUPS model significantly expands beyond published personalization models by capturing relevant aspects of current and future personalization systems that personalize physical reality as well. In this section, we critically evaluate our proposal, and present a validation of the RUPS model, demonstrating its coverage and soundness as well as its descriptive and predictive power with respect to personalization systems.

Selection of Personalization Systems. To compile a list of diverse personalization systems, we searched the ACM Digital Library using the following search terms to obtain a preliminary list of papers: `[[Title: personalization] OR [Title: personalized]] AND [[Abstract: personalization] OR [Abstract: personalized]] AND [CCS 2012: Human-centered computing]`. We then manually surveyed the abstracts of the resulting list of 735 articles with respect to whether they adhered to the following criteria: The paper needs to propose or describe a system or prototype that includes personalization by using personal data in the sense we describe in Section 2. We furthermore excluded papers that report only on an evaluation of a system or prototype. In the next step (N=120), we made sure that the selected systems are sufficiently different

Table 1: Mapping of selected publications to the components of the UP model (Part 1 of 2; Up Rec. = UP Recipients). The *UP Data Sources* (see Sect. 4.2) also include the device(s) that provide the data, and *User* and *Situational Data* differentiate between real-time and static (historic) data.

Source	UP Rec.	UP Data Sources					
		User Data			*Situational Data*		
		Device(s)	*Real-Time User Data*	*Static User Data*	*Device(s)*	*Real-Time Situational Data*	*Static Situational Data*
[37]	-	-	Mood, Situation (e.g., work or freetime), Physical activity level	Preferences, Rating of past activities and tips	-	Weather, time	-
[53]	-	Smartwatch	-	Photos/videos taken by the user, skiing tracks	-	Current location	-
[71]	-	WiFi router	Workout activity detection	Past workout activity	-	-	-
[72]	-	Sensors, Ergometer	Biological signals, Physiological signals (e.g. speed on ergometer)	Historical gameplay data	-	-	-
[86]	-	Camera, Microphone	User identification, Mood, emotion recognition and engagement analysis	Past interactions, Order history, Preferences	-	-	-
[100]	-	Fitness Tracker, Online calender	Current heart rate logs, bed time, wake-up time, step count	Historic heart rate logs, bed time, wake-up time, step count, Calendar events (past and future)	-	-	-
[101]	-	-	-	Historical bicycling route logs	-	-	-
[102]	-	MiBand3 smart bracelet	Mood, Biological signals and activities	Music preferences	MiBand3 smart bracelet	Location, weather, time	-
[122]	-	Vibrotactile breathing pacer	Current respiratory rate, BPM.	Past respiratory rate	-	-	-
[137]	-	-	User input (text)	Profile (e.g. country and language/locale), viewing history	-	Target title features, User-item affinity	Target title features (e.g., metadata)
[142]	-	Vehicle mounted smart devices, smartphone	Current driving behavior	Historic driving behavior at home city	-	Location	-
[155]	-	Learning application on a tablet	Student's performance	-	-	-	-
[168]	-	IDE	Interactions with the IDE	-	-	-	-
[205]	-	Thermal camera, RGB camera, Web/mobile interface	Facial temperature, Face identification	Thermal comfort preferences	Environmental sensors	Temperature, Humidity	-
[211]	-	Microsoft Kinect v2	Human mesh reconstruction	-	-	-	-
[212]	-	Questionnaire, Smartphone	-	Name, Gender, Interest areas, Personality traits	Visual Object Detection	Location, Points of Interest (POI)	-

Table 2: Mapping of selected publications to the components of the UP model (Part 2 of 2; IC = Information Collection)

...	UP Data Sources		UP Creation		UP Sharing		UP Delivery		UP Mediators
	Content Data	*Additional Data Sources*	*IC*	*Personalization Algorithm*	*IC*	*Sharing Algorithm*	*IC*	*Delivery Medium*	
[37]	Smartphone application interface	Compendium of Physical Activities	-	Physical activity and tip recommendation	-	-	-	Smartphone	-
[53]	Physical skiing map	-	-	AR overlay on physical map	-	Followers see primary user's AR content overlayed on a physical map	-	Head-mounted mobile phone	-
[71]	Workout recommendations	-	-	Personalized Workout Interpretation and Smart Workout Assessment	-	-	-	-	-
[72]	Virtual game content	User/player models, training plans	-	Adaption of the gameplay	-	-	-	Computer screen	-
[86]	Drinks, News	Other users' preferences, OpenPose library, Twitter API	-	Drink recommendation system, News recommendation	-	-	-	Robot	-
[100]	Sleep schedules	-	-	User- and time-specifc sleep schedules	-	-	-	Online calendar and web-app	-
[101]	Surface of a leather bag	-	-	Patterns of personal bicycle logs to be printed on a leather bag	-	-	-	Leather Bag	-
[102]	Songs	MillionSong Datase, Spotify, ComParE2013 acoustic feature set	-	Multi-task Ubiquitous Music Recommendation Model	-	-	-	-	-
[122]	Pace, frequency, and amplitude of vibrations of the breathing pacer	-	-	personalized patterns of pace, frequency, and amplitude of vibrations	-	-	-	Vibrotactile breathing pacer	-
[137]	Target titles	-	-	Personalized search result	-	-	-	-	-
[142]	Driving guidelines	Driving environment for the home city and the visiting city	-	Driving safety recommendations while driving in a new city	-	-	-	-	-
[155]	Education content of an e-learning application	-	-	Non-task break activity timing	-	-	-	Robot, Tablet	-
[168]	IDE Interface	High-level task labels	-	Task recommendations for interactions with the IDE	-	-	-	IDE	-
[205]	-	-	-	Thermal comfort estimation model	-	-	-	-	AHSRAE Standard 55 – Thermal Environmental Conditions for Human Occupancy
[211]	Furniture (chair, desk, keyboard, monitor)	-	-	Ergonomic furniture configurations	-	-	-	Personal workspace (chair, desktop monitor, keyboard, and desk)	Ergonomics standards and guidelines
[212]	Information about Points-of-Interest (POI) in AR	-	-	Personalized top facts about POI	-	-	-	Smartphone (AR)	-

from each other to guarantee a heterogeneous selection. We further strove towards including papers from a broad range of venues (when possible) to further increase the diversity of systems. Finally, the remaining papers (N=55) were screened to select those that provided a sufficiently deep description to permit analysis which yielded a selection of 13 papers suitable for our validation ([37, 71, 72, 100–102, 122, 142, 155, 168, 205, 211, 212]).

Three papers were added to this survey even though they were not returned in response to the original query: We included an example of personalization that is used by millions of users, i.e., search personalization at Netflix [137], another one because it is one of the very few works that include the sharing of personalized content [53], and we added one as a more detailed description of a system that was found in the query [86]. The resulting list of 16 personalization systems is not meant to exhaustively describe the spectrum of systems the RUPS model covers, but rather as an exemplary range of different systems that our model is *able to cover*.

During our screening, we generally observed that most of these works and others we encountered usually do not define *personalization* (e.g., [102, 211]), or name a system as personalized while it is simply adaptive based on non-personal parameters. The term is also often used in a context where *customization* would be more appropriate (see Section 2), e.g., because the adaptation is user- and not system-initiated such as Tao et al.'s adaptable food shapes [191]. Furthermore, we found surprisingly few personalized systems or prototypes for AR or VR.

Methodology of the Validation. We analyzed all articles on the final list [37, 53, 71, 72, 86, 100–102, 122, 137, 142, 155, 168, 205, 211, 212] in detail and mapped each presented system to the RUPS model (see Table 1 and 2). On an abstract level, and for all articles, the mapping of each system to our model was straightforward and unanimous between the co-authors. Towards demonstrating coverage, we show that we can map the components of each published system; towards demonstrating soundness, we show that each mapping can be done naturally and does not require bending or reinterpreting components of the system nor of the RUPS model. Finally, we evaluate the RUPS model's *explanatory* and *predictive* power: *Full coverage* of our model implies that each aspect of a published personalization system can be mapped to a component of our model. However, we expect that parts of the RUPS model have not been considered in each system, leaving mapping gaps. Studying these mapping gaps permits us to evaluate whether our model generates *new insights* about existing personalization systems, illustrating its ability to explain existing systems and to predict future developments on top of the state of the art.

We also attempted to map a selection of systems to the two personalization models discussed in Section 3, however this showed that these models do not adequately cover all important aspects of the surveyed personalization systems (see Appendix B). For instance, the Presentation Algorithm in Bozdag's model [22] does not include the delivery devices and therefore, e.g., the smart furniture in [211] cannot be mapped here. Also, Lee et al.'s work [101] cannot be fully mapped to Eirinaki and Vazirgiannis' model, since the *Usage Logs* here come from a person's bicycle routes while the to-be-personalized content (i.e. the surface of a leather bag) is from a different environment. The model however assumes that the usage logs originate from the same environment that is personalized. Also, Bozdag's model conflates different information sources for the personalization, e.g., in Yang et al.'s work [212] it is unclear from the mapping, which information is about the user, situation, or content and which is from additional sources. In the following, we discuss details of all mappings we performed to the RUPS model and how it more adequately covers current personalization systems than existing models.

5.1 UP Recipients and UP Data Sources

Each surveyed system has (typically human) users, but specific user characteristics of a system's expected users or the expected context of use are usually not discussed. This confirms the notion that personalization research and practice often has a technology-centered focus that does not pay sufficient attention to the actual users and their context [97]. Yet, the RUPS model is able to capture the systems' users in the component *UP Recipients*, as well as potential bystanders and objects in the environment.

Concerning the *UP Data Sources*, we differentiate between static (e.g., user profiles) and real-time data in our analysis to accommodate for the increased availability and diversity of data sources (e.g., environmental or physiological sensors) in MR and UbiComp environments. We found that most of the surveyed systems use at least one type of real-time user and situational data, including activity recognition, physiological or environmental data, detected objects, or interactions with an interface. While most systems involve static user data (e.g., demographics or preferences), almost none of the surveyed systems include static situational data—the exception is [137]. Furthermore, only few systems explicitly list *Additional Data Sources* (e.g., from a public dataset [86] or a third-party [102]). Our mapping furthermore supports the identification of potential additional data sources which be could be useful to enhance the system's personalization For instance, our mapping shows that Lee et al.'s system [100] could benefit from adding the calendars of other people, who, e.g., share a household with the primary user, to better align the personalized sleeping schedules with their social context; while Yang et al.'s system [212] could profit from including public, real-time data from public tourism offices with their personalized AR tourist guide.

Furthermore, the RUPS model emphasizes the closed-loop-nature of personalization where the displayed content continuously influences the iterative data collection of personalization. If a user of Schmidmaier et al.'s adaptive IDE [168], e.g., follows the personalized task suggestions, their future interactions with the IDE would be affected and thus, a feedback effect occurs since the input for the personalization are the user's interaction with the IDE. However, such feedback effects cannot be accurately mapped using current personalization models due to the linear modeling of personalization processes, demonstrating our proposal's higher explanatory power.

5.2 UP Creation and UP Delivery

Regarding *UP Creation*, only some of the surveyed papers discuss potential implications stemming from the use of their system, such as Pargal et al. who investigated the effect of their personalized recommendations while driving on driving safety [142]. However, the surveyed papers do not consider potential implications of their systems outside of their systems' core domain, e.g., on how individuals perceive reality, or on how these may effect society. Additionally, we noted that the surveyed articles usually conflate *UP Creation* and *UP Delivery* as these are typically done by the same party. However, the RUPS model is capable of including systems where these components are separated, for instance, for privacy reasons (e.g., if the *Delivery Medium* has access to information locally that the *UP Creation* may not make use of), and can at the same time be used to verify the modularity of a proposed personalization system (e.g., towards enabling the reuse of individual components in other systems). Furthermore, our survey shows that current personalization systems also tend to conflate the *Information Collection* and the *Personalization Algorithm*. The RUPS model highlights the separation of concerns between information collection and personalization, which is particularly relevant when different information types or sources are used in creating, sharing and delivering personalized content. This enables a modular UP system design, where these three parts of a UP system are executed by different parties that each may have their own information collection with access to different data sources.

Through the *UP Delivery* component, the RUPS model is capable of describing heterogeneous *Delivery Media* we found in our survey (e.g., breathing pacer [122], leather bag [101], or robot and tablet [155]), and also permits a comparison across these different systems. Specifically, this shows that the model is effective to not only capture virtual delivery media, but it also intuitively applies to systems that personalize physical reality. The RUPS model facilitates exploring the design space of *UP Delivery* that stretches beyond the possibilities discussed in the surveyed papers, especially when the *Delivery Medium* is unspecified (e.g., [71, 102, 137, 142, 205]). Concretely, it shows an interesting cascade of insights, where a dedicated discussion of the *Delivery Medium* leads to insights regarding inclusion of other information sources, *UP Sharing*, and *UP Mediators*. Li et al.'s system [102], for instance, does not discuss the device with which the users consume the recommended music. If private delivery media such as headphones are used to realize this personalized music recommendation system, bystanders remain oblivious to the music, while the usage of loudspeakers would affect them directly. In this case, *UP Sharing* might be helpful to ensure that the recommended music matches the taste of all people who share this environment. Loudspeakers as *Delivery Medium* also point towards additional mediators, since there might be, e.g., cultural constraints that define which types of music or volume levels are socially acceptable in a given environment.

5.3 UP Sharing

While some of the surveyed papers mention the *implicit* sharing of personalized content, only Fedosov et al. [53] discuss the *explicit* sharing of personalized content among multiple users. Yet, the authors do not explore this aspect further towards studying individual and societal implications of the shared personalized content. Also, most of the surveyed papers do not take multi-user scenarios into account, they seem to treat personalization as a phenomenon that affects only individual users without considering the wider implications (cf. [97]). The *UP Sharing* component of the RUPS model highlights the broader impact of personalization that is essential for providing multiple users and bystanders with a common understanding of physical reality, e.g., through providing bridges between fragmented perceptions of reality. With respect to the predictive power of the RUPS model, mapping previous work to our model unveils not only issues of (non-shared) personalized content, but also shows possibilities for managing the sharing of personalized content in a way that would benefit the users of the systems we studied. More specifically, the RUPS model is applicable—but has little explanatory and predictive power—when applied to personalization systems that already implicitly share personalized content (e.g., [86, 211]), since there is little information asymmetry in this case. However, when applied to those surveyed systems where personalized content is not shared implicitly (i.e., [100, 102, 122, 137, 212]), mapping to the RUPS model demonstrates that these systems are prone to information asymmetry between the *UP Users* and bystanders, and that this asymmetry could (and possibly should) be managed to benefit all involved. We hence argue that the provision of a distinct component that focuses on the sharing of UP Content represents a valuable advancement of our model with respect to other models of personalization in the literature.

5.4 UP Mediators

Our RUPS model shows all components of a UP system as being influenced by system-external mediators, however, such *UP Mediators* are not explicitly discussed in the majority of surveyed articles (except for [205, 211]). Yet, applying the RUPS model highlights that such mediators are present across the different system components: The personalized search results at Netflix [137], for instance, might be mediated by Netflix's company policies (e.g., ranking their in-house productions higher) while the personalized breathing pacer [122] might be mediated by medical guidelines (e.g., the maximum recommended usage duration of a breathing pacer). By making the mediation factors of personalization more explicit than in the past models, the RUPS model contributes to a more systemic and societally integrated understanding of personalization systems.

6 Using the Model for the Structured Analysis of UP Systems

Through the validation of the RUPS model, we showed that it can fully describe past and current personalization systems and is able to highlight design gaps and opportunities for relevant extension possibilities. We propose that our model will not only become relevant once our *physical* realities are being significantly personalized as well, but that it already today provides a valuable tool to evaluate current and emerging UP systems. The RUPS model hence enables researchers and practitioners to consider the broader implications of the personalization systems they conceptualize, implement, and deploy (e.g., with respect to the loss of shared worlds [24]).

We thus propose an exemplary structure that may be used as a starting point for the structured evaluation of such a system: First, the parts of an existing personalization system are mapped to the components of the RUPS model. This shows, whether the considered system should indeed be categorized as a UP system if it uses *personal data* for system-initiated adaptations, and contains the *required* components of our model: *UP Recipients*, *UP Data Sources*, *UP Creation*, and *UP Delivery*. In the next step, each individual component is inspected based on the descriptions that we present in Section 4. What are the ethical/societal/technological implications of using a certain component? Who is affected how by the specific design choices in each component? Next, the connections and interface between the components are analyzed. The conceptual model (see Figure 2) can serve as a high-level overview here. This helps to address questions such as: How does the data and content flow between the components? Which data is needed where and how is it acquired based on the involved people's consent? Following, the components and their connections are analyzed with regards to possible UP Mediators. What regulations, biases, and socio-cultural assumptions apply not just to the handling of (personal) user data, but also to the other components and their connections? Finally, the *overall* system and *each individual component* can be analyzed regarding extension possibilities. For instance, additional delivery media or further data sources might be included or exchanged to extend the systems functionalities, e.g., if the analysis before showed problematic implications stemming from a certain component.

We contend that an analysis structured in this way is valuable for different groups of stakeholders: A researcher might use the RUPS model to study the implications of an UP system and its components on individuals and society. A designer of an UP system may use our model to check which types of mediation (negatively or beneficially) affect what part of a proposed personalization system. The system's *UP Data Sources* and the *Personalization Algorithm* itself might be evaluated with respect to biases that these components may exhibit, and which regulations apply to the *UP Delivery* and the *UP Recipients* within the intended usage context. For a regulator, the RUPS model thereby becomes a policy advice tool: Which components of a system should be regulated to achieve a specific policy goal? A company may use the RUPS model to verify whether its personalization solutions are in-line with company values, e.g., when the company announces to mitigate biases.

6.1 Future Work

Currently, the RUPS model provides these stakeholders a starting point and overview for engaging in a meaningful discourse with people affected by UP systems, e.g., through workshops or focus groups. To further support such a discourse on and engagement with UP systems, we plan to create materials with more detailed hands-on guidelines such as Design Cards [109] in the future, and evaluate them with relevant stakeholders. Such materials would, e.g., enable designers to use the RUPS model within the different stages of a personalization system's design process.

Besides using the RUPS model as a basis for further materials, also the model itself and its validation could be further studied. Cross-domain applications of the model could be helpful to investigate whether it is uniformly applicable in different domains (e.g., health, industry, entertainment) and to compare personalization systems across these domains systematically. Additionally, the model could be further assessed based on systems that combine personalization and customization (cf. *mixed-initiative* approaches [26]), or that include human-AI collaboration (cf. [34]). Further, the model may even be applied to systems that personalize content towards non-human animals (cf.[150]) or autonomous software agents (cf. [199]). While the model currently considers some multi-user scenarios through its *UP Recipients* and *UP Sharing* components, it would be interesting to investigate in more detail how well the model applies to scenarios where multiple users have their own PR which is implicitly (e.g., through the same delivery media) shared among them. Also, the individual and societal implications of multi-user PR scenarios could be investigated in more detail. Additionally, the RUPS model could be applied to more systems that make use of emerging technologies such as AR/MR/VR scenarios to further investigate the model's future viability. Furthermore, since *UP Mediators* are influential for each model component, future research could explore how these can be identified and made more tangible.

7 Conclusion

In this work, we presented the RUPS model that captures and contextualizes the components of UP systems in physical, hybrid, and virtual environments. Our model responds to the need for a new personalization model created by the vast expansion of personalization through MR and UbiComp technologies, as existing personalization models are not fit anymore to describe current UP systems. Through a literature-based validation, we demonstrated that the RUPS model covers diverse personalization systems, can be used to analyze these systems, and is effective in highlighting design gaps and opportunities, thereby naturally proposing relevant extension possibilities. Furthermore, the RUPS model answers the need for a systematic and holistic description of UP systems, and we showed that it induces a structured analysis process of UP systems. A *responsible* implementation of UP systems therefore ensures that PR is indeed beneficial for individuals and society, as it mitigates potential harmful implications. Such benevolent PR has the potential to make people's lives more inclusive, convenient, and efficient, while at the same time strengthening society by focusing on nourishing shared worlds and common experiences.

Finally, we argue that the RUPS model is necessary and timely. While we are today at a point where personalization is increasingly interwoven with people's realities but only visions of all-encompassing PR exist (e.g., [107, 110, 115]), our model provides a common way to enable a structured discourse of emerging UP systems. This is necessary since known harms of personalization and pervasive MR are prone to be amplified if combined in PR, as we discussed in Sections 2 and 3. We thus urge researchers and practitioners to study UP and its possible implications, including potential harms as well as of practical strategies for ensuring the responsible and ethical implementation of PR. As the technologies to create PR through MR and UbiComp technologies are increasingly available and capable, this topic should receive considerable attention over the next decade.

Acknowledgments

This work was partially funded by the Swiss Innovation Agency Innosuisse (#48342.1 IP-ICT) and the Basic Research Fund of the University of St.Gallen.

References

[1] Melvin Abraham, Pejman Saeghe, Mark McGill, and Mohamed Khamis. 2022. Implications of XR on Privacy, Security and Behaviour: Insights from Experts. In *Nordic Human-Computer Interaction Conference (NordiCHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3546155.3546691

[2] Sara Abri, Rayan Abri, and Salih Çetin. 2020. Group-based personalization using topical user profile. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 181–186. https://doi.org/10.1145/3386392.3399559

[3] Alessandro Acquisti, Laura Brandimarte, and Jeff Hancock. 2022. How privacy's past may shape its future. *Science* 375, 6578 (Jan. 2022), 270–272. https://doi.org/10.1126/science.abj0826

[4] Imtiaz Ahmad, Rosta Farzan, Apu Kapadia, and Adam J. Lee. 2020. Tangible privacy: Towards user-centric sensor designs for bystander privacy. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2 (Oct. 2020), 1–28. https://doi.org/10.1145/3415187

[5] Sun Joo (Grace) Ahn, Jooyoung Kim, and Jaemin Kim. 2023. The future of advertising research in virtual, augmented, and extended realities. *International Journal of Advertising* 42, 1 (Jan. 2023), 162–170. https://doi.org/10.1080/02650487.2022.2137316

[6] Saifeddin Alimamy and Juergen Gnoth. 2022. I want it my way! The effect of perceptions of personalization through augmented reality and online shopping on customer intentions to co-create value. *Computers in Human Behavior* 128 (March 2022), 107105. https://doi.org/10.1016/j.chb.2021.107105

[7] Oscar Alvarado, Hendrik Heuer, Vero Vanden Abeele, Andreas Breiter, and Katrien Verbert. 2020. Middle-Aged Video Consumers' Beliefs About Algorithmic Recommendations on YouTube. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (Oct. 2020), 121:1–121:24. https://doi.org/10.1145/3415192

[8] Oscar Alvarado, Nyi Nyi Htun, Yucheng Jin, and Katrien Verbert. 2022. A systematic review of interaction design strategies for group recommendation systems. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (Nov. 2022), 1.52. https://doi.org/10.1145/3555161

[9] Oscar Alvarado and Annika Waern. 2018. Towards Algorithmic Experience: Initial Efforts for Social Media Contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. https://doi.org/10.1145/3173574.3173860

[10] Sumit Asthana and Jane Im. 2024. "I know even if you don't tell me": Understanding Users' Privacy Preferences Regarding AI-based Inferences of Sensitive Information for Personalization. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 21.

[11] Ronald T. Azuma. 2019. The road to ubiquitous consumer augmented reality systems. *Human Behavior and Emerging Technologies* 1, 1 (2019), 26–32. https://doi.org/10.1002/hbe2.113

[12] Laura Bajorunaite, Stephen Brewster, and Julie R. Williamson. 2022. "Reality anchors": Bringing cues from reality into VR on public transport to alleviate safety and comfort concerns. In *Extended abstracts of the 2022 CHI conference on human factors in computing systems (Chi ea '22)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3491101.3519696

[13] Michael Barz, Sebastian Kapp, Jochen Kuhn, and Daniel Sonntag. 2021. Automatic Recognition and Augmentation of Attended Objects in Real-time using Eye Tracking and a Head-mounted Display. In *ACM Symposium on Eye Tracking Research and Applications*. ACM, Virtual Event Germany, 1–4. https://doi.org/10.1145/3450341.3458766

[14] Morten Bay. 2023. Arendt in the Metaverse: Four properties of eXtended Reality that imperil factual truth and democracy. *Convergence* 29, 6 (Dec. 2023), 1698–1712. https://doi.org/10.1177/13548565231199957

[15] Kenan Bektaş. 2020. Toward A Pervasive Gaze-Contingent Assistance System: Attention and Context-Awareness in Augmented Reality. In *ACM Symposium on Eye Tracking Research and Applications (ETRA '20 Adjunct)*. ACM, New York, NY, USA, 1–3. https://doi.org/10.1145/3379157.3391657

[16] Kenan Bektaş, Jannis Strecker, Simon Mayer, and Kimberly Garcia. 2024. Gaze-enabled activity recognition for augmented reality feedback. *Computers & Graphics* 119 (March 2024), 103909. https://doi.org/10.1016/j.cag.2024.103909

[17] Kenan Bektaş, Jannis Strecker, Simon Mayer, Kimberly Garcia, Jonas Hermann, Kay Erik Jenss, Yasmine Sheila Antille, and Marc Elias Solèr. 2023. GEAR: Gaze-enabled augmented reality for human activity recognition. In *ACM Symposium on Eye Tracking Research and Applications (ETRA '23)*. ACM, New York, NY, USA, 9. https://doi.org/10.1145/3588015.3588402

[18] Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. 2001. MagicBook: Transitioning between reality and virtuality. In *CHI '01 extended abstracts on human factors in computing systems (Chi ea '01)*. Association for Computing Machinery, New York, NY, USA, 25–26. https://doi.org/10.1145/634067.634087

[19] Grant Blank and Bianca C. Reisdorf. 2012. The participatory web. *Information, Communication & Society* 15, 4 (2012), 537–554. https://doi.org/10.1080/1369118X.2012.665935

[20] Camille M Blondin, Ekaterina Ivanova, Jonathan Eden, and Etienne Burdet. 2021. Perception and performance of electrical stimulation for proprioception. In *2021 43rd annual international conference of the IEEE engineering in medicine & biology society (EMBC)*. IEEE, Mexico, 4550–4554. https://doi.org/10.1109/EMBC46164.2021.9630186

[21] Balázs Bodó, Natali Helberger, Sarah Eskens, and Judith Möller. 2019. Interested in Diversity. *Digital Journalism* 7, 2 (Feb. 2019), 206–229. https://doi.org/10.1080/21670811.2018.1521292

[22] Engin Bozdag. 2013. Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15, 3 (Sept. 2013), 209–227. https://doi.org/10.1007/s10676-013-9321-6

[23] Léa Briand, Guillaume Salha-Galvan, Walid Bendada, Mathieu Morlon, and Viet-Anh Tran. 2021. A Semi-Personalized System for User Cold Start Recommendation on Music Streaming Apps. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD '21)*. Association for Computing Machinery, New York, NY, USA, 2601–2609. https://doi.org/10.1145/3447548.3467110

[24] Maria Brincker. 2021. Disoriented and Alone in the "Experience Machine" – On Netflix, Shared World Deceptions and the Consequences of Deepening Algorithmic Personalization. *SATS* 22, 1 (July 2021), 75–96. https://doi.org/10.1515/sats-2021-0005

[25] Michael Bub, Evgeny Volynsky, and Wolfgang Wörndl. 2018. User acceptance of proactive recommendations on smartphone and smartwatch. In *Adjunct publication of the 26th conference on user modeling, adaptation and personalization (Umap '18)*. Association for Computing Machinery, New York, NY, USA, 143–150. https://doi.org/10.1145/3213586.3226210

[26] Andrea Bunt, Cristina Conati, and Joanna McGrenere. 2007. Supporting interface customization using a mixed-initiative approach. In *Proceedings of the 12th international conference on Intelligent user interfaces (IUI '07)*. Association for Computing Machinery, New York, NY, USA, 92–101. https://doi.org/10.1145/1216295.1216317

[27] L. Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth Vishnoi. 2019. Controlling Polarization in Personalization: An Algorithmic Framework. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, Atlanta, GA, USA, 160–169. https://doi.org/10.1145/3287560.3287601

[28] David J. Chalmers. 2023. *Reality+: Virtual Worlds and the Problems of Philosophy*. Penguin Books Ltd, S.l.

[29] Bingqing Chen, Yue Li, Amber Hu, and Yilan Elan Tao. 2024. Awkward or Acceptable? Understanding the Bystander Perspective on the Ubiquity of Cross Reality in Ambiguous Social Situations. In *2nd Joint Workshop on Cross Reality, co-located with ISMAR 2024*. IEEE, Seattle, Washington, United States, 5. https://doi.org/10.1109/ISMAR-Adjunct64951.2024.00041

[30] Lu Cheng, Kush R. Varshney, and Huan Liu. 2021. Socially Responsible AI Algorithms: Issues, Purposes, and Challenges. *J. Artif. Int. Res.* 71 (Sept. 2021), 1137–1181. https://doi.org/10.1613/jair.1.12814

[31] Standardization Administration of China. 2020. Information security technology - Personal information (PI) security-specification (English Translation). , 42 pages. https://www.tc260.org.cn/upload/2020-09-18/1600432872689070371.pdf

[32] Ji Won Chung, Xiyu Jenny Fu, Zachary Deocadiz-Smith, Malte F Jung, and Jeff Huang. 2023. Negotiating Dyadic Interactions through the Lens of Augmented Reality Glasses. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. ACM, Pittsburgh PA USA, 493–508. https://doi.org/10.1145/3563657.3595967

[33] Lena Cohen, Heila Precel, Harold Triedman, and Kathi Fisler. 2021. A New Model for Weaving Responsible Computing Into Courses Across the CS Curriculum. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, New York, NY, USA, 858–864. https://doi.org/10.1145/3408877.3432456

[34] Cristina Conati. 2024. AI-Driven Personalization to Support Human-AI Collaboration. In *Companion Proceedings of the 16th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '24 Companion)*. Association for Computing Machinery, New York, NY, USA, 5–6. https://doi.org/10.1145/3660515.3661324

[35] Robbe Cools and Adalberto Simeone. 2022. Mobile displays for cross-reality interactions between virtual and physical realities. In *Proceedings of the 20th international conference on mobile and ubiquitous multimedia (Mum '21)*. Association for Computing Machinery, New York, NY, USA, 217–219. https://doi.org/10.1145/3490632.3497838

[36] Robbe Cools, Renée Venema, Augusto Esteves, and Adalberto L. Simeone. 2024. The Impact of Near-Future Mixed Reality Contact Lenses on Users' Lives via an Immersive Speculative Enactment and Focus Groups. In *ACM Symposium on Spatial User Interaction*. ACM, Trier Germany, 1–13. https://doi.org/10.1145/3677386.3682099

[37] Ine Coppens, Luc Martens, and Toon De Pessemier. 2023. Motivating people to move more with personalized activity and tip recommendations: a randomized controlled trial. In *Companion proceedings of the 28th international conference on intelligent user interfaces (IUI '23 companion)*. Association for Computing

Machinery, New York, NY, USA, 123–126. https://doi.org/10.1145/3581754.3584149

[38] Siddhartha Datta. 2022. Cross-Reality Re-Rendering: Manipulating between Digital and Physical Realities. http://arxiv.org/abs/2211.08005

[39] Nigel Davies, Marc Langheinrich, Sarah Clinch, Ivan Elhart, Adrian Friday, Thomas Kubitza, and Bholanathsingh Surajbali. 2014. Personalisation and privacy in future pervasive display networks. In *Proceedings of the SIGCHI conference on human factors in computing systems (Chi '14)*. Association for Computing Machinery, New York, NY, USA, 2357–2366. https://doi.org/10.1145/2556288.2557287

[40] Jaybie A. De Guzman, Kanchana Thilakarathna, and Aruna Seneviratne. 2020. Security and Privacy Approaches in Mixed Reality: A Literature Survey. *Comput. Surveys* 52, 6 (Nov. 2020), 1–37. https://doi.org/10.1145/3359626

[41] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI conference on human factors in computing systems (Chi '14)*. Association for Computing Machinery, New York, NY, USA, 2377–2386. https://doi.org/10.1145/2556288.2557352

[42] Anind K. Dey. 2010. Context-Aware Computing. In *Ubiquitous computing fundamentals*, John Krumm (Ed.). Chapman & Hall/CRC Press, Boca Ragon, FL, USA, 321–352.

[43] Anind K. Dey and Gregory D. Abowd. 1999. *Towards a Better Understanding of Context and Context-Awareness*. Technical Report. Georgia Institute of Technology. 12 pages. https://smartech.gatech.edu/handle/1853/3389

[44] Directorate-General for Research and Innovation and European Commission. 2013. *Options for strengthening responsible research and innovation - Report of the Expert Group on the State of Art in Europe on Responsible Research and Innovation*. Publications Office of the European Union, Luxembourg. https://data.europa.eu/doi/10.2777/46253

[45] Gilad Edelman. 2023. Why Don't We Just Ban Targeted Advertising? https://www.wired.com/story/why-dont-we-just-ban-targeted-advertising/

[46] Niklas Eder. 2021. Beyond Automation: Machine Learning-Based Systems and Human Behavior in the Personalization Economy. https://papers.ssrn.com/abstract=3979840

[47] Chloe Eghtebas, Gudrun Klinker, Susanne Boll, and Marion Koelle. 2023. Co-Speculating on Dark Scenarios and Unintended Consequences of a Ubiquitous(ly) Augmented Reality. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. ACM, Pittsburgh PA USA, 2392–2407. https://doi.org/10.1145/3563657.3596073

[48] Magdalini Eirinaki and Michalis Vazirgiannis. 2003. Web mining for web personalization. *ACM Transactions on Internet Technology* 3, 1 (Feb. 2003), 1–27. https://doi.org/10.1145/643477.643478

[49] Isamu Endo, Kazuki Takashima, Maakito Inoue, Kazuyuki Fujita, Kiyoshi Kiyokawa, and Yoshifumi Kitamura. 2021. A reconfigurable mobile head-mounted display supporting real world interactions. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems (Chi ea '21)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3411763.3451765

[50] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* 59, L119 (April 2016), 1–88. http://data.europa.eu/eli/reg/2016/679/oj/eng

[51] European Parliament and Council of the European Union. 2023. Regulation (EU) 2023/988 of the European Parliament and of the Council of 10 May 2023 on general product safety, amending Regulation (EU) No 1025/2012 of the European Parliament and of the Council and Directive (EU) 2020/1828 of the European Parliament and the Council, and repealing Directive 2001/95/EC of the European Parliament and of the Council and Council Directive 87/357/EEC. *Official Journal of the European Union* 66, L 135 (May 2023), 51. http://data.europa.eu/eli/reg/2023/988/oj/eng

[52] Haiyan Fan and Marshall Scott Poole. 2006. What Is Personalization? Perspectives on the Design and Implementation of Personalization in Information Systems. *Journal of Organizational Computing and Electronic Commerce* 16, 3 (2006), 179–202. https://doi.org/10.1207/s15327744joce1603&4_2

[53] Anton Fedosov, Ivan Elhart, Evangelos Niforatos, Alexander North, and Marc Langheinrich. 2016. SkiAR: Wearable Augmented Reality System for Sharing Personalized Content on Ski Resort Maps. In *Proceedings of the 7th Augmented Human International Conference 2016 (AH '16)*. ACM, New York, NY, USA, 1–2. https://doi.org/10.1145/2875194.2875234

[54] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. https://doi.org/10.1145/3411764.3445148

[55] Roy T. Fielding and Richard N. Taylor. 2002. Principled design of the modern web architecture. *ACM Trans. Internet Technol.* 2, 2 (May 2002), 115–150. https://doi.org/10.1145/514183.514185

[56] Markus Funk, Andreas Bächler, Liane Bächler, Thomas Kosch, Thomas Heidenreich, and Albrecht Schmidt. 2017. Working with augmented reality? A long-term analysis of in-situ instructions at the assembly workplace. In *Proceedings of the 10th International Conference on PErvasive Technologies Related to Assistive Environments (Petra '17)*. Association for Computing Machinery, New York, NY, USA, 222–229. https://doi.org/10.1145/3056540.3056548

[57] Krzysztof Z. Gajos, Jacob O. Wobbrock, and Daniel S. Weld. 2008. Improving the performance of motor-impaired users with automatically-generated, ability-based interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Chi '08)*. Association for Computing Machinery, New York, NY, USA, 1257–1266. https://doi.org/10.1145/1357054.1357250

[58] Andrea Gallardo, Chris Choy, Jaideep Juneja, Efe Bozkir, Camille Cobb, Lujo Bauer, and Lorrie Cranor. 2023. Speculative Privacy Concerns About AR Glasses Data Collection. *Proceedings on Privacy Enhancing Technologies Symposium* 2023, 4 (2023), 416–435. https://doi.org/10.56553/popets-2023-0117

[59] Susan Gauch, Mirco Speretta, Aravind Chandramouli, and Alessandro Micarelli. 2007. User profiles for personalized information access. In *The adaptive web: Methods and strategies of web personalization*. Springer-Verlag, Berlin, Heidelberg, 54–89.

[60] Dmitri Goldenberg, Kostia Kofman, Javier Albert, Sarai Mizrachi, Adam Horowitz, and Irene Teinemaa. 2021. Personalization in practice: Methods and applications. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (Wsdm '21)*. Association for Computing Machinery, New York, NY, USA, 1123–1126. https://doi.org/10.1145/3437963.3441657

[61] Ileana Maria Greca and Marco Antonio Moreira. 2000. Mental models, conceptual models, and modelling. *International Journal of Science Education* 22, 1 (2000), 1–11. https://doi.org/10.1080/095006900289976

[62] Andy Green, Germ Janmaat, and Helen Cheng. 2011. Social Cohesion: Converging and Diverging Trends. *National Institute Economic Review* 215 (Jan. 2011), R6–R22. https://doi.org/10.1177/0027950111401140

[63] Saul Greenberg, Sebastian Boring, Jo Vermeulen, and Jakub Dostal. 2014. Dark patterns in proxemic interactions: a critical perspective. In *Proceedings of the 2014 Conference on Designing Interactive Systems (Dis '14)*. Association for Computing Machinery, New York, NY, USA, 523–532. https://doi.org/10.1145/2598510.2598541

[64] Travis Greene, Copenhagen Business School, Galit Shmueli, National Tsing Hua University, Soumya Ray, and National Tsing Hua University. 2023. Taking the Person Seriously: Ethically Aware IS Research in the Era of Reinforcement Learning-Based Personalization. *Journal of the Association for Information Systems* 24, 6 (2023), 1527–1561. https://doi.org/10.17705/1jais.00800

[65] Tobias Grosshauser and Thomas Hermann. 2009. Augmented Haptics — An Interactive Feedback System for Musicians. In *Proceedings of the 4th International Conference on Haptic and Audio Interaction Design (HAID '09)*. Springer-Verlag, Berlin, Heidelberg, 100–108. https://doi.org/10.1007/978-3-642-04076-4_11

[66] Jens Grubert, Tobias Langlotz, Stefanie Zollmann, and Holger Regenbrecht. 2017. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 23, 6 (June 2017), 1706–1724. https://doi.org/10.1109/TVCG.2016.2543720

[67] Andrew Guess, Brendan Nyhan, Benjamin Lyons, and Jason Reifler. 2018. *Why selective exposure to like-minded political news is less prevalent than you think*. Technical Report. Knight Foundation. 26 pages.

[68] Jan Gugenheimer, Christian Mai, Mark McGill, Julie Williamson, Frank Steinicke, and Ken Perlin. 2019. Challenges Using Head-Mounted Displays in Shared and Social Spaces. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3290607.3299028

[69] Jan Gugenheimer, Evgeny Stemasov, Julian Frommel, and Enrico Rukzio. 2017. ShareVR: Enabling Co-Located Experiences for Virtual Reality between HMD and Non-HMD Users. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 4021–4033. https://doi.org/10.1145/3025453.3025683

[70] Jan Gugenheimer, Evgeny Stemasov, Harpreet Sareen, and Enrico Rukzio. 2018. FaceDisplay: Towards asymmetric multi-user interaction for nomadic virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Chi '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173628

[71] Xiaonan Guo, Jian Liu, Cong Shi, Hongbo Liu, Yingying Chen, and Mooi Choo Chuah. 2018. Device-free personalized fitness assistant using WiFi. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4 (Dec. 2018), 1–23. https://doi.org/10.1145/3287043

[72] Stefan Göbel, Sandro Hardy, Viktor Wendel, Florian Mehm, and Ralf Steinmetz. 2010. Serious games for health: Personalized exergames. In *Proceedings of the 18th ACM International Conference on Multimedia (Mm '10)*. Association for Computing Machinery, New York, NY, USA, 1663–1666. https://doi.org/10.1145/1873951.1874316

[73] Daniel Görlich. 2022. Societal XR—A Vision Paper. *ParadigmPlus* 3, 2 (Aug. 2022), 1–10. https://doi.org/10.55969/paradigmplus.v3n2a1

[74] Jörg Habich and Verena Nowotny. 2019. *Fragmented Realities – Searching for a Common Understanding of Truth.* Technical Report. Bertelsmann Stiftung, Salzburg. 17 pages.

[75] Hilda Hadan, Lydia Choong, Leah Zhang-Kennedy, and Lennart E. Nacke. 2024. Deceived by Immersion: A Systematic Analysis of Deceptive Design in Extended Reality. *ACM Comput. Surv.* 56, 10 (May 2024), 250:1–250:25. https://doi.org/10.1145/3659945

[76] Julia Hanson, Miranda Wei, Sophie Veys, Matthew Kugler, Lior Strahilevitz, and Blase Ur. 2020. Taking Data Out of Context to Hyper-Personalize Ads: Crowdworkers' Privacy Perceptions and Decisions to Disclose Private Information. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* ACM, Honolulu HI USA, 1–13. https://doi.org/10.1145/3313831.3376415

[77] Marieke Heemskerk, Karen Wilson, and Mitchell Pavao-Zuckerman. 2003. Conceptual Models as Tools for Communication Across Disciplines. *Conservation Ecology* 7, 3 (May 2003), 13. https://doi.org/10.5751/ES-00554-070308

[78] Eelco Herder and Claus Atzenbeck. 2023. 7 Feedback loops and mutual reinforcement in personalized interaction. In *Personalized Human-Computer Interaction*, Mirjam Augstein, Eelco Herder, and Wolfgang Wörndl (Eds.). De Gruyter Oldenbourg, Berlin, Boston, 153–172. https://www.degruyter.com/document/doi/10.1515/9783110988567-007/html

[79] Erik Hermann. 2022. Artificial intelligence and mass personalization of communication content—An ethical and literacy perspective. *New Media & Society* 24, 5 (May 2022), 1258–1277. https://doi.org/10.1177/14614448211022702

[80] Alex Hern. 2022. How TikTok's algorithm made it a success: 'It pushes the boundaries'. https://www.theguardian.com/technology/2022/oct/23/tiktok-rise-algorithm-popularity Last accessed January 10, 2025..

[81] Damian Hostettler. 2022. Increasing the Acceptance of Industrial Robots by Adapting Movement Behavior to Individual User Differences. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22).* Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3491101.3503812

[82] Damian Hostettler, Simon Mayer, and Christian Hildebrand. 2022. Human-like movements of industrial robots positively impact observer perception. *International Journal of Social Robotics* 15 (2022), 1399–1417. https://doi.org/10.1007/s12369-022-00954-2

[83] Fabian Hutmacher and Markus Appel. 2022. The Psychology of Personalization in Digital Environments: From Motivation to Well-Being – A Theoretical Integration. *Review of General Psychology* 27, 1 (June 2022), 108926802211056. https://doi.org/10.1177/10892680221105663

[84] Jane Im, Jill Dimond, Melody Berton, Una Lee, Katherine Mustelier, Mark S. Ackerman, and Eric Gilbert. 2021. Yes: Affirmative Consent as a Theoretical Framework for Understanding and Imagining Social Platforms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* ACM, Yokohama Japan, 1–18. https://doi.org/10.1145/3411764.3445778

[85] Dietmar Jannach and Markus Zanker. 2024. A Survey on Intent-aware Recommender Systems. *ACM Transactions on Recommender Systems* 3, 2 (Oct. 2024), 1–32. https://doi.org/10.1145/3700890

[86] Nitha Elizabeth John, Alessandra Rossi, and Silvia Rossi. 2022. Personalized Human-Robot Interaction with a Robot Bartender. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22 Adjunct).* ACM, New York, NY, USA, 155–159. https://doi.org/10.1145/3511047.3537686

[87] Ross Johnstone, Neil McDonnell, and Julie R. Williamson. 2022. When Virtuality Surpasses Reality: Possible Futures of Ubiquitous XR. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts.* ACM, New Orleans LA USA, 1–8. https://doi.org/10.1145/3491101.3516396

[88] Natascha Just and Michael Latzer. 2017. Governance by algorithms: reality construction by algorithmic selection on the Internet. *Media, Culture & Society* 39, 2 (March 2017), 238–258. https://doi.org/10.1177/0163443716643157

[89] Judy Kay and Bob Kummerfeld. 2012. Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. *ACM Transactions on Interactive Intelligent Systems* 2, 4 (Dec. 2012), 1–42. https://doi.org/10.1145/2395123.2395129

[90] Holger Klapperich, Alarith Uhde, and Marc Hassenzahl. 2020. Designing everyday automation with well-being in mind. *Personal and Ubiquitous Computing* 24, 6 (Dec. 2020), 763–779. https://doi.org/10.1007/s00779-020-01452-w

[91] Tobias Koopmann, Alexander Dallmann, Lena Hettinger, Thomas Niebler, and Andreas Hotho. 2019. On the right track! Analysing and Predicting Navigation Success in Wikipedia. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media.* ACM, Hof Germany, 143–152. https://doi.org/10.1145/3342220.3343650

[92] Nima Kordzadeh and Maryam Ghasemaghaei. 2022. Algorithmic bias: review, synthesis, and future research directions. *European Journal of Information Systems* 31, 3 (2022), 388–409. https://doi.org/10.1080/0960085X.2021.1927212

[93] Veronika Krauss, Pejman Saeghe, Alexander Boden, Mohamed Khamis, Mark McGill, Jan Gugenheimer, and Michael Nebeling. 2024. What Makes XR Dark? Examining Emerging Dark Patterns in Augmented and Virtual Reality through Expert Co-Design. *ACM Trans. Comput.-Hum. Interact.* 31, 3 (April 2024), 39. https://doi.org/10.1145/3660340

[94] John Krumm (Ed.). 2010. *Ubiquitous computing fundamentals.* Chapman & Hall/CRC Press, Boca Ragon, FL, USA.

[95] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking. In *Privacy and Identity Management. Data for Better Living: AI and Privacy*, Michael Friedewald, Melek Önen, Eva Lievens, Stephan Krenn, and Samuel Fricker (Eds.). Vol. 576. Springer International Publishing, Cham, 226–241. http://link.springer.com/10.1007/978-3-030-42504-3_15

[96] Iryna Kuksa, Tom Fisher, and Anthony Kent. 2023. *Understanding personalisation: new aspects of design and consumption.* Chandos Publishing, Cambridge, United Kingdom.

[97] Iryna Kuksa, Michael Skinner, Tony Kent, and Tom Fisher. 2023. The contemporary phenomenon of personalisation. In *Understanding Personalisation*, Iryna Kuksa, Prof Tom Fisher, and Prof Anthony Kent (Eds.). Chandos Publishing, Cambridge, United Kingdom, 3–25.

[98] Amruth N. Kumar, Rajendra K. Raj, Sherif G. Aly, Monica D. Anderson, Brett A. Becker, Richard L. Blumenthal, Eric Eaton, Susan L. Epstein, Michael Goldweber, and Pankaj Jalote. 2024. Multiple Approaches for Teaching Responsible Computing. In *Computer Science Curricula 2023.* Association for Computing Machinery, New York, NY, USA.

[99] Mei-Kei Lai and Yan Yan Cao. 2019. Designing Interactive Olfactory Experience in Real Context and Applications. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction.* ACM, Tempe Arizona USA, 703–706. https://doi.org/10.1145/3294109.3295659

[100] Lik-Hang Lee, Tristan Braud, Simo Hosio, and Pan Hui. 2022. Towards Augmented Reality Driven Human-City Interaction: Current Research on Mobile Headsets and Future Challenges. *Comput. Surveys* 54, 8 (Nov. 2022), 1–38. https://doi.org/10.1145/3467963

[101] Moon-Hwan Lee, Oosung Son, and Tek-Jin Nam. 2016. Patina-inspired personalization: Personalizing products with traces of daily use. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (Dis '16).* Association for Computing Machinery, New York, NY, USA, 251–263. https://doi.org/10.1145/2901790.2901812

[102] Jiayu Li, Zhiyu He, Yumeng Cui, Chenyang Wang, Chong Chen, Chun Yu, Min Zhang, Yiqun Liu, and Shaoping Ma. 2022. Towards Ubiquitous Personalized Music Recommendation with Smart Bracelets. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (Sept. 2022), 125:1–125:34. https://doi.org/10.1145/3550333

[103] Xiangdong Li, Yunzhan Zhou, Wenqian Chen, Preben Hansen, Weidong Geng, and Lingyun Sun. 2019. 8. Towards personalized virtual reality touring through cross-object user interfaces. In *Personalized Human-Computer Interaction*, Mirjam Augstein, Eelco Herder, and Wolfgang Wörndl (Eds.). De Gruyter, Berlin, Boston, 201–222.

[104] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21).* Association for Computing Machinery, New York, NY, USA, 1054–1063. https://doi.org/10.1145/3404835.3462966

[105] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-Aware Online Adaptation of Mixed Reality Interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology.* ACM, New York, NY, USA, 147–160. https://doi.org/10.1145/3332165.3347945

[106] David Lindlbauer and Andy D. Wilson. 2018. Remixed reality: Manipulating space and time in augmented reality. In *Proceedings of the 2018 CHI conference on human factors in computing systems (Chi '18).* Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173703

[107] Liquid City and Niantic. 2022. Stories from the Real-World Metaverse. https://liquid.city/stories/ Last accessed January 13, 2025.

[108] Danielle Lottridge, Katie Quehl, Frank Bentley, Max Silverman, Melissa Ong, Michael Dickard, Brooke White, and Ethan Plaut. 2022. Ubiquitous News Experienced Alone: Interviews with Americans and their Devices. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (March 2022), 1–29. https://doi.org/10.1145/3512975

[109] Andrés Lucero, Peter Dalsgaard, Kim Halskov, and Jacob Buur. 2016. Designing with Cards. In *Collaboration in Creative Design: Methods and Tools*, Panos Markopoulos, Jean-Bernard Martens, Julian Malins, Karin Coninx, and Aggelos Liapis (Eds.). Springer International Publishing, Cham, 75–95. https://doi.org/10.1007/978-3-319-29155-0_5

[110] Damien Lutz. 2023. The Maya Lenz - A speculative design concept for future AR. https://www.damienlutz.com.au/thelenz/ Last accessed January 12, 2025.

[111] Z. Maamar, S.K. Mostefaoui, and Q.H. Mahmoud. 2005. Context for personalized web services. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences.* IEEE, Big Island, HI, USA, 166b–166b. https://doi.org/10.1109/HICSS.2005.164

[112] Michael Madary. 2023. Mediated Reality. In *Exploring Extended Realities* (1 ed.), Andrew Kissel and Erick José Ramirez (Eds.). Routledge, New York, 19.

[113] Bernard Mallia. 2023. Designing and Implementing Socially Beneficial Recommender Systems: An Interdisciplinary Approach. In *NORMalize 2023: The First Workshop on the Normative Design and Evaluation of Recommender Systems*. CEUR Workshop Proceedings, Singapore, 10. https://ceur-ws.org/Vol-3639/paper2.pdf

[114] Florian Mathis. 2024. Everyday Life Challenges and Augmented Realities: Exploring Use Cases For, and User Perspectives on, an Augmented Everyday Life. In *Proceedings of the Augmented Humans International Conference 2024 (AHs '24)*. Association for Computing Machinery, New York, NY, USA, 52–62. https://doi.org/10.1145/3652920.3652921

[115] Keiichi Matsuda. 2016. Hyper-Reality. http://hyper-reality.co/ Last accessed January 12, 2025.

[116] Neil McDonnell and Nathan Wildman. 2019. Virtual Reality: Digital or Fictional? *Disputatio* 11, 55 (Dec. 2019), 371–397. https://doi.org/10.2478/disp-2019-0004

[117] Mark McGill, Daniel Boland, Roderick Murray-Smith, and Stephen Brewster. 2015. A dose of reality: Overcoming usability challenges in VR head-mounted displays. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems (Chi '15)*. Association for Computing Machinery, New York, NY, USA, 2143–2152. https://doi.org/10.1145/2702123.2702382

[118] Abraham Hani Mhaidli and Florian Schaub. 2021. Identifying Manipulative Advertising Techniques in XR Through Scenario Construction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3411764.3445253

[119] Mateusz Mikusz, Peter Shaw, Nigel Davies, Petteri Nurmi, Sarah Clinch, Ludwig Trotter, Ivan Elhart, Marc Langheinrich, and Adrian Friday. 2021. A Longitudinal Study of Pervasive Display Personalisation. *ACM Transactions on Computer-Human Interaction* 28, 1 (Jan. 2021), 2:1–2:45. https://doi.org/10.1145/3418352

[120] Paul Milgram and Fumio Kishino. 1994. A Taxonomy of Mixed Reality Visual Displays. *IEICE Trans. Information Systems* vol. E77-D, no. 12 (Dec. 1994), 1321–1329.

[121] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. 1995. Augmented reality: a class of displays on the reality-virtuality continuum. In *Proc. SPIE 2351*. SPIE, Boston, MA, United States, 282–292. https://doi.org/10.1117/12.197321

[122] Pardis Miri, Robert Flory, Andero Uusberg, Heather Culbertson, Richard H. Harvey, Agata Kelman, Davis Erik Peper, James J. Gross, Katherine Isbister, and Keith Marzullo. 2020. PIV: Placement, pattern, and personalization of an inconspicuous vibrotactile breathing pacer. *ACM Trans. Comput.-Hum. Interact.* 27, 1 (Jan. 2020), 1–44. https://doi.org/10.1145/3365107

[123] Hayoun Moon, Mia Saade, Daniel Enriquez, Zachary Duer, Hye Sung Moon, Sang Won Lee, and Myounghoon Jeon. 2024. Mixed-reality art as shared experience for cross-device users: Materialize, understand, and explore. *International Journal of Human-Computer Studies* 190 (Oct. 2024), 103291. https://doi.org/10.1016/j.ijhcs.2024.103291

[124] Cataldo Musto, Amra Delic, Oana Inel, Marco Polignano, Amon Rapp, Giovanni Semeraro, and Jürgen Ziegler. 2022. Workshop on Explainable User Models and Personalised Systems (ExUM). In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 160–162. https://doi.org/10.1145/3511047.3536350

[125] Florian Müller, Martin Schmitz, Markus Funk, Sebastian Günther, Niloofar Dezfuli, and Max Mühlhäuser. 2018. Personalized User-Carried Single Button Interfaces as Shortcuts for Interacting with Smart Devices. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, 1–6. https://doi.org/10.1145/3170427.3188661

[126] Ken Nakagaki, Artem Dementyev, Sean Follmer, Joseph A. Paradiso, and Hiroshi Ishii. 2016. ChainFORM: A Linear Integrated Modular Hardware System for Shape Changing Interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 87–96. https://doi.org/10.1145/2984511.2984587

[127] Alaeddin Nassani, Huidong Bai, Gun Lee, Mark Billinghurst, Tobias Langlotz, and Robert W. Lindeman. 2018. Filtering Shared Social Data in AR. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–6. https://doi.org/10.1145/3170427.3188609

[128] Anton Nijholt. 2021. Experiencing Social Augmented Reality in Public Spaces. In *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers*. ACM, Virtual USA, 570–574. https://doi.org/10.1145/3460418.3480157

[129] Anton Nijholt. 2023. Toward an Ever-present Extended Reality: Distinguishing Between Real and Virtual. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*. ACM, Cancun, Quintana Roo Mexico, 396–399. https://doi.org/10.1145/3594739.3610726

[130] Donald A. Norman. 1983. Some Observations on Mental Models. In *Mental Models* (1 ed.), Dedre Gentner and Albert L. Stevens (Eds.). Psychology Press, New York.

[131] Office of the California Attorney General. 2018. California Consumer Privacy Act of 2018. https://www.oag.ca.gov/privacy/ccpa

[132] Joseph O'Hagan, Jan Gugenheimer, Jolie Bonner, Florian Mathis, and Mark McGill. 2023. Augmenting People, Places & Media: The Societal Harms Posed by Everyday Augmented Reality, and the Case for Perceptual Human Rights. In *Proceedings of the 22nd International Conference on Mobile and Ubiquitous Multimedia (MUM '23)*. Association for Computing Machinery, New York, NY, USA, 219–229. https://doi.org/10.1145/3626705.3627782

[133] Joseph O'Hagan, Pejman Saeghe, Jan Gugenheimer, Daniel Medeiros, Karola Marky, Mohamed Khamis, and Mark McGill. 2023. Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (Jan. 2023), 177:1–177:35. https://doi.org/10.1145/3569501

[134] Reinhard Oppermann. 1994. Adaptively supported adaptability. *International Journal of Human-Computer Studies* 40, 3 (March 1994), 455–472. https://doi.org/10.1006/ijhc.1994.1021

[135] Lorenzo Orlandi, Giulia Martinelli, Francesco Laiti, Davide Lobba, Niccoló Bisagno, and Nicola Conci. 2023. Meta-Trainer: An Augmented Reality Trainer for Home Fitness with Real-Time Feedback. In *2023 IEEE International Workshop on Sport, Technology and Research (STAR)*. IEEE, Trento, Italy, 90–93. https://doi.org/10.1109/STAR58331.2023.10302670

[136] Jason Orlosky, Misha Sra, Kenan Bektaş, Huaishu Peng, Jeeeun Kim, Nataliya Kos'myna, Tobias Höllerer, Anthony Steed, Kiyoshi Kiyokawa, and Kaan Akşit. 2021. Telelife: The Future of Remote Living. *Frontiers in Virtual Reality* 2 (2021), 19. https://www.frontiersin.org/articles/10.3389/frvir.2021.763340

[137] Vito Ostuni, Christoph Kofler, Manjesh Nilange, Sudarshan Lamkhede, and Dan Zylberglejd. 2023. Search personalization at netflix. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 756–758. https://doi.org/10.1145/3543873.3587675

[138] Aviv Ovadya and Luke Thorburn. 2023. Bridging Systems: Open Problems for Countering Destructive Divisiveness across Ranking, Recommenders, and Governance. http://arxiv.org/abs/2301.09976

[139] Richard Owen, Phil Macnaghten, and Jack Stilgoe. 2012. Responsible research and innovation: From science in society to science for society, with society. *Science and Public Policy* 39, 6 (Dec. 2012), 751–760. https://doi.org/10.1093/scipol/scs093

[140] Joseph O'Hagan, Jan Gugenheimer, Florian Mathis, Jolie Bonner, Richard Jones, and Mark McGill. 2024. A Viewpoint On... The Societal Impact of Everyday Augmented Reality, and the Need for Perceptual Human Rights. *IEEE Security and Privacy Magazine* 22, 07 (2024), 6. https://doi.org/10.1109/MSEC.2023.3333988

[141] Dimitris Paraschakis. 2017. Towards an ethical recommendation framework. In *2017 11th International Conference on Research Challenges in Information Science (RCIS)*. IEEE, Brighton, UK, 211–220. https://doi.org/10.1109/RCIS.2017.7956539

[142] Sugandh Pargal, Debasree Das, Bikash Sahoo, Bivas Mitra, and Sandip Chakraborty. 2023. GRIDS: Personalized guideline recommendations while driving through a new city. *ACM Trans. Recomm. Syst.* 2, 2 (July 2023), 28. https://doi.org/10.1145/3609337

[143] Taiwoo Park, Mi Zhang, and Youngki Lee. 2018. When Mixed Reality Meets Internet of Things: Toward the Realization of Ubiquitous Mixed Reality. *GetMobile: Mobile Computing and Communications* 22, 1 (May 2018), 10–14. https://doi.org/10.1145/3229316.3229320

[144] Dorian Peters, Karina Vold, Diana Robinson, and Rafael A. Calvo. 2020. Responsible AI—Two Frameworks for Ethical Design Practice. *IEEE Transactions on Technology and Society* 1, 1 (March 2020), 34–47. https://doi.org/10.1109/TTS.2020.2974991

[145] Filip Petrushevski. 2012. Personalized lighting control based on a space model. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. Association for Computing Machinery, New York, NY, USA, 568–571. https://doi.org/10.1145/2370216.2370311

[146] Ken Pfeuffer, Yasmeen Abdrabou, Augusto Esteves, Radiah Rivu, Yomna Abdelrahman, Stefanie Meitner, Amr Saadi, and Florian Alt. 2021. ARtention: A design space for gaze-adaptive user interfaces in augmented reality. *Computers & Graphics* 95 (April 2021), 1–12. https://doi.org/10.1016/j.cag.2021.01.001

[147] Robin Piening, Ken Pfeuffer, Augusto Esteves, Tim Mittermeier, Philippe Schroder, and Florian Alt. 2021. Looking for Info: Evaluation of Gaze Based Information Retrieval in Augmented Reality. In *Human-Computer Interaction – INTERACT 2021*. Springer International Publishing, Cham, 544–565. https://doi.org/10.1007/978-3-030-85623-6_32

[148] James Pierce, Claire Weizenegger, Parag Nandi, Isha Agarwal, Gwenna Gram, Jade Hurrle, Hannah Liao, Betty Lo, Aaron Park, Aivy Phan, Mark Shumskiy, and Grace Sturlaugson. 2022. Addressing adjacent actor privacy: Designing for bystanders, co-users, and surveilled subjects of smart home cameras. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference (DIS '22)*. Association for Computing Machinery, New York, NY, USA, 26–40. https://doi.org/10.1145/3532106.3535195

[149] Kathrin Pollmann and Daniel Ziegler. 2020. Personal quizmaster: a pattern approach to personalized interaction experiences with the MiRo robot. In *Proceedings of Mensch und Computer 2020 (MuC '20)*. ACM, New York, NY, USA, 485–489. https://doi.org/10.1145/3404983.3410414

[150] Patricia Pons, Javier Jaen, and Alejandro Catala. 2017. Towards Future Interactive Intelligent Systems for Animals: Study and Recognition of Embodied Interactions. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, Limassol Cyprus, 389–400. https://doi.org/10.1145/3025171.3025175

[151] Dominic Potts, Martynas Dabravalskis, and Steven Houben. 2022. TangibleTouch: A Toolkit for Designing Surface-based Gestures for Tangible Interfaces. In *Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '22)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3490149.3502263

[152] Jeffrey Prince and Scott Wallsten. 2021. How Much is Privacy Worth Around the World and Across Platforms?. In *TPRC48: The 48th Research Conference on Communication, Information and Internet Policy*. TPRC, Virtual Event, 70. https://doi.org/10.2139/ssrn.3528386

[153] Wajid Rafique, Abdelhakim Senhaji Hafid, and Junaid Qadir. 2023. Developing smart city services using intent-aware recommendation systems: A survey. *Transactions on Emerging Telecommunications Technologies* 34, 4 (2023), e4728. https://doi.org/10.1002/ett.4728

[154] Shwetha Rajaram, Chen Chen, Franziska Roesner, and Michael Nebeling. 2023. Eliciting Security & Privacy-Informed Sharing Techniques for Multi-User Augmented Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–17. https://doi.org/10.1145/3544548.3581089

[155] Aditi Ramachandran, Chien-Ming Huang, and Brian Scassellati. 2017. Give me a break! Personalized timing strategies to promote learning in robot-child tutoring. In *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction (Hri '17)*. Association for Computing Machinery, New York, NY, USA, 146–155. https://doi.org/10.1145/2909824.3020209

[156] Nimesha Ranasinghe and Ellen Yi-Luen Do. 2016. Virtual Sweet: Simulating Sweet Sensation Using Thermal Stimulation on the Tip of the Tongue. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16 Adjunct)*. ACM, New York, NY, USA, 127–128. https://doi.org/10.1145/2984751.2985729

[157] Philipp A. Rauschnabel, Reto Felix, Chris Hinsch, Hamza Shahab, and Florian Alt. 2022. What is XR? Towards a Framework for Augmented and Virtual Reality. *Computers in Human Behavior* 133 (Aug. 2022), 107289. https://doi.org/10.1016/j.chb.2022.107289

[158] Holger Regenbrecht, Alistair Knott, Jennifer Ferreira, and Nadia Pantidi. 2024. To See and be Seen—Perceived Ethics and Acceptability of Pervasive Augmented Reality. *IEEE Access* 12 (2024), 32618–32636. https://doi.org/10.1109/ACCESS.2024.3366228

[159] Holger Regenbrecht, Sander Zwanenburg, and Tobias Langlotz. 2022. Pervasive Augmented Reality—Technology and Ethics. *IEEE Pervasive Computing* 21, 3 (2022), 1–8. https://doi.org/10.1109/MPRV.2022.3152993

[160] Jan Ole Rixen, Christian Funk, Enrico Rukzio, and Jan Gugenheimer. 2023. May I Still Define Myself? Exploring How Dissonance in Displaying Personal Information Through Head-Mounted Augmented Reality Can Affect Personal Information Sovereignty. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–7. https://doi.org/10.1145/3544549.3585821

[161] Franziska Roesner, Tamara Denning, Bryce Clayton Newell, Tadayoshi Kohno, and Ryan Calo. 2014. Augmented reality: hard problems of law and policy. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, Seattle Washington, 1283–1288. https://doi.org/10.1145/2638728.2641709

[162] Joan Sol Roo and Martin Hachet. 2017. One Reality: Augmenting How the Physical World is Experienced by combining Multiple Mixed Reality Modalities. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, Québec City QC Canada, 787–795. https://doi.org/10.1145/3126594.3126638

[163] Louis B. Rosenberg. 2022. Augmented Reality: Reflections at Thirty Years. In *Proceedings of the Future Technologies Conference (FTC) 2021, Volume 1 (Lecture Notes in Networks and Systems)*, Kohei Arai (Ed.). Springer International Publishing, Cham, 1–11. https://doi.org/10.1007/978-3-030-89906-6_1

[164] Kimberly Ruth, Tadayoshi Kohno, and Franziska Roesner. 2019. Secure {Multi-User} Content Sharing for Augmented Reality Applications. In *Proceedings of the 28th USENIX Security Symposium*. USENIX Association, Santa Clara, CA, 141–158. https://www.usenix.org/conference/usenixsecurity19/presentation/ruth

[165] Daniel Salber, Anind K. Dey, and Gregory D. Abowd. 1999. The context toolkit: aiding the development of context-enabled applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. Association for Computing Machinery, New York, NY, USA, 434–441. https://doi.org/10.1145/302979.303126

[166] Ville Salonen and Heikki Karjaluoto. 2016. Web personalization: The state of the art and future avenues for research and practice. *Telematics and Informatics* 33, 4 (Nov. 2016), 1088–1104. https://doi.org/10.1016/j.tele.2016.03.004

[167] Laura Schelenz, Avi Segal, Oduma Adelio, and Kobi Gal. 2023. Transparency-Check: An Instrument for the Study and Design of Transparency in AI-based Personalization Systems. *ACM Journal on Responsible Computing* 1, 1 (Dec. 2023), 1–18. https://doi.org/10.1145/3636508

[168] Matthias Schmidmaier, Zhiwei Han, Thomas Weber, Yuanting Liu, and Heinrich Hußmann. 2019. Real-time personalization in adaptive ides. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization (UMAP'19 adjunct)*. Association for Computing Machinery, New York, NY, USA, 81–86. https://doi.org/10.1145/3314183.3324975

[169] Albrecht Schmidt, Michael Beigl, and Hans-W Gellersen. 1999. There is more to context than location. *Computers & Graphics* 23, 6 (1999), 893–901. https://doi.org/10.1016/S0097-8493(99)00120-X

[170] Clive Seale (Ed.). 2018. *Researching society and culture* (4th edition ed.). Sage, Los Angeles; London.

[171] Arne Seeliger, Raphael P. Weibel, and Stefan Feuerriegel. 2022. Context-adaptive visual cues for safe navigation in augmented reality using machine learning. *International Journal of Human–Computer Interaction* 0, 0 (2022), 1–21. https://doi.org/10.1080/10447318.2022.2122114

[172] Adalberto L. Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3307–3316. https://doi.org/10.1145/2702123.2702389

[173] Rick Skarbez, Missie Smith, and Mary Whitton. 2023. It Is Time to Let Go of 'Virtual Reality'. *Commun. ACM* 66, 10 (Oct. 2023), 41–43. https://doi.org/10.1145/3590959

[174] Richard Skarbez, Missie Smith, and Mary C. Whitton. 2021. Revisiting Milgram and Kishino's Reality-Virtuality Continuum. *Frontiers in Virtual Reality* 2 (March 2021), 647997. https://doi.org/10.3389/frvir.2021.647997

[175] Mel Slater, Cristina Gonzalez-Liencres, Patrick Haggard, Charlotte Vinkers, Rebecca Gregory-Clarke, Steve Jelley, Zillah Watson, Graham Breen, Raz Schwarz, William Steptoe, Dalila Szostak, Shivashankar Halan, Deborah Fox, and Jeremy Silver. 2020. The Ethics of Realism in Virtual and Augmented Reality. *Frontiers in Virtual Reality* 1 (March 2020), 13. https://doi.org/10.3389/frvir.2020.00001

[176] Brent Smith and Greg Linden. 2017. Two decades of recommender systems at amazon.com. *IEEE Internet Computing* 21, 3 (2017), 12–18. https://doi.org/10.1109/MIC.2017.72

[177] Maximilian Speicher, Brian D. Hall, and Michael Nebeling. 2019. What is Mixed Reality?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–15. https://doi.org/10.1145/3290605.3300767

[178] Janick Spirig, Kimberly Garcia, and Simon Mayer. 2021. An Expert Digital Companion for Working Environments. In *Proceedings of the 11th International Conference on the Internet of Things (IoT '21)*. Association for Computing Machinery, New York, NY, USA, 25–32. https://doi.org/10.1145/3494322.3494326

[179] Misha Sra, Abhinandan Jain, and Pattie Maes. 2019. Adding Proprioceptive Feedback to Virtual Reality Experiences Using Galvanic Vestibular Stimulation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 1–14. https://doi.org/10.1145/3290605.3300905

[180] Jaideep Srivastava, Robert Cooley, Mukund Deshpande, and Pang-Ning Tan. 2000. Web usage mining: discovery and applications of usage patterns from Web data. *ACM SIGKDD Explorations Newsletter* 1, 2 (Jan. 2000), 12–23. https://doi.org/10.1145/846183.846188

[181] Christopher Starke, Janine Baleis, Birte Keller, and Frank Marcinkowski. 2022. Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society* 9, 2 (2022), 20539517221115189. https://doi.org/10.1177/20539517221115189

[182] Nick Statt. 2018. Google personalizes search results even when you're logged out, new study claims. https://www.theverge.com/2018/12/4/18124718/google-search-results-personalized-unique-duckduckgo-filter-bubble

[183] Jannis Strecker, Luka Bekavac, Kenan Bektaş, and Simon Mayer. 2025. Change Your Perspective, Widen Your Worldview! Societally Beneficial Perceptual Filter Bubbles in Personalized Reality. In *Proceedings of Purposeful XR: Affordances, Challenges, and Speculations for an Ethical Future at CHI 2025 (Purposeful XR '25)*. Yokohama, Japan, 6. https://doi.org/10.48550/arXiv.2504.10271

[184] Jannis Strecker, Kimberly García, Kenan Bektaş, Simon Mayer, and Ganesh Ramanathan. 2022. SOCRAR: Semantic OCR through Augmented Reality. In *Proceedings of the 12th International Conference on the Internet of Things (IoT '22)*. ACM, New York, NY, USA, 8. https://doi.org/10.1145/3567445.3567453

[185] Jannis Strecker, Simon Mayer, and Kenan Bektas. 2024. Personalized Reality: Challenges of Responsible Ubiquitous Personalization. In *Mensch und Computer 2024 - Workshopband*. Gesellschaft für Informatik e.V., Karlsruhe, Germany, 10.18420/muc2024. https://doi.org/10.18420/muc2024-mci-ws11-200

[186] Jannis Strecker, Simon Mayer, and Kenan Bektaş. 2023. Sharing Personalized Mixed Reality Experiences. In *Mensch und Computer 2023 – Workshopband*. Gesellschaft für Informatik e.V., Rapperswil, SG, 3. https://doi.org/10.18420/

muc2023-mci-ws12-263
[187] Jannis Strecker, Jing Wu, Kenan Bektaş, Conrad Vaslin, and Simon Mayer. 2024. ShoppingCoach: Using Diminished Reality to Prevent Unhealthy Food Choices in an Offline Supermarket Scenario. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3613905.3650795
[188] Naomi Stricker, Jan Liam Albert, Md Masoom Rabbani, Simon Mayer, and Andres Gomez. 2022. Secure Communication with Batteryless Sensors. In *2022 11th Mediterranean Conference on Embedded Computing (MECO)*. IEEE, Budva, Montenegro, 1–4. https://doi.org/10.1109/MECO55406.2022.9797104
[189] S. Shyam Sundar and Sampada S. Marathe. 2010. Personalization versus Customization: The Importance of Agency, Privacy, and Power Usage. *Human Communication Research* 36, 3 (June 2010), 298–322. https://doi.org/10.1111/j.1468-2958.2010.01377.x
[190] Hilary Sutcliffe. 2011. *A report on Responsible Research & Innovation*. Technical Report. MATTER, European Commission.
[191] Ye Tao, Youngwook Do, Humphrey Yang, Yi-Chin Lee, Guanyun Wang, Catherine Mondoa, Jianxun Cui, Wen Wang, and Lining Yao. 2019. Morphlour: Personalized Flour-based Morphing Food Induced by Dehydration or Hydration Method. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (UIST '19)*. Association for Computing Machinery, New York, NY, USA, 329–340. https://doi.org/10.1145/3332165.3347949
[192] Yujie Tao and Pedro Lopes. 2022. Integrating Real-World Distractions into Virtual Reality. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–16. https://doi.org/10.1145/3526113.3545682
[193] Ludovic Terren and Rosa Borge-Bravo. 2021. Echo Chambers on Social Media: A Systematic Review of the Literature. *Review of Communication Research* 9 (March 2021), 99–118. https://doi.org/10.12840/ISSN.2255-4165.028
[194] Luke Thorburn, Maria Polukarov, and Carmine Ventre. 2024. Societal Sorting as a Systemic Risk of Recommenders. In *18th ACM Conference on Recommender Systems*. ACM, Bari Italy, 951–956. https://doi.org/10.1145/3640457.3688175
[195] G. S. Thyagaraju, M. M. Math, Umakant P. Kulkarni, and A. R. Yardi. 2009. Conflict Resolving Algorithms to Resolve Conflict in Multi-user Context-Aware Environments. In *2009 IEEE International Advance Computing Conference*. IEEE, Patiala, India, 202–208. https://doi.org/10.1109/IADCC.2009.4809007
[196] Horst Treiblmaier and Irene Pollach. 2007. Users' Perceptions of Benefits and Costs of Personalization. In *ICIS 2007 Proceedings*. Association for Information Systems, Montreal, QC, Canada, 16. http://aisel.aisnet.org/icis2007/141
[197] Chung Hsien Tsai and Jiung Yao Huang. 2014. A Mobile Augmented Reality Based Scaffolding Platform for Outdoor Fieldtrip Learning. In *2014 IIAI 3rd International Conference on Advanced Applied Informatics*. IEEE, Kokura, Japan, 307–312. https://doi.org/10.1109/IIAI-AAI.2014.70
[198] UNESCO. 2023. *Ethical impact assessment. A tool of the Recommendation on the Ethics of Artificial Intelligence.* Technical Report. UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000386276
[199] Danai Vachtsevanou, Andrei Ciortea, Simon Mayer, and Jérémy Lemée. 2023. Signifiers as a first-class abstraction in hypermedia multi-agent systems. In *Proceedings of the 2023 international conference on autonomous agents and multiagent systems (Aamas '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1200–1208.
[200] Jari Vesanen. 2007. What is personalization? A conceptual framework. *European Journal of Marketing* 41, 5/6 (June 2007), 409–418. https://doi.org/10.1108/03090560710737534
[201] Jari Vesanen and Mika Raulas. 2006. Building bridges for personalization: A process model for marketing. *Journal of Interactive Marketing* 20, 1 (Feb. 2006), 5–20. https://doi.org/10.1002/dir.20052
[202] Federica Lucia Vinella, Amra Delić, Francesco Barile, Ioanna Lykourentzou, and Judith Masthoff. 2022. GMAP 2022: Workshop on Group Modeling, Adaptation and Personalization. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22 adjunct)*. Association for Computing Machinery, New York, NY, USA, 280–282. https://doi.org/10.1145/3511047.3536351
[203] René von Schomberg. 2013. A Vision of Responsible Research and Innovation. In *Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society*, Richard Owen, John Bessant, and Maggy Heintz (Eds.). John Wiley & Sons, Ltd, Chichester, West Sussex, 51–74. https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118551424.ch3
[204] Julius von Willich, Markus Funk, Florian Müller, Karola Marky, Jan Riemann, and Max Mühlhäuser. 2019. You Invaded my Tracking Space! Using Augmented Virtuality for Spotting Passersby in Room-Scale Virtual Reality. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. Association for Computing Machinery, New York, NY, USA, 487–496. https://doi.org/10.1145/3322276.3322334
[205] Peter Wei, Yanchen Liu, Hengjiu Kang, Chenye Yang, and Xiaofan Jiang. 2021. A Low-Cost and Scalable Personalized Thermal Comfort Estimation System in Indoor Environments. In *Proceedings of the First International Workshop on Cyber-Physical-Human System Design and Implementation (CPHS21)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3458648.3460006
[206] Mark Weiser. 1991. The Computer for the 21st Century. *SCIENTIFIC AMERICAN* 265, 3 (1991), 12. https://doi.org/10.1038/scientificamerican0991-94
[207] Michele Willson. 2017. Algorithms (and the) everyday. *Information, Communication & Society* 20, 1 (Jan. 2017), 137–150. https://doi.org/10.1080/1369118X.2016.1200645
[208] Philipp Wintersberger, Ambika Shahu, Johanna Reisinger, Fatemeh Alizadeh, and Florian Michahelles. 2022. Self-Balancing Bicycles: Qualitative Assessment and Gaze Behavior Evaluation. In *Proceedings of the 21st International Conference on Mobile and Ubiquitous Multimedia (MUM '22)*. Association for Computing Machinery, New York, NY, USA, 189–199. https://doi.org/10.1145/3568444.3568451
[209] Marty J. Wolf, Frances Grodzinsky, and Keith Miller. 2016. Augmented reality all around us: power and perception at a crossroads. *ACM SIGCAS Computers and Society* 45, 3 (Jan. 2016), 126–131. https://doi.org/10.1145/2874239.2874257
[210] Chuhan Wu, Fangzhao Wu, Yongfeng Huang, and Xing Xie. 2023. Personalized News Recommendation: Methods and Challenges. *ACM Transactions on Information Systems* 41, 1 (Jan. 2023), 24:1–24:50. https://doi.org/10.1145/3530257
[211] Yu-Chian Wu, Te-Yen Wu, Paul Taele, Bryan Wang, Jun-You Liu, Pin-sung Ku, Po-En Lai, and Mike Y. Chen. 2018. ActiveErgo: Automatic and Personalized Ergonomics using Self-actuating Furniture. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 1–8. https://doi.org/10.1145/3173574.3174132
[212] Felix Yang, Saikishore Kalloori, Ribin Chalumattu, and Markus Gross. 2022. Personalized Information Retrieval for Touristic Attractions in Augmented Reality. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (Wsdm '22)*. Association for Computing Machinery, New York, NY, USA, 1613–1616. https://doi.org/10.1145/3488560.3502194
[213] Jing Yang, Amit Barde, and Mark Billinghurst. 2022. Audio Augmented Reality: A Systematic Review of Technologies, Applications, and Future Research Directions. *Journal of the Audio Engineering Society* 70, 10 (Nov. 2022), 788–809. https://doi.org/10.17743/jaes.2022.0048
[214] Karen Yeung. 2017. 'Hypernudge': Big Data as a Mode of Regulation by Design. *Information, Communication & Society* 20, 1 (2017), 118–136. https://doi.org/10.1080/1369118X.2016.1186713
[215] Yiran Zhao, Yujie Tao, Grace Le, Rui Maki, Alexander Adams, Pedro Lopes, and Tanzeem Choudhury. 2023. Affective touch as immediate and passive wearable intervention. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–23. https://doi.org/10.1145/3569484

A An Earlier Personalization Model

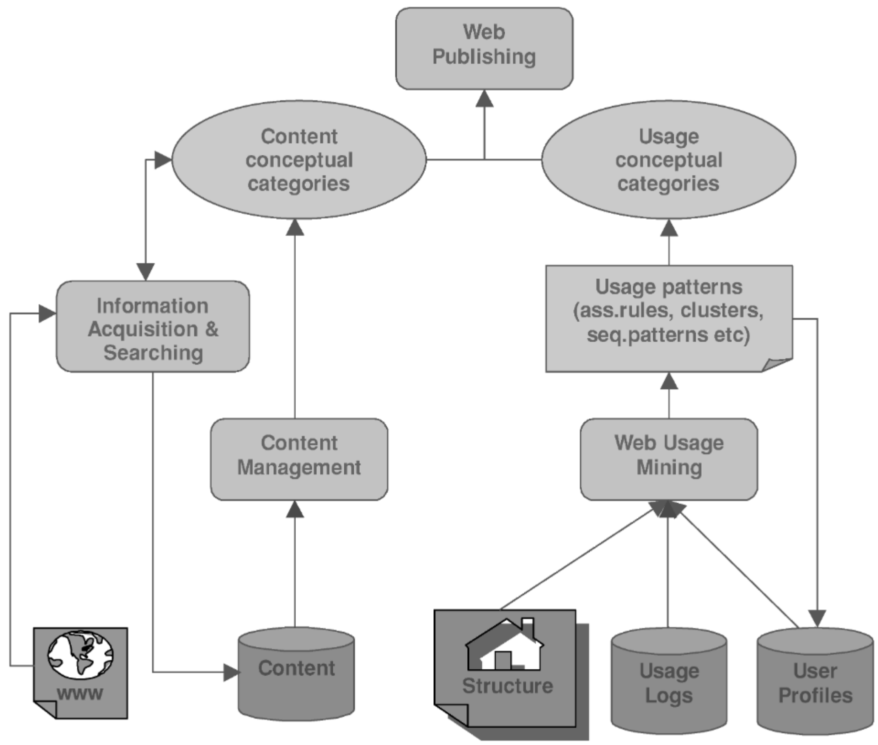

Figure 5: The early Web personalization model from Eirinaki and Vazirgiannis: "Modules of a Web personalization system" (from [48, p.5]). Their model shows the system components for personalization (e.g., content, or user profiles) while the human element (e.g., biases, or context) is missing. See Section 3.2 for a discussion of this and other models.

B Mapping of Selected Works to Existing Personalization Models

In Table 3 and Table 4 we mapped selected works to the two personalization models discussed in Section 3. See Section 5 for more details on the mapping.

Table 3: The mapping of selected works to Eirinaki and Vazirgiannis' Model [48] (see Figure 5).

Source	WWW	Content	Structure	Usage Logs	User Profiles	Web Usage Mining	Web Publishing
[211]	-	Furniture (chair, desk, keyboard, monitor)	-	Human mesh reconstruction	-	-	-
[101]	-	Surface of a leather bag	-	-	-	-	-
[212]	-	Information about Points-of-Interest (POI) in AR	-	-	Name, Gender, Interest areas, Personality traits	-	-

Table 4: The mapping of selected works to Bozdag's Model [48] (see Figure 1).

Source	Information	Source Selection, and Collection Algorithm	Information Selection/ Prioritizing Algorithm	Human Operator	Personalization Algorithm	Presentation Algorithm	Receiver
[211]	Human mesh reconstruction	-	-	ergonomics standards and guidelines	Ergonomic furniture configurations	-	-
[101]	Historical bicycling route logs	-	-	-	Patterns of personal bicycle logs to be printed on a leather bag	-	-
[212]	Information about Points-of-Interest (POI) in AR, Name, Gender, Interest areas, Personality traits, Location, Points of Interest (POI)	-	-	-	Personalized top facts about POI	-	-