

Player-Centered AI for Automatic Game Personalization: Open Problems

Jichen Zhu
jichen.zhu@gmail.com
Drexel University
Philadelphia, PA

Santiago Ontañón*
so367@drexel.edu
Drexel University
Philadelphia, PA

ABSTRACT

Computer games represent an ideal research domain for the next generation of personalized digital applications. This paper presents a player-centered framework of AI for game personalization, complementary to the commonly used system-centered approaches. Built on the Structure of Actions theory, the paper maps out the current landscape of game personalization research and identifies eight open problems that need further investigation. These problems require deep collaboration between technological advancement and player experience design.

CCS CONCEPTS

- Computing methodologies → Artificial intelligence; • Human-centered computing → User models.

KEYWORDS

Player Modeling, Game Personalization, Adaptive Games

ACM Reference Format:

Jichen Zhu and Santiago Ontañón. 2020. Player-Centered AI for Automatic Game Personalization: Open Problems. In *International Conference on the Foundations of Digital Games (FDG '20), September 15–18, 2020, Bugibba, Malta*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3402942.3402951>

1 INTRODUCTION

Automated personalization is becoming an integral part of everyday life. We consume products suggested to us by recommendation systems. We find out what is going on with friends and in the world through content feeds curated for us individually. More and more, our games and digital apps can figure out our needs and preferences and adapt accordingly. In education, adaptive technology allows students to learn at their own speed [70]. Personalization is not limited to the digital world; it is also transforming the manufacturing industry [30, 65]. These advances propel the business world to anticipate the next evolution: the Internet of Me, where mass personalization is driven by users' individual characteristics such as biology [2, 30, 31].

*Currently at Google.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FDG '20, September 15–18, 2020, Bugibba, Malta

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8807-8/20/09...\$15.00
<https://doi.org/10.1145/3402942.3402951>

A significant amount of research has been devoted to automatic personalization in digital applications, especially in Internet applications [8]. As the content of the Internet services grows, personalized applications such as recommendation systems help to mitigate information overload and decision fatigue [8]. This body of work ranges from relatively simple changes on a web page (e.g., using the name of each user) to complex customization using deeper models of user needs and behavior [29].

Computer games are a relatively new domain for personalization. Compared to the classic personalization domains of information seeking and e-commerce, people play games for a broader range of reasons (e.g., challenge, exploration, aesthetic experience, and social activity). It is hence more difficult to identify individual players' needs and preferences that the games should adapt to. Besides, computer games generally involve more complex content and user interaction than other digital applications such as websites. Typical gameplay is multi-sensory (e.g., visual, auditory, and tactile) and contains multiple layers of meanings (e.g., formal rules and stories). To personalize games thus requires further technological advancement (in how to procedurally adapt more complex game content) and new design principles (of how to personalize for various player needs) than what we have learned from the classic domains of personalization. Computer games, therefore, are an excellent domain for researching the next stage of personalization technology. In this paper, we adopt Bakkes, Tan, and Pisan's definition [4] that personalized games are those games that adapt themselves based on information about the current player, e.g. by determining the difficulty level appropriate to the current player automatically.

So far researchers have made significant inroads in personalized games. For example, player modeling has been used to capture player behavior [12, 21, 23, 67], player types [5, 22], user preferences [52, 62], skill level [26, 36, 77], knowledge [27] and player goals [19]. Those player models have been used to adjust different aspects of player experience, such as game narrative [52, 62], level difficulty [25], or generated game content [51, 56, 63, 68].

This paper's key argument is that existing AI research on personalized games can benefit from more player-centered perspectives. Despite their technical contributions, most existing work in this area is primarily geared towards more sophisticated algorithms and system capabilities. This system-centered approach to personalization has been attempted in early research of personalized web applications. It led to practices that tried "to find uses for the tools, and deploying the coolest new features" and made these applications less useful to the population they were supposed to serve [29]. To avoid similar drawbacks, the game AI research community can benefit from further aligning the technical research with deep models of player needs and behavior.

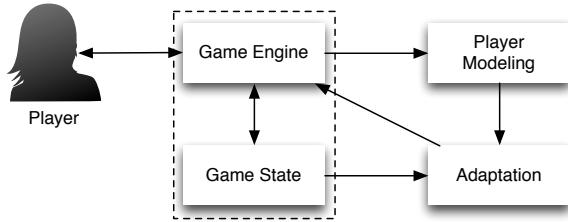


Figure 1: A typical system-centered view to an adaptive personalized game system.

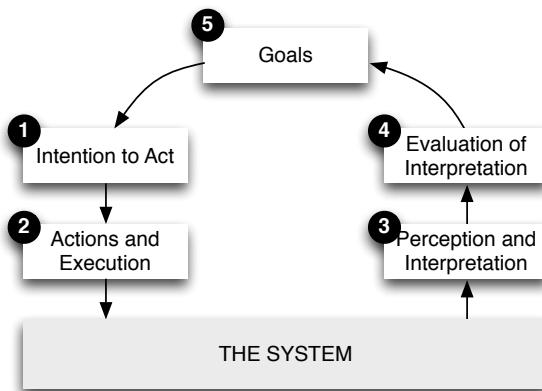


Figure 2: A condensed framework based on Norman's Structure of the Action cycle. [38]

In this paper, we present our initial work to strengthen the above alignment by mapping the current landscape of game AI research in personalization through a cognitive science theory. In particular, we use Norman's cognitive theory on Stages of Actions [38] to examine the state of the art research along each stage of what a player goes through and identify open problems for further investigation. Our goal is to present the current state of game AI research on adaptive personalized games that complement technology-centered overviews. The paper's central position is that to advance player-centered AI for game personalization, we need to combine technological advancement and player experience/game design.

There have been several overviews on the subject of game personalization. For example, Bakkes, Tan, and Pisan [4] offered a taxonomy of personalization in games research. More recently, Snodgrass et al. [55], combining a systematic survey on game design and AI algorithms on personalization, proposed a framework for personalization through the adaptation of the Player, Environment, Agents, and System. Existing related work analyzes personalization from the game developers' point of view (e.g., AI researchers, game designers). By contrast, this paper is among the first attempts to understand game personalization from the vantage view of players' cognitive processes.

2 THEORETICAL BACKGROUND

The idea of personalization originated as a reaction to mass production. In 1956, the same year as the Dartmouth summer school

of AI, Wendell R. Smith coined the term "market segmentation" to capture trends in marketing [54]. He challenged the long-held assumption of heterogeneity in both supply and demand sides of the market. Instead, Smith proposed to segment the market into smaller homogeneous ones to better capture the "diversity and variations in consumer demand." If we take Smith's idea to the extreme, each person becomes a complete segmentation — personalization.

With few exceptions, game AI publications on adaptive games typically represent adaptive games through system architecture diagrams similar to Figure 1. While this approach to thinking of adaptive games shed light on the connections between different AI components to achieve game personalization, it black-boxes players' needs and requirements. While it is necessary for technological advancements, it risks developing technology for its own sake.

This paper attempts to examine the current state of game personalization AI techniques through a novel player-centered lens. We use Don Norman's cognitive theory on the structure of actions [38, p.47], which is foundational work in the field of human-computer interaction. It describes the underlying cognitive processes a person goes through when taking action in the world into seven stages: goals, intention to act, actions, executions, perception, interpretation, and evaluation of interpretation. For this paper, we condense the process into five stages, combining actions with execution, and perception and interpretation. We also replace "the world" in Norman's original diagram with the focus of our analysis: Adaptive game systems (Figure 2. Notice Figure 2 is the complementary view of Figure 1.

Goals refer to something to be achieved, often vaguely stated, such as "find a movie to watch." A goal is then translated to an **intention to act** (marked 1 in Fig 2). The intention is more specific, such as "find a good movie on *Netflix*". A goal can be satisfied with different intentions and sequences of actions. In our example, the same goal could also lead to the intention of calling a friend for recommendations. The next stage is **actions and execution** (2) where the person specifies an action sequence and executes it — "log into *Netflix* and browse the *Top Picks* for me."

Once actions are taken in the system, the stage of **perception and interpretation** (3) allows the user to perceive the state of the system and interpret that state (e.g., the list of movies that *Netflix* recommends to me). Then the user can form her **evaluation of interpretation** (4) by asking the question, in our example, "is there something I like in this list?" If she determines the outcome of her actions has not satisfied her goal, the user may choose a different intention or modify her goal and go through the cycle again.

Stages (1) and (2) are about execution: what do we do to the world/system. If the user encounters difficulties in either stage, it creates *the gulf of execution* where the user does not know what to do (e.g., when the user cannot find *Top Picks* list). Stages (3) and (4) are about evaluation: comparing what happened with what we wanted to happen. Issues in either stage lead to *the gulf of evaluation* where the user cannot tell what happened or whether their goal was met (e.g., the system did not respond after pressing the scroll button in *Top Picks*).

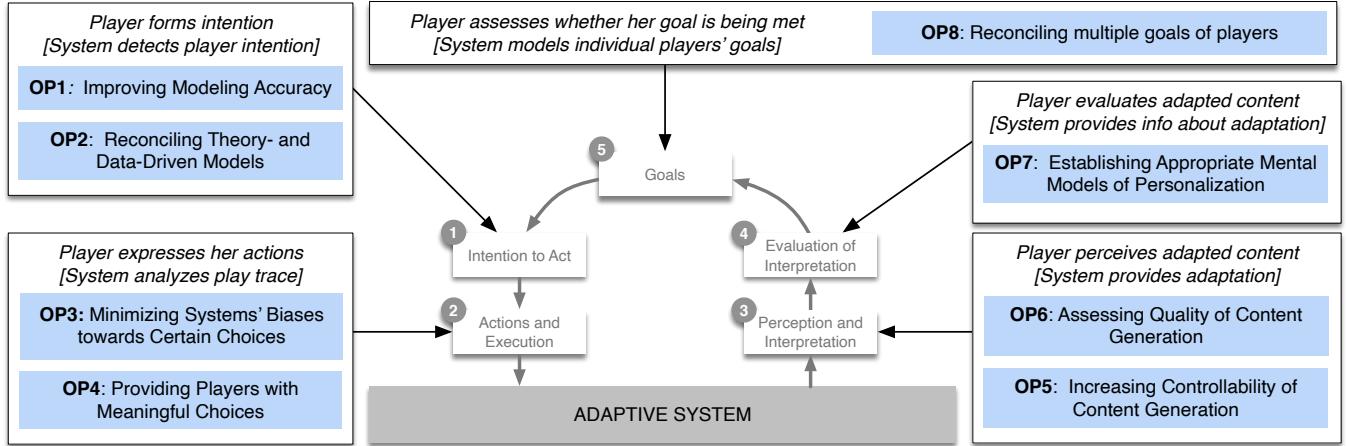


Figure 3: Eight open problems of current personalized game research in relation to the different stages of players' action cycle.

3 STATE OF THE ART AND OPEN PROBLEMS

According to Norman's theory, the stages of actions are the fundamental cognitive process each person goes through many times when interacting with any system, including adaptive personalized games. This framework sheds light on the gulf of execution and the gulf of evaluation, the underlying causes for the most common usability issues. It also differentiates the stages that are observable by an external spectator (Stages 2 & 3) from those that are latent and not directly observable (Stages 1, 4 & 5), highlighting critical challenges for adaptive game AI to model player intent, goals, and appraisal. Therefore, it provides a useful framework for understanding the current state of game personalization research and identifying critical open problems for developing player-centered approaches. Figure 3 shows how a player goes through the stages in the action cycle in an adaptive game and how the system functions accordingly. It also highlights the eight open problems discussed in this section.

3.1 Intention to Act

Currently, the gap between actual player intent and what an adaptive game “thinks” as player intent can be vast. Using algorithms to identify user intent is particularly challenging because the player’s intention is not directly observable and therefore has to be inferred based on observable features. Compared to Web applications with comparatively simple user actions, games often feature a more complex set of player choices and player experience. To demonstrate the difficulty of the problem, a player describes his experience of getting lost in one area of *Silent Hill: Shattered Memories*, an adaptive horror survival game. The game misunderstood it as an indication of his intent to choose to stay there and hence adapt the game to give the player more related content¹.

A direct approach to predict player intent is through player goal recognition. Researchers have used different machine techniques such as Bayesian models [1, 3, 60], Markov logic networks [19], and long short-term memory networks (LSTMs) [35] for plan, activity,

and intent recognition in computer games. These approaches have yielded reasonable results, given the complexity of the problem.

An indirect approach to approximate intent is player modeling. The underlying assumption is that if we can model certain relevant aspects of the player (e.g., gameplay preference), we can infer their individual intent and thus provide better personalization. Player modeling has been extensively studied. For example, work exists on modeling player types [5, 22], preferences [52, 62], player experience [41], skill level [26, 36, 77] or player behavior [12, 21, 23, 67], among other aspects. Despite the vast body of literature on this topic, two main open challenges still remain in player modeling.

3.1.1 Open Problem (OP 1): Improving Modeling Accuracy. While some of the work cited above shows it is possible to model general aspects of player groups with significant accuracy, modeling aspects of particular players is still a big challenge. One of the main problems is that to model an individual player, a considerable amount of data about that player would be needed. However, even if such data was available, labeling such data is a challenging problem. Even if the researchers could directly observe player behavior for an extended period of time, it is not always clear what the player’s intention is at any given time.

Some approaches, like collaborative filtering [44], commonly used to model user preferences for product recommendation, avoid the need for labeling data. Collaborative filtering makes predictions about the preferences of an individual by comparing her known preferences to those of other individuals. The assumption is that individuals with similar preferences about a product will likely have the same choices about others. However, the price they pay is not having models tuned to a specific individual, as preferences of an individual for a given product are assumed to be similar to other individuals who had similar preferences for other products.

3.1.2 OP 2: Reconciling Theory-Driven and Data-Driven Models. As reported several times [73], there is a tension between top-down (theory-driven) and bottom-up (data-driven) approaches. Some work on player modeling is based on existing theories, for example, those from psychology (e.g., Self Determination Theory [7, 47]).

¹Retrieved from https://www.reddit.com/r/silenthill/comments/3btew0/how_accurate_were_your_psych_profiles_for/ on January 31st, 2020.

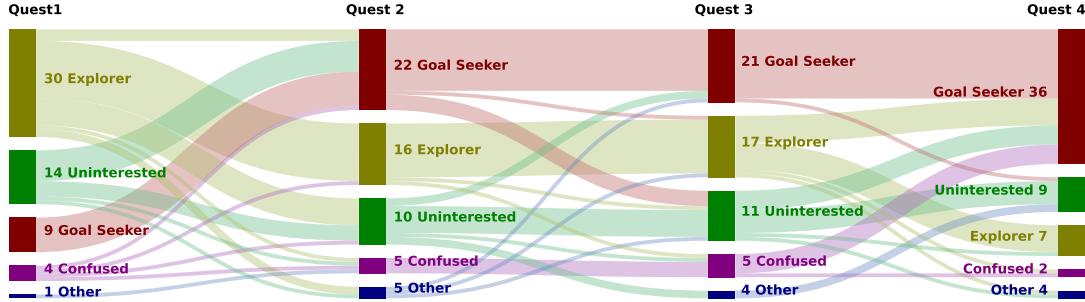


Figure 4: Alluvial diagram of how participants (n=58) shift their observed play styles between the four quests in the *Solving the Incognitum* game (reproduced with permission from [67]).

Others are directly based on the data itself (e.g., [67]). For a recent systematic review of data-driven player modeling, readers can refer to [24]. However, the research community has no consensus on what to do when those approaches disagree. For example, the learning science theory upon which the work of the TAEMILE project was based on assumed that the property being modeled was a trait, but the data showed that it shifted regularly (Figure 4). When this situation arises, it is unclear what is the right path forward: was the theory incorrect, or was the modeling approach used to interpret the data not accurate?

3.2 Actions and Executions

Since the AI infers a player's intention based on her observable actions, player modeling accuracy is thus directly tied to the range and quality of the choices available to the player. The quality of choices in this paper refers to the distance between the “true actions” a player would like to take and the set of actions at her disposal in the game. In most games, there is a gap between the “true actions” and what is offered by the game. This is partly because choices in games are expensive, especially when they involve art assets, animations, and voice acting. Game designers and researchers used creative ways such as false choices [10] and fold-back branching structure [9] to increase players' perceived agency without significantly growing game development costs. Another reason that shapes the options available to players is game designers' design intent for player experience.

Existing research on the quality of player choices mainly focuses on player agency and narrative immersion [15, 20, 37]. However, as a research field, we do not fully understand how to design game choices that can help the adaptive system more accurately predict player goals and other player characteristics.

3.2.1 OP 3: Minimizing Systems' Biases towards Certain Choices. The prerequisite of using player choices to infer their intention and preferences is that players can choose, relatively freely, options truly represent them. In other words, they need to be able to express their intent/preferences/other traits with little external influence. Compared to other interactive systems, designing player choices in games represents additional difficulties because of their built-in reward structure. As most players are motivated to win, they may choose options that help them win instead of options that better represent them. For instance, in our *Avian* project [40], an

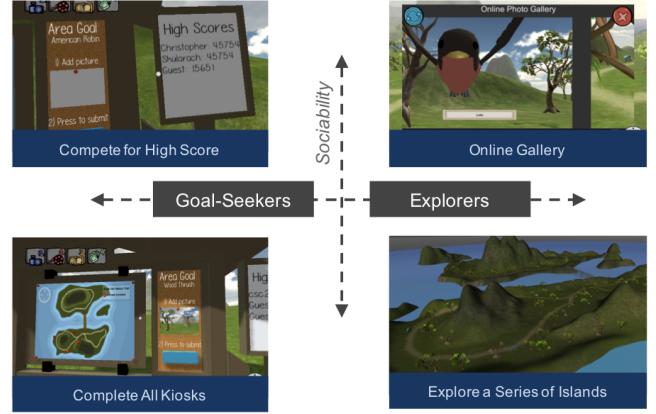


Figure 5: Reward mechanisms for each play style in *Avian* (reproduced with permission from [40]).

educational game on citizen science and bird watching, we intended to predict players' play style based on how they play the game. In the first version of the game, we found that most players chose to take actions associated with one playstyle. After looking into it further, we realized players selected these actions because that is how to win the game and where most gameplay content exists. With that revelation, we expanded the game design so that different types of players can find meaningful ways to satisfy their needs while progressing in the game (Figure 5). For explorers, for example, we added more in-games space to discover and an online gallery for those who are social. As a result, we observed a more balanced distribution of different player types.

This example further illustrates how the success of game AI is heavily tied to game design. When using theory-driven player modeling approaches, it is important to ensure that the different player attributes that the game AI is modeling are sufficiently supported in the player experience. However, when using data-driven approaches, it is not clear how this problem can be addressed, as there is no pre-existing set of categories to inspect during game design before data collection. Further research is needed on the design of choices for effective player modeling.

3.2.2 OP 4: Providing Players with Meaningful Choices. For a choice to be truly meaningful, what the player chooses should lead to

consequential changes. To do so requires a significant amount of game content and risks what is known as the *authorial bottleneck problem* [17] in the context of interactive stories.

Several recent research efforts aim to alleviate this authorial bottleneck. For example, the work of Valls et al. explored how to automatically extract domain knowledge from existing stories written in natural language and then use such domain knowledge to generate new stories [69], or even game worlds [66]. Another example is by Li et al. [33] to use crowd-sourcing to obtain plot graphs that can then create interactive experiences. A more recent example is the experimental game *AI Dungeon*², where a state of the art language model (GPT-2 [43]) is trained to generate text based on arbitrary user input in a MUD-style text-based game. Finally, in general, the field of procedural content generation (PCG) [50] can also have a significant impact on alleviating this authorial bottleneck.

On a related issue, game designers have long known the dilemma that people who lack choices seem to want them and often will fight for them. Yet at the same time, people with too many options intensely dislike it [48]. For adaptive games to effectively capture high-quality player choices that can be used for player modeling, we need more work to understand how to increase the efficiency of player choices and avoid decision fatigue.

3.3 Perception and Interpretation

Assume we have a perfect model of the player and the balanced options from which they can select. Now, the player selects her “true action”, how can now the game adapt in a way that the player can perceive?

Several areas of research have proposed different ideas that contribute to answering this question. For example, the narrative planning community has developed several planning-based approaches to adapt a story based on a player’s actions to maximum flexibility, while still ensuring the story advances in the direction desired by the designers. Riedl and Stern [46] presented the *Automated Story Director*, which can detect when the target story goal cannot be reached given the player actions, and re-plan accordingly to achieve the desired narrative effect. This is an instance of what is more generally called *drama management* or *experience management* [22, 34, 61, 72, 76]. Experience Management studies AI systems that automatically adapt interactive experiences such as computer games to serve specific users better and to fulfill specific design goals. For example, experience management techniques have been designed to adapt interactive game experiences to follow the desired story arc [72], or to adjust the difficulty of a game dynamically [25]. Experience management techniques have been explored for several decades. Although its effectiveness has been demonstrated in commercial deployment in games such as *Left 4 Dead 2*, current approaches still have limitations. For example, detecting if the current user is being served, or the goals of the experience manager have been achieved is very challenging, as this might involve automatically detecting, for instance, if the user is engaged, which is an open problem.

All these approaches need to balance *player autonomy* (the freedom of the player to behave as they prefer in the game), with

authorial intent (the goal that the game designer had in mind, and that is what the AI adapting the game is aiming for). Example authorial intents include trying to adapt a story to make it conform to a proper Aristotelian arc, or trying to make sure the player is exposed to a pre-specified set of concepts in educational games.

Another promising direction for algorithmic adaptation in games is procedural content generation (PCG). PCG work aimed at adaptive games includes automatic difficulty adjustment via PCG [26], generating levels specific to player style [51, 56, 63, 68], or the general idea of experience-driven PCG [74]. Currently, existing approaches have their limitations. For instance, in PCG, controllability and quality assessment are two key open challenges.

3.3.1 OP 5: Increasing Controllability of Content Generation. Although there has been some work in this direction (such as some reported above), the general problems of how to algorithmically specify the desired properties of the content we want a PCG algorithm to generate, and how to guide the algorithms to generate it is still unsolved. Additionally, how to exploit player models in general to create personalized content is still an open problem, even though some work exists [51, 56, 63, 64, 68]. More generally, how can we generate new adaptations of a game to achieve the desired effect without having to pre-author all of them. For example, in the work of Valls et al. [68], a graph-grammar approach is proposed to generate new puzzles for an educational game about learning parallel and concurrent programming called *Parallel*. In this approach, the set of desired concepts the player is supposed to practice (e.g., *mutual exclusion*) and the set of concepts the player is not supposed to have mastered yet is given to the PCG component, which can generate a level with those specifications by activating and deactivating rules in the grammar that create problem instances with specific concepts. Although reported results show promise, unanticipated rule interactions can still generate a level that contains an undesired concept.

3.3.2 OP 6: Assessing Quality of Content Generation. A second major problem in this area is how to assess the quality of the generated adaptations/content. Which are the metrics to evaluate this content? How do we avoid “catastrophic failures”? Some work exists in this direction [49, 57], trying to propose objective metrics to assess the quality of content, but to date, no automated metric has shown good enough correlation with human reported evaluations [58].

Using the same example of the *Parallel* game above, even if it is possible to automatically identify if a certain concept (e.g., *deadlock*) appears in a given puzzle, doing so is NP-complete, as it often requires exploring the set of all possible execution orders of the threads in the level. Additionally, even if we could detect the presence or not of a concept, it is unclear how to assess if the puzzle is a “good example” of the concept for pedagogical purposes.

3.4 Evaluation of Interpretation

For players to fully evaluate their personalized experience, they need to be able to form the appropriate mental models of whether and how the system adapts to them. In traditional user experience design literature, the doctrine of seamless design [6, 11, 38, 59] would advocate that players should not be aware that content has been personalized for them. After all, “a good tool is an invisible

²<https://aidungeon.io>

tool” [71]. However, this design guideline has been increasingly questioned in the context of adaptive personalized systems. Evidence starts to show that when users are not aware of the adaptive nature of the system, they often cannot utilize it fully. For instance, recent studies show that more than half of *Facebook* users are not aware that their news feed is personalized by algorithms [42]. Even for those who are aware, they can’t make sense of how exactly it works. As a consequence, these users build incorrect folk theories that are very different from how the algorithm actually works. This misconception can sometimes negatively impact real-world relationships [14]. Furthermore, when users are completely shielded from knowing how the personalization mechanism works, they lose the ability to inspect it. The phenomenon of filter bubbles [39] is an example of what could go wrong.

The issue of how and to what extent players should be able to evaluate their personalized gameplay experience has not received sufficient attention. Similarly, in an adaptive user interface, researchers found that when the users are aware that the personalized interface changed or why it changed, they simply ignored it [16]. Another example is the work of Lau [32], who analyzed a collection of AI-powered adaptive text editors (SMARTedit, Sheepdog, and CoScripter) and why they failed. Out of the five reasons uncovered, one of them was not having “a model users can understand”, highlighting one more time that helping the user create appropriate mental models of what an adaptive system is doing is vital for its success. Finally, another interesting work in this direction in the context of recommender systems is that of Ekstrand et al. [13], who allowed users to choose which recommendation algorithm to use, thus making them aware of the algorithms. This resulted in 25% of users using the feature.

3.4.1 OP 7: Establishing Appropriate Mental Models of Personalization. Despite the growing evidence that users need to develop the proper mental model to take full advantage of personalization, this topic has not received enough attention from the game research community. Most overview papers of player modeling [53, 73, 75] are written from an algorithmic perspective and do not include these player-centered aspects. Along with the rising interest in trust, transparency, and ethics of AI [18, 45], the topics such as how to establish the right mental model of personalization and how much information about the personalization mechanism should be revealed to players should be further studied.

3.5 Goals

Personalization aims to model the player in order to changes the player experience accordingly. Simultaneously, personalization is used to influence players and guide their choices and actions. This is also true outside of games, in most apps and websites we use every day. Consider, for example, Amazon and the products it shows you upon entering a search query. Amazon uses user modeling to identify products that might suit your profile, but at the same time, your choices are influenced by the personalization algorithms. Therefore, personalization also works in the other direction: it also changes the player to match a specific digital experience.

This feedback loop has been referred to as the *paradox of personalization* [28]. As a consequence of all previous gaps, and this

paradox, and as pointed out by Koponen [28], “personalization remains unfulfilling and incomplete. It leaves us with a feeling that it serves someone else’s interests better than our own.”

3.5.1 OP 8: Reconciling multiple goals of players. Humans are complex and have different co-existing goals and preferences. Most existing work simplifies the problem by focusing on modeling one aspect of player characteristics or player experience. As personalized adaptive games become more sophisticated, another open problem is to algorithmically reconcile multiple goals and preferences of an individual, some of which may be contradictory to one another.

An example of work towards addressing this open problem can be found on the TAEMILE [67] project mentioned above. TAEMILE is an interactive learning environment for earth science. The game attempts to model goal seeker vs. explorer learning behavior. However, in the context of learning, the player’s short-term learning preferences of goal-seeking or exploring often do not match her long-term goal of mastering the subject. This is because an effective pedagogical experience requires both. To cater to the player’s long-term needs, TAEMILE did something different from many similar projects: once a player model is built, instead of adapting the game to the player’s short-term learning behavior, it is used to nudge the player to try the opposite behavior. In a way, this project uses player modeling to identify the play style of a player, and then use the information to expand how she learns.

As multi-player games become popular among players, further research is also needed to personalize a shared gameplay experience between multiple people. One approach is to expand the framework of experience manager to incorporate models of multiple players [76].

4 CONCLUSIONS

In conclusion, computer games represent an ideal research domain for the next generation of personalized digital applications. To reach the full potential of personalized games, we argue that a player-centered approach for personalization AI is necessary. We hence presented a novel player-centered framework, complementary to the commonly used system-centered approaches, to synthesize existing AI research on personalized games. Built on the cognitive science theory on the structure of actions, we mapped current game personalization research based on the cognitive process a player goes through. As a result, we identified eight open problems that need further investigation. These problems require deep collaboration between AI researchers and game designers.

ACKNOWLEDGMENTS

This paper is partially based on Jichen Zhu’s invited talk at the 13th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE’17). This work is partially supported by the National Science Foundation under Grant Number IIS-1816470.

REFERENCES

[1] David W Albrecht, Ingrid Zukerman, and An E Nicholson. 1998. Bayesian models for keyhole plan recognition in an adventure game. *User modeling and user-adapted interaction* 8, 1-2 (1998), 5–47.

[2] Matt Anderson. 2018. The Internet of Me: When the Consumer Becomes the Electronics. *Wired Magazine* (2018). <https://www.wired.com/brandlab/2018/01/internet-consumer-becomes-electronics/>

[3] Alok Baikadi, Jonathan Rowe, Bradford Mott, and James Lester. 2014. Generalizability of goal recognition models in narrative-centered learning environments. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 278–289.

[4] Sander Bakkes, Chek Tien Tan, and Yusuf Pisan. 2012. Personalised gaming: a motivation and overview of literature. In *Proceedings of The 8th Australasian Conference on Interactive Entertainment: Playing the System*. 1–10.

[5] Richard Bartle. 1996. Hearts, clubs, diamonds, spades: Players who suit MUDs. *Journal of MUD research* 1, 1 (1996), 19.

[6] Jay David Bolter and Diane Gromala. 2003. *Windows and mirrors: Interaction design, digital art, and the myth of transparency*. MIT press.

[7] Patrice Bouvier, Karim Séhaba, and Élise Lavoué. 2014. A trace-based approach to identifying users' engagement and qualifying their engaged-behaviours in interactive systems: application to a social game. *User Modeling and User-Adapted Interaction* 24, 5 (2014), 413–451.

[8] Elizabeth F. Churchill. 2013. Putting the person back into personalization. *Interactions* 20, 5 (2013), 12–15.

[9] Chris Crawford. 2004. *Chris Crawford on Interactive Storytelling*. New Riders, Berkeley, CA.

[10] Timothy Day and Jichen Zhu. 2017. Agency informing techniques: communicating player agency in interactive narratives. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*. 1–4.

[11] Alan Dix, Alan John Dix, Janet Finlay, Gregory D Abowd, and Russell Beale. 2003. *Human-computer interaction*. Pearson Education.

[12] Anders Drachen, Alessandro Canossa, and Georgios N Yannakakis. 2009. Player modeling using self-organization in Tomb Raider: Underworld. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*. IEEE, 1–8.

[13] Michael D Ekstrand, Daniel Kluver, F Maxwell Harper, and Joseph A Konstan. 2015. Letting users choose recommender algorithms: An experimental study. In *Proceedings of the 9th ACM Conference on Recommender Systems*. 11–18.

[14] Motaahare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I always assumed that I wasn't really that close to [her]" Reasoning about Invisible Algorithms in News Feeds. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 153–162.

[15] Matthew William Fendt, Brent Harrison, Stephen G Ware, Rogelio E Cardona-Rivera, and David L Roberts. 2012. Achieving the illusion of agency. In *International Conference on Interactive Digital Storytelling*. Springer, 114–125.

[16] Anushay Furqan, Chelsea Myers, and Jichen Zhu. 2017. Learnability through AdaptiveDiscovery Tools in Voice UserInterfaces. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI 2017)*. Denver, CO, forthcoming.

[17] Pablo Gervás. 2009. Computational approaches to storytelling and creativity. *AI Magazine* 30, 3 (2009), 49–49.

[18] Bryce Goodman and Seth Flaxman. 2017. European Union regulations on algorithmic decision-making and a right to explanation. *AI magazine* 38, 3 (2017), 50–57.

[19] Eun Young Ha, Jonathan P Rowe, Bradford W Mott, and James C Lester. 2011. Goal recognition with Markov logic networks for player-adaptive games. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

[20] D Fox Harrell and Jichen Zhu. 2009. Agency Play: Dimensions of Agency for Interactive Narrative Design. In *Proceedings of AAAI 2009 Spring Symposium on Intelligent Narrative Technologies II*. AAAI Press, Stanford, CA, 44–52.

[21] Brent Harrison and David L Roberts. 2011. Using sequential observations to model and predict player behavior. In *Proceedings of the 6th International Conference on Foundations of Digital Games*. ACM, 91–98.

[22] Carrie Heeter, Yu-Hao Lee, Ben Medller, and Brian Magerko. 2011. Beyond player types: gaming achievement goal. In *ACM SIGGRAPH 2011 Game Papers*. ACM, 7.

[23] Christoffer Holmgård, Antonios Liapis, Julian Togelius, and Georgios N Yannakakis. 2014. Evolving personas for player decision modeling. In *2014 IEEE Conference on Computational Intelligence and Games*. IEEE, 1–8.

[24] Danial Hooshyar, Moslem Yousefi, and Heuiseok Lim. 2018. Data-driven approaches to game player modeling: a systematic literature review. *ACM Computing Surveys (CSUR)* 50, 6 (2018), 1–19.

[25] Robin Hunicke. 2005. The case for dynamic difficulty adjustment in games. In *2005 ACM SIGCHI International Conference on Advances in computer entertainment technology*. ACM, 429–433.

[26] Martin Jennings-Teats, Gillian Smith, and Noah Wardrip-Fruin. 2010. Polymorph: dynamic difficulty adjustment through level generation. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*. ACM, 11.

[27] Pavan Kantharaju, Katelyn Alderfer, Jichen Zhu, Bruce Char, Brian Smith, and Santiago Ontañón. 2018. Tracing player knowledge in a parallel programming educational game. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[28] Jarno M. Koponen. 2015. The Future Of Algorithmic Personalization. *TechCrunch* (2015). <https://techcrunch.com/2015/06/25/the-future-of-algorithmic-personalization>

[29] Joseph Kramer, Sunil Noronha, and John Vergo. 2000. A User-Centered Design Approach to Personalization. *Commun. ACM* 43, 8 (2000), 44–48.

[30] Ashok Kumar. 2007. From mass customization to mass personalization: a strategic transformation. *International Journal of Flexible Manufacturing Systems* 19, 4 (2007), 533.

[31] Frederic Lardinois. 2013. The Next Frontier For Google Maps Is Personalization. *TechCrunch* (2013). <https://techcrunch.com/2013/02/01/the-next-frontier-for-google-maps-is-personalization/>

[32] Tessa Lau. 2009. Why programming-by-demonstration systems fail: Lessons learned for usable ai. *AI Magazine* 30, 4 (2009), 65–65.

[33] Boyang Li, Stephen Lee-Urban, George Johnston, and Mark Riedl. 2013. Story generation with crowdsourced plot graphs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.

[34] Michael Mateas and Andrew Stern. 2002. A behavior language for story-based believable agents. *IEEE Intelligent Systems* 17, 4 (2002), 39–47.

[35] Wookhee Min, Bradford Mott, Jonathan Rowe, Barry Liu, and James Lester. 2016. Player goal recognition in open-world digital games with long short-term memory networks. *IJCAI International Joint Conference on Artificial Intelligence 2016-January* (2016), 2590–2596.

[36] Olana Missura and Thomas Gärtner. 2009. Player modeling for intelligent difficulty adjustment. In *International Conference on Discovery Science*. Springer, 197–211.

[37] Janet H Murray. 1997. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. Simon & Schuster/Free Press, New York.

[38] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Basic books.

[39] Eli Pariser. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.

[40] Justin Patterson. 2014. *Avian: Game Design and Player Metrics for Player Modeling in Educational Games*. Ph.D. Dissertation, Drexel University.

[41] Chris Pedersen, Julian Togelius, and Georgios N Yannakakis. 2009. Modeling player experience in super mario bros. In *2009 IEEE Symposium on Computational Intelligence and Games*. IEEE, 132–139.

[42] Emilee Rader and Rebecca Gray. 2015. Understanding user beliefs about algorithmic curation in the Facebook news feed. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 173–182.

[43] Alex Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019), 9.

[44] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. 1994. GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*. 175–186.

[45] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.

[46] Mark O Riedl and Andrew Stern. 2006. Believable agents and intelligent story adaptation for interactive storytelling. In *International Conference on Technologies for Interactive Digital Storytelling and Entertainment*. Springer, 1–12.

[47] Robert Sawyer, Jonathan Rowe, Roger Azevedo, and James Lester. 2018. Modeling Player Engagement with Bayesian Hierarchical Models. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[48] Jesse Schell. 2008. *The Art of Game Design: A book of lenses*. AK Peters/CRC Press.

[49] Noor Shaker, Gillian Smith, and Georgios N Yannakakis. 2016. Evaluating content generators. In *Procedural Content Generation in Games*. Springer, 215–224.

[50] Noor Shaker, Julian Togelius, and Mark J Nelson. 2016. *Procedural content generation in games*. Springer.

[51] Noor Shaker, Georgios Yannakakis, and Julian Togelius. 2010. Towards automatic personalized content generation for platform games. In *Sixth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[52] Manu Sharma, Santiago Ontañón, Manish Mehta, and Ashwin Ram. 2010. Drama management and player modeling for interactive fiction games. *Computational Intelligence* 26, 2 (2010), 183–211.

[53] Adam M Smith, Chris Lewis, Kenneth Hullett, Gillian Smith, and Anne Sullivan. 2011. An inclusive taxonomy of player modeling. *University of California, Santa Cruz, Tech. Rep. UCSC-SOE-11-13* (2011).

[54] Wendell R Smith. 1956. Product differentiation and market segmentation as alternative marketing strategies. *Journal of marketing* 21, 1 (1956), 3–8.

[55] Sam Snodgrass, Omid Mohaddesi, Jack Hart, Guillermo Romero Rodriguez, Christoffer Holmgård, and Casper Harteveld. 2019. Like PEAS in PoDS: the player, environment, agents, system framework for the personalization of digital systems. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*. 1–15.

[56] Sam Snodgrass and Santiago Ontañón. 2016. Controllable Procedural Content Generation via Constrained Multi-Dimensional Markov Chain Sampling.. In *IJCAI*. 780–786.

[57] Adam Summerville. 2018. Expanding expressive range: Evaluation methodologies for procedural content generation. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[58] Adam Summerville, Julian RH Mariño, Sam Snodgrass, Santiago Ontañón, and Levi HS Lelis. 2017. Understanding mario: an evaluation of design metrics for platformers. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*. 1–10.

[59] Kirsten Sweareigen and Rashmi Sinha. 2002. Interaction design for recommender systems. In *Designing Interactive Systems*, Vol. 6. 312–334.

[60] Gabriel Synnaeve and Pierre Bessiere. 2011. A Bayesian model for plan recognition in RTS games applied to StarCraft. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

[61] David Thue and Vadim Bulitko. 2018. Toward a Unified Understanding of Experience Management. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[62] David Thue, Vadim Bulitko, and Marcia Spetch. 2008. PaSSAGE: A Demonstration of Player Modeling in Interactive Storytelling.. In *AIIDE*.

[63] Julian Togelius, Renzo De Nardi, and Simon M Lucas. 2006. Making racing fun through player modeling and track evolution. (2006).

[64] Julian Togelius, Mike Preuss, Nicola Beume, Simon Wessing, Johan Hagelbäck, Georgios N Yannakakis, and Corrado Grappiolo. 2013. Controllable procedural map generation via multiobjective evolution. *Genetic Programming and Evolvable Machines* 14, 2 (2013), 245–277.

[65] MM Tseng, RJ Jiao, and C Wang. 2010. Design for mass personalization. *CIRP annals* 59, 1 (2010), 175–178.

[66] Josep Valls-Vargas, Santiago Ontañón, and Jichen Zhu. 2013. Towards story-based content generation: From plot-points to maps. In *2013 IEEE Conference on Computational Intelligence in Games (CIG)*. IEEE, 1–8.

[67] Josep Valls-Vargas, Santiago Ontañón, and Jichen Zhu. 2015. Exploring player trace segmentation for dynamic play style prediction. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

[68] Josep Valls-Vargas, Jichen Zhu, and Santiago Ontañón. 2017. Graph grammar-based controllable generation of puzzles for a learning game about parallel programming. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*. 1–10.

[69] Josep Valls-Vargas, Jichen Zhu, and Santiago Ontañón. 2017. Towards End-to-End Natural Language Story Generation Systems. In *Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[70] Pengcheng Wang, Jonathan P Rowe, Wookhee Min, Bradford W Mott, and James C Lester. 2017. Interactive Narrative Personalization with Deep Reinforcement Learning.. In *IJCAI*. 3852–3858.

[71] Marc Weiser. 1994. The world is not a desktop. *interactions* 1, 1 (1994), 7–8.

[72] Peter William Weyhrauch. 1997. *Guiding interactive drama*. Carnegie Mellon University.

[73] Georgios N Yannakakis, Pieter Spronck, Daniele Loiacono, and Elisabeth André. 2013. Player modeling. (2013).

[74] Georgios N Yannakakis and Julian Togelius. 2011. Experience-driven procedural content generation. *IEEE Transactions on Affective Computing* 2, 3 (2011), 147–161.

[75] Georgios N Yannakakis and Julian Togelius. 2018. *Artificial intelligence and games*. Vol. 2. Springer.

[76] Jichen Zhu and Santiago Ontañón. 2019. Experience management in multi-player games. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–6.

[77] Alexander E Zook and Mark O Riedl. 2012. A temporal data-driven player model for dynamic difficulty adjustment. In *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*.