



Improving Fairness in Adaptive Social Exergames via Shapley Bandits

Robert C. Gray
Department of Digital Media
Drexel University
Philadelphia, PA, USA
robert.c.gray@drexel.edu

Diane H. Dallal
Department of Psychiatry
University of Pennsylvania
Philadelphia, PA, USA
dianehdallal@gmail.com

Jennifer Villareale
Department of Digital Media
Drexel University
Philadelphia, PA, USA
jmv85@drexel.edu

Santiago Ontañón
Google Research
Philadelphia, PA, USA
santi.ontanon@gmail.com

Thomas B. Fox
Department of Digital Media
Drexel University
Philadelphia, PA, USA
seafordtbf@gmail.com

Danielle Arigo
Department of Psychology
Rowan University
Glassboro, NJ, USA
arigo@rowan.edu

Shahin Jabbari
Department of Computer Science
Drexel University
Philadelphia, PA, USA
shahin@drexel.edu

Jichen Zhu
Department of Digital Design
IT University of Copenhagen
Copenhagen, Denmark
jichen.zhu@gmail.com

ABSTRACT

Algorithmic fairness is an essential requirement as AI becomes integrated in society. In the case of social applications where AI distributes resources, algorithms often must make decisions that will benefit a subset of users, sometimes repeatedly or exclusively, while attempting to maximize specific outcomes. How should we design such systems to serve users more fairly? This paper explores this question in the case where a group of users works toward a shared goal in a social exergame called *Step Heroes*. We identify adverse outcomes in traditional multi-armed bandits (MABs) and formalize the *Greedy Bandit Problem*. We then propose a solution based on a new type of fairness-aware multi-armed bandit, *Shapley Bandits*. It uses the Shapley Value for increasing overall player participation and intervention adherence rather than the maximization of total group output, which is traditionally achieved by favoring only high-performing participants. We evaluate our approach via a user study ($n=46$). Our results indicate that our Shapley Bandits effectively mediates the Greedy Bandit Problem and achieves better user retention and motivation across the participants.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Social and professional topics** → **Computing / technology policy**; • **Computing methodologies** → *Artificial intelligence*.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
IUI '23, March 27–31, 2023, Sydney, NSW, Australia
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0106-1/23/03.
<https://doi.org/10.1145/3581641.3584050>

KEYWORDS

multi-armed bandits, algorithmic fairness, shapley value, adaptive games, social exergames

ACM Reference Format:

Robert C. Gray, Jennifer Villareale, Thomas B. Fox, Diane H. Dallal, Santiago Ontañón, Danielle Arigo, Shahin Jabbari, and Jichen Zhu. 2023. Improving Fairness in Adaptive Social Exergames via Shapley Bandits. In *28th International Conference on Intelligent User Interfaces (IUI '23)*, March 27–31, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3581641.3584050>

1 INTRODUCTION

Many organizations are deploying AI systems to manage and allocate resources, such as assigning work tasks and distributing work schedules to personnel to improve productivity in dynamic environments [53, 73, 76]. However, these opportunities have also raised new challenges, most notably the potential for AI systems to unfairly allocate opportunities and resources among a group of users. In the case of social applications where users work together toward a shared goal, this challenge is particularly difficult; individual user contributions to the group may vary, and the AI may place a subset of users above others to maximize outcomes without considering user contributions. Hence, fairness has become an emerging and important research area within the AI and machine learning communities [13, 16]. Work in this area identified fairness problems in a wide range of AI techniques such as classification [48, 105], regression [15], online learning [54, 63], or reinforcement learning [52]. Despite a growing body of work, we do not yet have good solutions for fairness-aware AI systems that need to distribute resources in social applications where groups of users work together towards a joint goal [109, 110], such as social exergames.

In this paper, based on prior work in an AI-based adaptive social fitness application [83, 103, 106], we found that commonly used metrics to capture individual users' contributions to the group were

insufficient due to fairness issues. We noticed an AI tasked with adapting the experience to maximize users' physical activity (PA) repeatedly placed some users' preferences above others, leading to non-adherence by marginalized users. Rather than focusing purely on the amount of PA (e.g., daily steps) achieved by the group, this paper examines the potential benefits of a fairness-aware approach targeting player adherence to the group activity as an alternative metric for the AI to maximize. Specifically, we present a new type of fairness-aware multi-armed bandit approach that we call the *Shapley Bandit*, as it uses the Shapley Value [92] to enforce fairness constraints. With this approach, we explore how to design AI for social applications that can serve a group of users while ensuring fundamental expectations of "distributive fairness" are maintained [5, 30]. Compared to prior work in fairness in online learning and, in particular, multi-armed bandit problems [54, 60, 63], which assume that the choices made by the AI only affect one of the individuals, our work studies the setting where the choices made by the AI can affect more than one individual.

We conducted a mix-methods user study to evaluate the effectiveness and tradeoffs of our Shapley Bandits compared to traditional bandits in the social exergame *Step Heroes*. Our user study results ($n=46$) show that our Shapley Bandits achieve higher motivation scores for physical activity than a control group, whereas a traditional bandit approach did not. Additionally, it is better at retaining users, where users are more likely to maintain adherence to the intervention. Based on our findings, we propose implications for *behavioral fairness* in AI-driven social applications in terms of considering users' contributions to the group as a factor for fairness.

Our primary contribution is a notion of fairness integrated into an AI strategy that makes its selections based on the disparity between a user's exerted effort within a team endeavor and the degree of favorable treatment they receive from the AI in a multiplayer environment. Specifically, this work brings the following contributions:

- To the best of our knowledge, this is the first attempt in the game AI research community to model group preferences and the related fairness issues.
- We present a new fairness-aware multi-armed bandit strategy, the *Shapley Bandit*, that uses the Shapley Value to estimate a notion of *disparity* and then aims to reduce it while at the same time maximizing expected reward.
- We present empirical evidence of the effectiveness of our approach in the context of exergames via a user study ($n=46$).

Our research indicates the importance of integrating fairness concerns in multi-armed bandits literature, especially in the context of adaptive social exergames that serve multiple players. Our empirical results show that if the AI only focuses on total group output (e.g., a group's total daily steps), it risks consistently favoring only high-performing participants while ignoring the needs of marginalized users. Over time, this traditional approach may lead to the latter dropping out, thus raising concerns about the ethics as well as the effectiveness of the intervention. We observed that even though introducing fairness considerations does not directly lead to higher PA levels, it contributes to user retention and adherence, which is critical to lasting behavior change. Therefore, holistic and

ethical considerations of the AI's utility function are key design issues in such applications and need further research.

2 RELATED WORK

2.1 Fairness in AI

With the rapid adaptation of AI systems in domains such as financial lending, sentencing, and healthcare, there have been growing concerns about societal aspects such as bias and transparency when deploying these systems. These concerns have led to an emerging research area within the AI and machine learning communities known as algorithmic fairness (see [13, 16] for recent surveys). The research in this field typically considers well-studied learning settings such as classification [48, 105], regression [15], online learning [54, 63], or reinforcement learning [52]. This research examines how bias or discrimination can happen in these settings, proposes formal definitions for what bias or unfairness means mathematically, and designs provably fair algorithmic frameworks. The main focus of the field has been on group fairness, where individuals are divided into groups using a combination of sensitive attributes such as race or gender and where fairness guarantees are designed to apply at the group level [15, 31, 48, 105]. However, group fairness does not provide any meaningful guarantee at the individual level, still allowing unfair treatment of individuals. Several works have studied individual fairness [36, 54], and it is known that individual fairness is often hard to enforce in practice [36].

Additionally, past research has shown that there is a disconnect between the needs and realities of high-stake applications (like algorithmically informed public decision-making in taxation, justice, or child protection) and tools and research in algorithmic fairness [98]. This remains true, despite recent efforts in this direction. For example, Richardson et al. studied the gap between the needs of practitioners and the tools offered by fairness research [86], and Nakao et al. studied how to directly engage end-users in identifying and addressing fairness issues in AI applications [75].

Our work studies individual fairness in online learning and, in particular, multi-armed bandit problems. Prior work [54, 60, 63] assumes that the choice of action in each round only affects one of the individuals favorably. Our work differs from prior work in that the choice of action can affect more than one individual either favorably or unfavorably.

2.2 Group Fairness in Resource Allocation

In applications where an AI adapts an experience for a group of human users, a body of work has been developed on using AI techniques to improve productivity through managing and allocating resources [11, 29, 43, 88, 93, 95]. In particular, an important research area is developing AI systems to mediate resources among a group of people, such as assigning workers to tasks [10, 76], allocating work schedules [35, 53, 73], and matching project requirements with employees' skill sets [76]. These decisions are usually affected by a set of *externally* defined objectives and constraints, such as meeting deadlines and maximizing resource utilization [10], hours an operator can work [29], and human factors, such as expertise and skill level [76].

Recently, researchers have discussed fairness-related harms that can occur, such as people’s individual experiences with AI systems [88] or how AI systems represent individuals in groups [4]. Of note, fairness in group AI (i.e., AI managing groups of people by distributing resources) is a particularly difficult challenge, as AI systems can unfairly allocate opportunities, resources, or information [17]. For example, an automated scheduling system may inadvertently withhold jobs from a particular individual due to how they are labeled or represented in the system.

Researchers have focused on fairness by mitigating disparities in the treatment of individuals based on categories such as racial, gender, or age groups [27, 40, 50]. In the case of this research, in which we provide a game intervention toward motivating people to walk more in their daily lives, targeting fairness based on such categories lies outside the scope of our domain and the visibility of our AI. Instead, we aim to address the fundamental issue of “distribution fairness” [5, 30], which considers users’ utility to the group as a factor for fairness. We address distribution fairness through the lens of Organizational Justice Theory [46, 47, 88] extending from J.S. Adams’ equity theory [1, 2]. The primary assertion of these theories is that there exists an implied contract that individuals participating in a group should receive a distribution of reward commensurate with their personal investment toward the goals of that group, with an emphasis on the negative responses of individuals that are likely to result from a failure to uphold this tenet.

2.3 Fairness in Group Recommender Systems

A closely related area to our research is group recommender systems. A group recommender system makes a decision, often by suggesting items to a group of people engaged in a *group activity* [57], such as watching a movie together. They have been used to recommend music [26, 33, 69], movies and TV [78, 104], and restaurant and travel [68, 90] to a group of users collectively. In order to capture the preference of the group as a whole, group recommender systems use group aggregation models to integrate the preferences of individual members [34]. Fairness is an essential consideration of group recommender systems because not all recommended items satisfy the group members equally. In other words, a group aggregation model decides whose preferences or needs should be given more weight at a given time.

Group aggregation models can be divided into three main strategies: 1) majority-based, 2) consensus-based, and 3) borderline strategies [56, 66, 91]. Generally speaking, majority-based strategies select the item preferred by most group members. Consensus-based strategies aggregate the preferences of all members somehow and then select items based on the common opinion of the group. Widely used models in this type include the average strategy, least misery strategy, Borda count, most pleasure strategy, average without misery strategy, and fairness strategy. Borderline strategies are in-between the previous two types. An example is the “dictatorship” strategy, where items are selected based on one selected member’s preferences.

Most of the above approaches assume that user preferences do not interact with each other. For example, at a party where people listen to the same music, what a person would like to hear next can be influenced by the songs recommended earlier. In this paper, we

explore a different way of aggregating user preferences or needs by minimizing the disparity between how much an individual contributes to the group (estimated via the Shapley Value) and the degree to which the AI caters to that player with its recommendations.

2.4 Personalized Digital Interventions for PA and Social Comparison

There is a broad base of existing literature on personalized interventions to promote physical activity (PA), especially to adapt PA goals for improving goal adherence [7, 18, 20, 81, 84]. Among existing work on PA personalization, an emergent approach is to use AI techniques to personalize digital interventions automatically [39, 80, 103]. A benefit of AI-based interventions is the ability to provide real-time tailoring of feedback and adaptation of goals based on continuous data monitoring on an individual basis.

Most notably, personalized digital interventions have shown to be more beneficial for behavior change over generic automated systems [83], such as helping individuals develop attainable step goals through more personalized recommendations [83, 103, 106]. However, relatively little work has leveraged AI systems to enhance PA for groups as a whole. Instead, personalization has prioritized the content and timing of individual activity recommendations using demographic information [21], behavioral patterns [62], time and location [21, 96], and personality [9].

Social comparison describes the process by which individuals evaluate themselves or their behavior relative to others [38]. In digital group interventions, applications typically use social features such as ranked leaderboards and competitive challenges to engage *social comparison* processes [102]. While these interventions have motivated health behavior change in weight control and PA promotion [59, 77], not all users may be equally motivated by these features. Research has shown that individuals prefer different people to compare with [14], and these preferences may vary over short periods [7]. Comparisons to others who seem “better off” than the comparer in a given domain are *upward comparisons*, and comparisons to others who seem “worse off” are *downward comparisons* [22]. Unfortunately, most personalized group interventions do not account for individual differences and provide all users with the same kind of generic exposure to social comparison targets [7], which can negate any need for behavior change efforts [8, 71, 100].

An alternative approach is to motivate PA through social gamified fitness applications or “social exergames” using similar group features described above. This active research area integrates physical exertion with digital games [65] to transform physical exercise in an engaging game environment and leverage social interactions through facilitating cooperative and competitive game settings to increase and sustain players’ motivation for PA [25, 28, 61]. Social interaction through the lens of cooperation and competition has been a clear motivator for PA activities in these games and has been shown to encourage comparisons between players [3, 28, 61]. However, social features such as ranked leaderboards and competitive challenges [7, 28] have been shown to also risk demotivating people due to how they *compare* themselves to others. This paper extends the literature on personalization in digital interventions for PA, especially social exergames, by developing an AI approach



Figure 1: Left: Screenshot of the study profile selection screen in our preliminary study. Right: Illustration of the decision the AI agent makes. The MAB-based AI determines the placement of the third teammate’s steps among three options (A, B, or C) to cater to the users’ social comparison orientation (SCO) and increase user motivation.

that automatically models the social comparison preferences of the group and provides appropriate social comparison opportunities.

3 RESEARCH PLATFORM

The following introduces our preliminary research platform (Section 3.1) that helped us identify the fairness issues that arise when optimizing for group performance (Section 3.2). We then introduce the research platform for our primary user study (Section 3.3) and the AI approach we designed to address fairness (Section 4).

3.1 Preliminary Design: Web-Based App with Standard MAB

Since research has established that social elements are a crucial motivating factor for PA [24, 37, 61, 67, 108], our preliminary work aims to promote PA using adaptive personalized AI inspired by *social comparison* theory (see Section 2.4). We designed a web-based app in which a group of users can compare their daily steps with other users’ PA-related profiles (e.g., daily steps, favorite types of exercises, and other PA-related interests) [107]. Users’ steps were captured by Fitbit and synced automatically with our platform. Every day, the user selects one out of two other users that the MAB recommends to them as (social) comparison targets (Figure 1 left). Each group consists of two randomly paired human users and one artificial user whose steps are controlled by the AI, similar to our previous work [45, 107]. The artificial user gives the MAB an additional chance to provide good comparison targets, where human users’ daily steps are unpredictable and thus may not provide the motivating comparison targets the two humans in the group need. Based on the changes in the users’ daily steps and their reported motivation for PA, the MAB updates its model for their real-time preference of social comparison direction and reaction to the comparison. It then recommends comparison targets to maximize the chance of increasing PA and motivation.

We conducted a 21-day user study ($n=53$) to investigate our personalization mechanism and how exposure to personalized social comparison targets may affect users’ PA and motivation to exercise. Our initial results indicated that our approach was able to automatically model and manipulate social comparison in the pursuit of PA

promotion. The detected effects achieved small-to-moderate effect sizes, illustrating the real-world implications of the intervention for enhancing motivation and PA [107].

3.2 Problem Formulation: The Greedy Bandit Problem

Further analysis of the results in our preliminary design raised questions regarding how the MAB distributes its “attention” to different users. The MAB selects an arm (in the form of an artificial user as the comparison target for two human users) to maximize the sum of steps and self-reported motivation of both users. We note that, in essence, our MAB operates similarly to the average strategy for group recommender systems (Section 2.3).

Let us consider Figure 1 (right) and imagine a case in which the model predicts Arm A will yield a highly positive effect on User 1 due to their preference for upward comparisons. However, in this example, the AI also predicts a negative effect on User 2 due to User 2’s dislike for upward comparisons. It may be that the *net* prediction of Arm A still yields the highest overall effect (i.e., the strength of User 1’s preference may be larger than that of User 2), so our traditional bandit strategy optimizing for the sum of steps of both users will select it. When faced with the same decision again later, the MAB is likely to make the same decision. Over repeated selections, a blind favoring of User 1 in a “greedy” pursuit of a specific metric may lead to overlooking User 2 regularly.

In a larger group, a similar case may arise when a particular arm appeals to the majority of users (e.g., 60% of users prefer one game setting that the other 40% dislike). This is similar to the majority-based group aggregation models in group recommender systems (Section 2.3). Over time, a greedy strategy may risk repeatedly neglecting the minority, leading to dissatisfaction among the users.

Disparity. Further analysis of our user study data indeed shows evidence that users whose social comparison preferences are not met by the MAB have a higher chance of missing sessions. In Figure 2, the blue bars show the *disparity* between each user’s effort and treatment, defined as follows. We define the *effort* E_i of a user i as their average daily total steps. We define a the *treatment* T_i of a user i as the total number of days where they received the best




Figure 2: Sorted disparity scores (blue) and miss likelihood (orange, secondary axis) for users in the preliminary design study. Users sorted to the left depict those who were given very good treatment relative to their effort (i.e., step performance), while users sorted to the right received poor treatment relative to their efforts. The graph reveals a positive correlation between a user’s disparity score and their number of missed sessions (Pearson’s $R=0.16$).

social comparison targets according to the MAB estimations. We now define $PR(E_i)$ as the percentile where user i falls in terms of effort across all users in the study (e.g., 1.0 if they are the one with the most steps, and 0.0 if they are the one with the least steps) and $PR(T_i)$ as the percentile where the user i falls in terms of treatment. The disparity of a user i is then defined as $D_i = (PR(E_i) - PR(T_i))$. This calculation results in a number between potential extremes of -1 (i.e., highest treatment with lowest effort) and +1 (i.e., lowest treatment with highest effort). Figure 2 sorts the users from least disparity to the most. The orange plot shows the corresponding user’s total number of missed sessions (days).

Our results show a positive correlation (Pearson’s $R=0.16$), where users who receive lower treatment relative to their effort percentile are more likely to miss sessions. In other words, high *disparity* is correlated with non-adherence behavior. When the AI does not make choices catering to a user’s preferences to a degree commensurate to that user’s investment of effort within the game to support their team, we see that those users are more likely to miss sessions. In other words, they are more likely to divest from the activity and the intervention it provides.

We argue that this result reveals two problems. The first is that when a user ceases to participate in the application in which the AI agent operates, any intervention attempted by the AI agent becomes completely ineffective. The second involves a limited notion of “fairness” derived from the example above. Although a user or group of users may be in the preferential minority, they still bring value to the game environment and therefore contribute to the overall game experience and community. Indeed, even a user’s mere presence or participation enables the multiuser environment to exist for others. Therefore, we would like to promote a best practice in which an AI agent operating in a group setting recognizes this contribution provided by each user and ensures it is rewarded.

We refer to this problem as the *Greedy Bandit Problem*. Our main technical contribution, the *Shapley Bandit* (Section 4), is designed to directly target this problem via the following approach: (1) using the Shapley Value to get an estimate of disparity (notice the definition for disparity above relies on having access to the data for the whole user study, and hence it cannot be directly used within in the bandit), and (2) using a fairness-aware bandit strategy targeted at reducing disparity.

3.3 Improved Design: Idle Social Exergame *Step Heroes*

For the main user study presented in this paper, we implemented the same mechanism of adapting the social comparison environment of two users by MAB, this time within the context of a mobile game called *Step Heroes*. There is a large body of work that shows incorporating game-like elements is effective in motivating physical activity [24, 28, 61]. In particular, *Step Heroes* was designed to include idle game features, which can be especially effective in behavior change games [6, 55, 99].

Step Heroes is a multiplayer exergame where teams of three players work together to combat invasive enemies who are draining the health of the forest. The goal of the game is, as a team, to clear as many forest areas as possible by defeating the enemies living there. Similar to many idle games, a main part of the gameplay in *Step Heroes* centers around upgrading avatars. A player needs to decide how to allocate her currency (based on her real-world steps and the experience points she earned from defeating past enemies) to develop the different combat abilities of her avatar. To be really successful, she also needs to work together with her teammates so that their avatars have complementary strengths in the joint battles (Figure 3, middle panel).



Figure 3: Overview of the user interface of *Step Heroes*. Part A illustrates the unique avatars representing each player in the game environment. Part B showcases the attack boost meter in which players can contribute “steps” to boost their avatar’s attack power. Part C shows the currency meter, consisting of “pollen” earned from defeating enemies and spent on avatar upgrades to their specific stats (Part E). Players may also purchase and equip a variety of available weapon options that each have unique effects and abilities for further customization of their avatar, shown in part D.

A critical design decision was to bring the avatar upgrade mechanics to the center of players’ collaboration and social comparison. Each avatar has a blossom over its head (Figure 3, part A). The blossoms visually expand and contract according to the number of steps each player has taken, providing a visual representation of each player’s contribution to the team for easy social comparison. In addition, the color of an avatar’s blossom reflects its combat abilities. For example, if a player invests heavily in her avatar’s speed and attack abilities, its blossom will show yellow and red. This design choice was intended to summarize a player’s gameplay decisions and help other players to decide how to upgrade their avatars for a more balanced team composition.

Step Heroes is designed for users to experience in regular short bursts of play every day over several weeks. As an idle game, *Step Heroes* is designed to encourage players to frequently return to examine their team’s progress and to spend their step-based currency on upgrading their avatars [99]. In addition, we designed the game to support a wide range of engagement to fit within users’ daily schedules. A player can accumulate all the steps and play a short session a day or log into the game and play as many times as they like.

In the rest of the section, we focus on how a player’s interaction with the profiles in *Step Heroes*—the main intervention that uses

social comparison—motivates players’ PA. It consists of a three-step process. Notice that in order to test our AI, each team consists of two human players and one AI-controlled artificial player. Human players are again not made aware that one of their teammates is an AI player.

Teammate Profile Selection. Each day when a player logs in, the game first displays the player’s daily step count from the previous day at the top of the interface (automatically retrieved from the player’s Fitbit account) and the profiles of both of their teammates (Figure 4, part A). We purposefully designed this selection page with minimal information so the player could focus on a single dimension for social comparison: the previous day’s total steps. The player is asked to select between their teammates to continue, with a notification that they can only choose one profile to review *in full* each day and cannot go back to choose another.

Full Player Profile. After the player confirms their selection, they are given more detailed information about the selected teammate. The profile reveals further information about that player’s game decisions, such as stamen upgrades, weapons, and abilities. For example, players may click on their teammate’s weapon or stamen icons (Figure 4, part B) to reveal details about their teammate’s actions. Below the game information, we provide users with

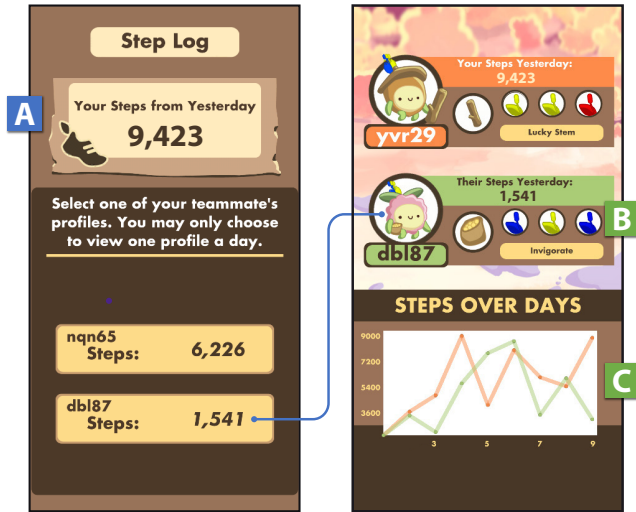


Figure 4: The Step Heroes Teammate Profile Selection screen (left) and the Full Game Profile screen (right) that is revealed when a teammate is chosen.

a step graph to examine and compare their step counts with the selected teammate over time (Figure 4, part C). This screen is the only time users can examine their teammate’s profiles. During gameplay, we did not provide access to their teammate’s profiles to ensure a meaningful selection at the start of each day.

Main Gameplay. After the user has explored the selected profile, they may continue to the game interface (center image in Figure 3, where players may begin to play *Step Heroes*). Players can return to the game as often as they would like throughout the day; however, the Teammate Profile Selection screen will only display once during their first login of each day.

4 SHAPLEY BANDITS: USING SHAPLEY-BASED FAIRNESS CONSTRAINTS TO GUIDE BANDITS

To address the Greedy Bandit Problem (Section 3.2), we propose a new type of multi-armed bandit [87, 97] with two new contributions. The first is a method for calculating disparity by making use of the Shapley Value equation (Equation 1) to estimate the contributions of each player (notice that the definition for disparity in Section 3.2 requires having access to the data for the whole user study, and hence can only be used for post-user study analysis). The second is a new type of fairness-aware bandit strategy designed to reduce disparity while at the same time maximizing reward. These two contributions result in what we call a *Shapley Bandit*. Let us first introduce the concepts of the Shapley Value and multi-armed bandits before describing Shapley Bandits.

4.1 Shapley Value

The core of our solution for allocating AI attention toward group members is rooted in economic theory, where we leverage the concept of the Shapley Value to determine fair attribution of a group’s

success to the individuals within that group. The Shapley Value [92] was proposed as a method for determining the costs and rewards that ought to be attributed to individuals working together in a group endeavor [101]. Though the individual performance within a group may be known, the output of the individual members as a result of joining together as a team is more difficult to assess (due to synergies and cross-influences among team members). For example, consider a group of people joining together to be physically active (measured in daily walking steps), where each individual walks an average of 10,000 steps daily. Then, a new team member joins who walks 12,000 steps daily, after which the average steps of the other people increase to 11,000 steps per person. The new member’s contribution to the team is not just their 12,000 steps but also part of the extra 1,000-step improvement gained by the others (which would not have happened otherwise).

The Shapley Value aids in assessing the total contribution of an individual member toward the group’s resulting aggregate performance. Shapley proposed that such a calculation would satisfy the following four axioms (and his solution has been shown to be the only one that will do so [64, 89]): 1) Symmetry, where two individuals are interpreted as equivalent if they prove to be completely interchangeable within the group; 2) Nullity, where an individual who provides no value to any part of the group beyond their individual achievement is given zero synergistic attribution; 3) Additivity, where if separating a game into smaller games (i.e., “rounds”), the contribution scores from the rounds add up to the score for the whole game; and 4) Efficiency, where the contributions calculated for the individuals add up to the full contribution of the group.

For a player i participating in a coalition of players N , we consider each hypothetical sub-coalition (S) of players of N that do not yet contain player i (i.e., $N \setminus \{i\}$). We then add i to each of these sub-coalitions, noting the marginal benefit that player i adds to the group, and average these values. The result is the expected value for the contribution that i brings to N as a whole. This marginal value contributed by i , denoted as $\phi_i(N, v)$, is summarized in the following equation [64]:

$$\phi_i(N, v) = \frac{1}{N!} \sum_{S \subseteq N \setminus \{i\}} |S|!(|N| - |S| - 1)! [v(S \cup \{i\}) - v(S)] \quad (1)$$

The Shapley Value has found application across many fields, including law [32], economics [74], political science [49], and machine learning [41, 72]. In this paper, we use the Shapley Value to estimate player *disparity* in a group during the execution of the bandit strategy.

4.2 Multi-Armed Bandits

A multi-armed bandit (MAB) problem is a type of decision problem [87, 97] in which an agent must repeatedly select from a set of mutually exclusive opportunities (called *arms*). The rewards that would be returned from these arms are not known to the agent at the time of selection; however, over time, through repeated selections (or *pulls*), the agent can build an understanding of the reward distributions available across the arms. Over its complete sequence of selections, the agent aims to maximize its total reward; therefore, it must continuously balance spending its limited pulls toward *exploring* the different arms to gather information about which arms

are most promising versus *exploiting* the arm it currently believes to be the best [58].

Popular strategies include those of the upper confidence bound family (e.g., UCB1 [12]) or the ϵ -family of strategies (e.g., ϵ -greedy [58]). In this paper, we formulate a player modeling problem as an MAB problem as in our previous work [44, 45], where different game adaptation options appealing to specific player types are made available as arms to an MAB-based AI agent. As the agent selects among the adaptation options, it observes response metrics from the players (such as their self-reported motivation or daily steps) to determine which adaptations are most effective at driving those metrics.

4.3 Shapley Bandits

In contrast to a typical bandit strategy, which only aims to maximize overall utility by selecting the arm predicted to perform the best, a Shapley Bandit tracks the contribution attributed to each player and prioritizes catering to players' preferences based on those relative contributions.

Shapley Disparity: Given a set of players N , a Shapley Bandit estimates the *Shapley Disparity* of each player i in the following way. At each step t , when a reward is received for an arm pull (i.e., the number of steps of a player is observed) the Shapley Value (Equation 1) for player i is computed and is added to a *Cumulative Shapley Value*, CSV_i . We then define the *CSV Ratio*, CSV_R_i , as the Cumulative Shapley Value of player i divided by the sum of that of all players. In parallel, the bandit maintains a *Treatment Counter*, TC_i , that counts the number of times up to time t that the bandit has selected the arm currently considered to be the best for player i . Analogously, we define the *TC Ratio*, TCR_i , as the sum of the Treatment Counter for player i divided by the sum of that of all players. We then define the *Shapley Disparity* of a player i at a given time t as $SD_i = |CSV_R_i - TCR_i|$. Intuitively, this measures the difference between the fraction of the contribution of a player to the group and the proportion of times the bandit has catered to that player.

Also, notice that in order to evaluate Equation 1, we need access to the performance of different sub-coalitions. Since not all of them are observable, in this study, we approximated them by the sum of the individual contributions (steps) of each member of the sub-coalition.

Shapley Bandit: Given a group of players N , at each iteration t , instead of aiming to select the arm that maximizes the expected sum of step counts (as a Greedy Bandit would do), a Shapley Bandit works as follows:

- **Exploration:** With probability ϵ ($\epsilon = 0.01$ in our experiments¹), the bandit would decide to *explore* and choose a random arm.
- **Exploitation:** Otherwise, the bandit decides to *exploit*. The bandit then identifies which player i would result in the lowest overall sum of Shapley disparities (sum of Shapley disparities of all players in N) if the bandit is to cater to i in this turn. Then, the arm currently estimated to be the best

for player i will be chosen. To predict which is the best arm, the bandit uses all the observed rewards from that player up to time t , to estimate the expected reward of each arm.

For example, if after 9 rounds Player 1 has $CSV_1 = 27,500$ steps with $TC_1 = 5$ and Player 2 has $CSV_2 = 32,800$ steps and $TC_2 = 4$, a Shapley Bandit would calculate a CSV ratio $CSV_R_1 = 27500 / (27500 + 32800) = 0.456$, and TC ratio $TCR_1 = 5 / (4 + 5) = 0.556$ for Player 1, resulting in a Shapley Disparity of $SD_1 = |0.456 - 0.556| = 0.1$. If we calculate the same for Player 2, we would also obtain $SD_2 = 0.1$ (with a current Shapley Disparity sum of $0.1 + 0.1 = 0.2$). If we were to cater to Player 1 (and TC_1 would become = 6), the sum of their disparities (doing similar calculations) would become: 0.288, and if we were to cater to Player 2 (and TC_1 would become = 5), the sum would become: 0.088. A Shapley Bandit predicts Arm A to be the best for Player 1 and Arm C to be the best for Player 2. Hence, if in the next iteration, if the bandit decides to *exploit*, it would choose to cater to Player 2 because it has a larger Shapley Disparity, and the MAB would select Arm C for the current round (and increment Player 2's TC value).

In this way, a Shapley Bandit prioritizes the treatments allocated to players in a manner commensurate with each player's level of contribution to the team. In our case, we targeted steps to represent this contribution because in *Step Heroes* a player's steps serve both as a proxy for the player's real-world effort and the increased in-game utility they provide to their overall team's capability in conquering the game's challenges.

Comparing Shapley Bandits with traditional stochastic bandits, we would like to make a few observations:

- A Shapley Bandit directly minimizes "Shapley disparity" rather than maximizing utility. There is still an element of utility maximization once the player with the highest disparity is selected and a Shapley Bandit selects the best arm for that player; however, this is still subject to the main disparity minimization objective. Notice that minimizing Shapley disparity might sometimes go against maximizing utility (e.g., selecting the best choice for a player with a high Shapley disparity might imply choosing an arm that is not the one that would result in the highest utility gain overall). For this reason, we also note that traditional performance metrics for stochastic bandits, such as *regret* [12], are not directly applicable and would need redefinition.
- A Shapley Bandit still needs to estimate which arms maximize utility for a given player in order to cater to them. Hence, it benefits from exploration in the same way as a traditional bandit.
- A Shapley Bandit needs to maintain a series of additional quantities that traditional stochastic bandits do not need, namely a Cumulative Shapley Value and Treatment Counter for each human player. However, given the small groups of users in our study (2 human players plus one AI player), the extra computational cost is negligible.

Therefore, we do not expect a Shapley Bandit to outperform a traditional, Greedy Bandit in terms of overall maximization of the target metric (in this case, daily steps). Instead, we anticipate incurring a cost to that total utility in favor of a strategy that encourages overall player retention. However, we anticipate that

¹Notice this is a relatively very low exploration rate, but we set it this way because we have a forced-exploration period at the beginning of the study where during the first 9 days the bandit always explores.

this fairness constraint will provide a better experience to a greater number of users (measured via adherence metrics) than focusing on purely maximizing total step counts.

5 METHODS

We conduct a user study in the context of *Step Heroes* to evaluate the effectiveness of using Shapley Bandits to personalize the social comparison environment for a small group of users in the context of a social exergame. Specifically, we aim to evaluate the following hypotheses:

- **H1:** The Greedy Bandit will outperform the Shapley Bandit in increasing players' average steps and self-reported motivation, the two metrics that the Greedy Bandit aims to maximize directly.
- **H2:** Our Shapley Bandit will outperform the Greedy Bandit in increasing player retention and adherence with the intervention.

We recruited 46 participants (36 women, 10 men) from a major university in the U.S. For our study, healthy adults were eligible if they reported PA as being somewhat or very important to them and had access to a Fitbit health tracker or a Fitbit-compatible smartphone. We did not include any applicants that reported medical conditions or advisement that would limit their PA. We compensated participants with either gift cards or equivalent extra course credit at the end of the study. The study procedure was reviewed and approved by the Institutional Review Board (IRB) at Drexel University.

5.1 Procedure

After providing their demographic information, participants were given a set of psychology surveys, including the Iowa Netherlands Comparison Orientation Measure (INCOM-23) [42] and the Identification/Contrast Scale (ICS) [23]. These instruments provided baselines regarding a participant's social comparison orientation. Participants were then provided instructions on accessing *Step Heroes* and linking their Fitbit accounts. Players were asked to wear their Fitbit device every day and log into the game at least once per day to synchronize their steps and participate in their daily sessions. The first three days of the study were considered a *baseline period*, where we collected players' daily steps before introducing our game.

At the end of this baseline collection period, participants were ranked according to their baseline results and separated into pairs of similarly-ranked players for the rest of the intervention. Each pair was joined by an artificial player (whose steps were controlled by the game AI) to play *Step Heroes* as a team. Each participant could only interact with the other human player and the AI player in their team. Each team was randomly assigned one of the following three conditions:

- (1) *Control Condition:* The participants play a version of *Step Heroes* where the AI makes a random decision about the artificial player's steps and gameplay information each day (n=14).
- (2) *Experimental Condition #1:* The participants play a version of *Step Heroes* where the AI makes decisions based on the

same Greedy Bandit strategy used in the preliminary design study (n=14).

- (3) *Experimental Condition #2:* The participants play a version of *Step Heroes* where the AI makes decisions based on a Shapley Bandit strategy (n=18).

We asked participants to play at least once per day for 21 days. We also asked them to fill out three weekly qualitative surveys to assess any technical issues and gain insight into player motivation for physical activity and their experience with the game. Finally, we administered an exit survey at the end of participation.

5.2 Data Analysis

For quantitative data, we use inferential statistics for analysis. Specifically, we analyzed several primary data sources resulting from the user study:

- (1) Data related to the artificial player, including MAB predictions and selections.
- (2) User metrics regarding their interaction with the platform, including daily steps, reported motivation in the pre- and post-session questionnaire, and user participation.
- (3) Open-response feedback surveys provided to each participant following their 7th, 14th, and 21st session, along with an exit survey provided to each participant on their completion of the study.

Qualitative feedback in the weekly and exit surveys was coded by two independent researchers who used thematic analysis [19] to identify themes in the dataset. After initial observations, two researchers conducted a preliminary open coding. During this analysis step, both researchers reviewed each response and noted initial labels separately (e.g., the user noted a positive impact, or the user reported an increased step awareness). Then, they discussed their initial codes to form a codebook. Both researchers then independently coded the full dataset. Finally, they collaboratively discussed disagreements and modified their codes until they reached full consensus [51].

6 RESULTS

6.1 Quantitative Results

We first consider H1, where we expect that the Greedy Bandit will outperform the Shapley Bandit in the two metrics that the Greedy Bandit aims to maximize directly (i.e., average player steps and self-reported motivation). The Greedy Bandit operates without considering the fairness constraints integrated into a Shapley Bandit, and therefore we expect to observe a cost for these constraints.

Table 1 reports in the two left columns results on average steps for players in all three conditions relative to their baseline steps collected in the three days prior to starting their *Step Heroes* game. We see that the players in the Shapley Bandit condition achieved fewer steps than the Greedy Bandit condition (which was optimizing for that step count) and the random control condition. However, no condition was found to be statistically significantly different (i.e., via two-tailed T-test) from any other.

The two right columns of Table 1 present the average self-reported motivation following each session in all three conditions. Although the Greedy Bandit aims to maximize this metric along with steps, we

Table 1: Average Steps vs. Baseline for players in the Control vs. Greedy and Shapley Bandit conditions

Condition	Steps vs. Baseline	P-Value vs. (C)	Post-Motivation	P-Value vs. (C)
Control (C)	-69		1.316	
Greedy (E1)	171	0.1868	1.400	0.095
Shapley (E2)	-313	0.1023	1.451	0.011

find that the players in the Shapley Bandit condition reported higher post-test motivation scores. Though there was no statistically significant difference between the two experimental conditions when compared directly, the Shapley Bandit condition demonstrates a statistically significant difference between itself and the random control condition ($p < 0.05$), whereas the Greedy Bandit did not.

We next consider H2, where we anticipate that the Shapley Bandit condition will mitigate the symptoms of the Greedy Bandit Problem. We believed this might result in lower missed sessions and lower variance in disparity scores among players in the Shapley Bandit condition. Though the data showed both of these expectations to be true—miss rates were 26% vs. 25% and average disparity scores were 0.070 vs. 0.067 in the Greedy Bandit and Shapley Bandit conditions, respectively—neither of these findings were statistically significant due to sample size.

However, we anticipated that H2 would be supported most strongly in the results of the disparity analysis that first illustrated the symptoms of the Greedy Bandit Problem. Indeed, Figure 5 conducts the same disparity analysis on the user study data we conducted on the preliminary design study data (Section 3.2), with the left graph illustrating results for the Greedy Bandit and the right showing those of the Shapley Bandit. Regarding the correlation between disparity and the player’s propensity to miss sessions, the Greedy Bandit performed exactly as it did in the preliminary design study, again demonstrating a Pearson’s R correlation of $R = 0.16$. In contrast, the Shapley Bandit demonstrated a negative correlation of $R = -0.39$. The difference between these coefficients [82] was found to be statistically significant ($p < 0.05$). In other words, the primary symptom of the Greedy Bandit Problem (i.e., higher disparity correlating with higher non-adherence) was confirmed in the Greedy Bandit condition of our user study scenario. However, in the Shapley Bandit condition, players experience the opposite relationship between disparity and non-adherence. Let us now look at our qualitative analysis, which provides further support for H2.

6.2 Qualitative Results

The research team analyzed the participants’ qualitative feedback using thematic analysis [19]. Two coders conducted a preliminary open coding to identify notable patterns and used the consensual qualitative approach [51] until they reached full agreement. Consensus coding is generally more suited for small samples and for considering multiple viewpoints [70] than inter-rater reliability (IRR).

Qualitative feedback was coded into 11 distinct themes related to participants’ reasons for how they believed their participation impacted *motivation for PA*, *level of PA*, and *group happiness*. Three

additional codes were derived to capture whether or not the participant reported a positive or no effect from their response to questions asking participants to describe how participation impacted their overall *motivation for PA* and *level of PA*. In addition, feedback was also coded into three themes relating to other ways (besides increasing PA) participation affected their daily activities (i.e., Awareness, Relief, Step Habit). Finally, five themes emerged for how participants chose their daily profile selection. An overview of all codes is presented in Table 2.

Overall, the qualitative feedback suggested that participants believed the team they were assigned to was comprised of other participants in the study. 34 of 46 participants provided the post-study survey responses, resulting in 150 responses from the five survey questions. The weekly survey was administered three times throughout the study (i.e., Day 7, Day 14, and Day 21). 32 of 46 participants provided responses, resulting in 332 responses from the four survey questions. We consolidated the weekly and post-study survey results for questions, asking participants to rate and describe how participation impacted their “motivation for PA” and “level of PA,” as these questions were the same for both surveys.

Motivation for PA: When asked to describe how participation impacted their overall motivation for physical activity, 42% of participants in the Shapley Bandit condition reported a positive impact versus 27% in the Greedy Bandit condition. A majority of participants explicitly acknowledged that participation impacted their motivation to walk more or increase their step count. For example, one participant wrote, “Yes. It makes me want to walk more so that I have more steps on my profile.” In contrast, 69% of participants in the Greedy Bandit condition reported a negative impact versus 42% of participants in the Shapley Bandit condition.

Participants elaborated on other ways participation affected their daily activities, such as increased PA awareness (13%) and the feeling of relief (3%). For example, one participant wrote, “I’ve actually been sick, so my steps are really low for me. Seeing other people around the same amount of steps as me right now, makes me think of when I feel better and get more steps in a day.”

Level of PA: When asked to describe how participation impacted their level of physical activity, 26% of participants in the Shapley Bandit condition reported a positive impact in the level of PA versus 16% in the Greedy Bandit condition. For example, one participant explained, “Yes, prior to the study I would talk about it but not really do it, after, I actually got out there and did it, even if it was just a little.” In contrast, 76% of participants in the Greedy Bandit condition reported no impact versus 65% in the Shapley Bandit condition. One participant reported, “Not really, most of my steps are gotten between classes.”




Figure 5: Comparison of Greedy Bandit (left) to Shapley Bandit (right) disparity analysis. Compare left to Figure 2, reflecting the same analysis from the Greedy Bandit’s deployment in the preliminary design study. As in its performance in the preliminary design study, the Greedy Bandit (left) again yields a correlation of $R=0.16$ between player disparity and non-adherence. In contrast, the Shapley Bandit (right) demonstrates a negative correlation coefficient of $R= -0.39$ that differs significantly ($p<0.05$) [82] from the Greedy Bandit.

Participants elaborated on other ways participation affected their level of PA. For example, one participant wrote, “...I’m more aware of the steps I’m getting and the days that I’m not as active.” 13% described changes in their daily step habits. For instance, a participant noted, “There were times when I chose to walk instead of finding a ride home because of the game.”

Group Happiness: Participants were asked questions regarding their experience with their team. Using a 5-point Likert scale, they rated how happy they were with the group they were assigned within the study (0 = very unhappy, 4 = very happy). Participants in the Shapley Value condition responded with an average Likert score of 2.4, while the Greedy Bandit condition scored 2.2.

When participants were prompted to elaborate on their answers, 63% provided a reason. For example, one participant simply wrote, “They were fine,” which we coded as not providing a reason. However, those who provided a reason reported step counts (42%) as a significant factor in their team happiness. For instance, a participant wrote, “I think that the players I was matched with took the same amount of steps so it wasn’t intimidating or demotivating.” In contrast, participants reported a lack of team connection (18%) and insufficient information (18%) as reasons for not being as happy with their team. For example, a participant who rated their team as “Neutral” wrote, “Did not have enough info on each player to really connect with them.”

Choosing Profiles: When participants were asked to describe how they chose their profile each day, participants reported steps being a significant factor in their decision, ranging from choosing based on the highest steps (44%), similar steps (14%), or fewer steps (10%) to their own. Alternatively, 34% reported they chose randomly, and 7% alternated between teammates.

Motivation to Login: Using a 5-point Likert scale, participants were also asked to rate their motivation to log in to the game everyday (0 = not motivated, 4 = highly motivated). Participants in the Shapley Value condition had a Likert score of 2.3, while the Greedy Bandit condition had a score of 2.1.

These results, in general, show support for H2, where participants in the Shapley Bandit condition reported higher motivation, higher positive impact in the level of PA, and higher group satisfaction.

7 DISCUSSION

The increasing popularity of using technology, especially AI, to personalize healthcare and to allocate resources among diverse social groups [79, 94] makes fairness a pressing issue. We believe that fairness-aware Shapley Bandit helped to mitigate the symptoms of the Greedy Bandit Problem observed in both our pre-test and our *Step Heroes* user study under the Greedy Bandit condition. Our hypotheses were partially supported by data resulting from our user study. We found that the Greedy Bandit outperformed the Shapley Bandit in terms of average step achievement by players, in line with our intuition that the fairness constraints would incur a penalty compared to a strategy that aimed to maximize this value at all costs (although this difference was not found to be statistically significant between the two conditions).

Our results showed that the Shapley Bandit outperformed the random control condition in terms of self-reported motivation to a statistically significant degree ($p<0.05$), where the Greedy Bandit was not able to achieve the same. In retrospect, this supports intuition regarding the aim of the Shapley Bandit’s design toward improving player affect toward the intervention (e.g., adherence and gameplay participation) as opposed to the Greedy Bandit’s aim toward maximizing specific metrics (i.e., daily steps). It may therefore be reasonable that player self-reported motivation is more closely related to the intended affect change than a behavior metric such as steps. Although this outcome was initially unexpected, it is a positive reflection on the Shapley Bandit’s performance.

Our second hypothesis was fully supported by the results of our *Step Heroes* user study, specifically in the repeating of the disparity analysis we had conducted on our pre-test results (where the symptoms of the Greedy Bandit Problem are observed). Our

Table 2: The list of all codes and associated examples related to participants' qualitative feedback.

Category	Theme	Example
Reported PA Effect (3)	Yes	P106: "It motivates me"
	No Effect	P47: "Sometimes it encouraged me to do more, other times the numbers are just there without much impact or effect on me."
	Sometimes	P217: "Seeing other profiles hasn't affected my motivation"
Other Reported Effects (3)	PA Awareness	P205: "A little bit, I've been more aware of the steps I'm getting and the days that I'm not as active"
	Relief	P26: "I've actually been sick, so my steps are really low for me. Seeing other people around the same amount of steps as me right now, makes me think of when I feel better and get more steps in a day."
	Step Habit	P250: "Yes, it has made me take the stairs instead of elevator at work in order to get more step count."
Reasoning for Effects (11)	Lower Team Steps	P168: "Normal. They've been slacking lately so I feel better about myself"
	Higher Team Steps	P229: "Usually I have more steps than the other players, but when I do not have as many steps as them it makes me more motivated."
	Similar Team Steps	P28: "I feel less bad about how many steps I take since they're about the same"
	Competition with Team	P91: "It is motivating to want to beat the player's number of steps."
	Cooperation with Team	P267: "More motivated to try strategies, who may be stronger, teamwork"
	Viewing Team Profiles	P91: "Seeing other players' profiles motivates me to want to do better."
	Lack of Team Connection	P214: "It doesn't really motivate me because I have no real connection to the person."
	Not Enough Info	P255: "I did not know anything about them other than the number of steps each day. I wondered why neither of their characters had a ranking."
	Underperforming Team	P110: "It was a little frustrating because I think I had higher stats than them the whole time"
Work Schedule		P187: "It has not impacted my activity. I only get as much as I can given my circumstances and busy schedule."
	Personal Motivation	P267: "No, still personally motivated, independent of others"
Choosing Profiles (5)	Most Steps	P97: "Usually whoever has higher steps"
	Least Steps	P120: "I look at ones who have less steps than mine so I feel better about myself. (That's mean but it's the truth)."
	Similar Steps	P20: "I choose the profile that has a step count closest to mine, or at random if they are all similar."
	Alternate	P168: "I switch between the two people every day."
	Random	P263: "I just choose at random."

disparity analysis in the *Step Heroes* user study data showed that 1) the Greedy Bandit performed consistently and with the same resulting correlation (Pearson's $R=0.16$) relating a player's disparity measurement to missed sessions even in the new context of the *Step Heroes* exergame, and 2) that the Shapley Bandit reduced this correlation (Pearson's $R=-0.39$) to a statistically significant degree ($p<0.05$) [82].

7.1 Implications

Our research has several implications. Concerning the literature of multi-armed bandits, the idea of introducing fairness constraints into MABs (in our case via the use of Shapley Values) in the context of group preference modeling and adaptation opens up many new avenues for future work. For example, further research is needed to study the theoretical properties of this new type of bandit and devise new measures to evaluate them beyond the traditional *regret* metrics, which do not fully capture fairness considerations. Notice that our idea of Shapley Bandits is not specific to exergames but is

a general approach that is applicable to any domain that involves groups or teams where the bandit makes choices that affect more than one individual. Consider, for example, the idea of music recommendation for groups and task allocations in a team, as discussed in the related work.

Additionally, we show that holistic and ethical considerations of the AI's utility function are a key design issue for personalized adaptive applications. Our results provide evidence that, even though introducing fairness considerations in adaptive group exergames does not directly lead to higher PA levels, it can contribute to user retention and adherence. Adherence and long-term engagement are the foundation of whether a behavioral change intervention can have a lasting positive impact. Our research shows that if the AI only focuses on short-term performance (e.g., a group's aggregated daily steps), it risks consistently ignoring the needs of marginalized users and, over time, losing their participation. This insight has implications for the future design of player-AI interaction [111] not only in adaptive group exergames but behavioral change apps in general.

8 CONCLUSIONS

In this paper, we address the problem of fairness in AI-based group personalization via multi-armed bandits. Specifically, we presented evidence of what we call the *Greedy Bandit Problem*, arising when deploying a bandit-based AI system to personalize the group experience in a social exergame, resulting in the system focusing primarily on high-performance players and causing adherence problems. We use these findings to construct a new fairness-aware bandit strategy approach based on minimizing disparity estimated via the Shapley Value, which aims to align the degree of AI attention a player receives to the player's measured investment toward helping their team. We deployed that strategy in a new user study based around a team-based, idle exergame and evaluated results against two hypotheses that assessed our Shapley Bandit's performance against the traditional greedy approach. Our results support our intuition regarding the Shapley Bandit's ability to mitigate the effects of the Greedy Bandit Problem.

We identify a few limitations of our approach, including the small sample size of this initial *Step Heroes* user study. Future studies with larger sample sizes may help to further establish some of the trends so far observed. Additionally, our participants are healthy, college students with relatively homogeneous health needs. Future research is needed to test our approach with a more diverse group of users. We also note that our application of the Shapley Value does not currently include all terms we wish to consider regarding user participation. The Shapley Value, and by extension the axioms that it implements, intends to capture the totality of value that an individual may bring to a group endeavor, including some benefits that may not be easily measurable. For example, we argue that a player's mere participation in a game yields value to others simply in the fact that it enables a multiplayer experience. However, this is not currently reflected in the Cumulative Shapley Value (CSV) discussed in Section 4, which is based solely on step performance. Solving this problem is beyond the scope of this paper, where the Shapley Value provided guidance in our effort to mitigate the effects

of the Greedy Bandit Problem. However, this remains an interesting problem for future work.

Finally, there are many alternatives for modeling fairness [13], where we have opted for an equality-enforcing individual notion of fairness. Equality enforcing notions of fairness can have undesired side-effects (such as leveling down) in cases where there is utility dependency between the players [85], as it is in our case. We leave the study of whether or not such side-effects can happen in our framework as well as studying notions of fairness within as future work.

ACKNOWLEDGMENTS

This work is partially supported by the National Science Foundation (NSF) under Grant Numbers IIS-1816470, IIS-1917855 and Danish Novo Nordisk Foundation under Grant Number NNF20OC0066119. The authors would like to thank all past and current members of the projects.

REFERENCES

- [1] J Stacy Adams. 1963. Towards an understanding of inequity. *The journal of abnormal and social psychology* 67, 5 (1963), 422.
- [2] J Stacy Adams. 1965. Inequity in social exchange. In *Advances in experimental social psychology*. Vol. 2. Elsevier, 267–299.
- [3] Aino Ahtinen, Minna Isomursu, Muzayun Mukhtar, Jani Mäntyjärvi, Jonna Häkkinen, and Jan Blom. 2009. Designing social features for mobile and ubiquitous wellness applications. In *Proceedings of the 8th international Conference on Mobile and Ubiquitous Multimedia*. 1–10.
- [4] Ifeoma Ajunwa and Daniel Greene. 2019. Platforms at work: Automated hiring platforms and other new intermediaries in the organization of work. In *Work and labor in the digital age*. Emerald Publishing Limited.
- [5] Sheldon Alexander and Marian Ruderman. 1987. The role of procedural and distributive justice in organizational behavior. *Social justice research* 1, 2 (1987), 177–198.
- [6] Sultan Alharthi, Olaa Alsaedi, Zachary Dugas Toups, Joshua Tanenbaum, and Jessica Hammer. 2018. Playing to Wait: A Taxonomy of Idle Games. 1–15.
- [7] Danielle Arigo, Megan M Brown, Kristen Pasko, and Jerry Suls. 2020. Social Comparison Features in Physical Activity Promotion Apps: Scoping Meta-Review. *Journal of Medical Internet Research* 22, 3 (2020), e15642.
- [8] Danielle Arigo, Leah M Schumacher, Emilie Pinkasavage, and Meghan L Butryn. 2015. Addressing barriers to physical activity among women: A feasibility study using social networking-enabled technology. *Digital health* 1 (2015), 2055207615583564.
- [9] Sonia M Arteaga, Mo Kudeki, Adrienne Woodworth, and Sri Kurniawan. 2010. Mobile system to motivate teenagers' physical activity. In *Proceedings of the 9th international conference on interaction design and children*. 1–10.
- [10] Homa Atabakhsh. 1991. A survey of constraint based scheduling systems using an artificial intelligence approach. *Artificial Intelligence in Engineering* 6, 2 (1991), 58–73.
- [11] Jean Aubin. 1992. Scheduling ambulances. *Interfaces* 22, 2 (1992), 1–10.
- [12] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47, 2-3 (May 2002), 235–256.
- [13] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org.
- [14] Femke TC Bennenbroek, Bram P Buunk, Karen I Van Der Zee, and Brigit Grol. 2002. Social comparison and patient information: what do cancer patients want? *Patient Education and Counseling* 47, 1 (2002), 5–12.
- [15] Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2017. A Convex Framework for Fair Regression. *CoRR* abs/1706.02409 (2017).
- [16] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in Criminal Justice Risk Assessments: The State of the Art. *Sociological Methods & Research* 50 (2018), 3–44. Issue 1.
- [17] Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. 2020. Fairlearn: A toolkit for assessing and improving fairness in AI. *Microsoft, Tech. Rep. MSR-TR-2020-32* (2020).
- [18] Gabriele Bleser, Daniel Steffen, Attila Reiss, Markus Weber, Gustaf Hendebey, and Laetitia Fradet. 2015. Personalized physical activity monitoring using wearable sensors. In *Smart health*. Springer, 99–124.

- [19] Virginia Braun and Victoria Clarke. 2012. *Thematic analysis*. American Psychological Association.
- [20] Fiona C Bull, Matthew W Kreuter, and Darcell P Scharff. 1999. Effects of tailored, personalized and general health messages on physical activity. *Patient education and counseling* 36, 2 (1999), 181–192.
- [21] Fabio Buttussi and Luca Chittaro. 2008. MOPET: A context-aware and user-adaptive wearable system for fitness training. *Artificial Intelligence in Medicine* 42, 2 (2008), 153–163.
- [22] Bram P Buunk. 1995. Comparison direction and comparison dimension among disabled individuals: Toward a refined conceptualization of social comparison under stress. *Personality and Social Psychology Bulletin* 21, 4 (1995), 316–330.
- [23] Bram P Buunk and Jan F Ybema. 1997. Social comparisons and occupational stress: The identification-contrast model. *Health, coping, and well-being: Perspectives from social comparison theory* (1997), 359–388.
- [24] Taj Campbell, Brian Ngo, and James Fogarty. 2008. Game design principles in everyday fitness applications. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. 249–252.
- [25] Karina Caro, Yuanyuan Feng, Timothy Day, Evan Freed, Boyd Fox, and Jichen Zhu. 2018. Understanding the effect of existing positive relationships on a social motion-based game for health. In *Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare*. 77–87.
- [26] Dennis L Chao, Justin Balthrop, and Stephanie Forrest. 2005. Adaptive radio: achieving consensus using negative preferences. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. 120–123.
- [27] Jiahao Chen, Nathan Kallus, Xiaojie Mao, Geoffrey Svacha, and Madeleine Udell. 2019. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *Proceedings of the conference on fairness, accountability, and transparency*. 339–348.
- [28] Yu Chen and Pearl Pu. 2014. HealthyTogether: exploring social incentives for mobile fitness applications. In *Proceedings of the second international symposium of chinese chi*. 25–34.
- [29] KP Chow and CK Hui. 1993. Knowledge-based system for rostering. *Expert Systems with Applications* 6, 3 (1993), 361–375.
- [30] Ronald L Cohen. 1987. Distributive justice: Theory and research. *Social justice research* 1, 1 (1987), 19–40.
- [31] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 797–806.
- [32] Bertrand Crettez and Régis Deloche. 2019. A law-and-economics perspective on cost-sharing rules for a condo elevator. *Review of Law & Economics* 15, 2 (2019).
- [33] Andrew Crossen, Jay Budzik, and Kristian J Hammond. 2002. Flytrap: intelligent group music recommendation. In *Proceedings of the 7th international conference on Intelligent user interfaces*. 184–185.
- [34] Sriharsha Dara, C. Ravindranath Chowdry, and Chintoo Kumar. 2020. A survey on group recommender systems. *Journal of Intelligent Information Systems* 54, 2 (2020), 271–295. <https://doi.org/10.1007/s10844-018-0542-3>
- [35] Kathryn A Dowsland. 1998. Nurse scheduling with tabu search and strategic oscillation. *European journal of operational research* 106, 2-3 (1998), 393–407.
- [36] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. 214–226.
- [37] Deborah L Feltz, Samuel T Forlenza, Brian Winn, and Norbert L Kerr. 2014. Cyber buddy is better than no buddy: A test of the Köhler motivation effect in exergames. *GAMES FOR HEALTH: Research, Development, and Clinical Applications* 3, 2 (2014), 98–105.
- [38] Leon Festinger. 1954. A Theory of Social Comparison Processes. *Hum. Relations* 7, 2 (May 1954), 117–140.
- [39] Evan M Forman, Stephanie G Kerrigan, Meghan L Butryn, Adrienne S Juarascio, Stephanie M Manasse, Santiago Ontañón, Diane H Dallal, Rebecca J Crochiere, and Danielle Moskow. 2019. Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss? *Journal of behavioral medicine* 42, 2 (2019), 276–290.
- [40] Sorelle A Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P Hamilton, and Derek Roth. 2019. A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the conference on fairness, accountability, and transparency*. 329–338.
- [41] Amirata Ghorbani and James Y Zou. 2020. Neuron shapley: Discovering the responsible neurons. *Advances in Neural Information Processing Systems* 33 (2020), 5922–5932.
- [42] Frederick X Gibbons and Bram P Buunk. 1999. Individual differences in social comparison: development of a scale of social comparison orientation. *Journal of personality and social psychology* 76, 1 (1999), 129.
- [43] L Gierl, B Pollwein, G Heyde, and H Kurt. 1993. Knowledge-based scheduling of duty rosters for physicians. *Medical Informatics* 18, 4 (1993), 355–366.
- [44] Robert C Gray, Jichen Zhu, Danielle Arigo, Evan Forman, and Santiago Ontañón. 2020. Player Modeling via Multi-armed Bandits. In *Proceedings of the 15th International Conference on the Foundations of Digital Games*.
- [45] Robert C Gray, Jichen Zhu, and Santiago Ontañón. 2021. Multiplayer Modeling via Multi-Armed Bandits. In *Proceedings of the 2021 IEEE Conference on Games*.
- [46] Jerald Greenberg. 1990. Organizational justice: Yesterday, today, and tomorrow. *Journal of management* 16, 2 (1990), 399–432.
- [47] Jerald Greenberg and Jason A Colquitt. 2013. *Handbook of organizational justice*. Psychology Press.
- [48] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems*. 3315–3323.
- [49] Sergiu Hart. 1989. Shapley value. In *Game theory*. Springer, 210–216.
- [50] Hoda Heidari, Michele Loi, Krishna P Gummadi, and Andreas Krause. 2019. A moral framework for understanding fair ml through economic models of equality of opportunity. In *Proceedings of the conference on fairness, accountability, and transparency*. 181–190.
- [51] Clara E Hill. 2012. *Consensual qualitative research: A practical resource for investigating social science phenomena*. American Psychological Association.
- [52] Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2017. Fairness in Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*. 1617–1626.
- [53] Brigitte Jaumard, Frederic Semet, and Tsevi Vovor. 1998. A generalized linear programming model for nurse scheduling. *European journal of operational research* 107, 1 (1998), 1–18.
- [54] Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2016. Fairness in learning: classic and contextual bandits. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems*. 325–333.
- [55] Jesper Juul. 2010. *A casual revolution: Reinventing video games and their players*. MIT press.
- [56] Ondrej Kaššák, Michal Kompan, and Mária Bieliková. 2016. Personalized hybrid recommendation for group of users: Top-N multimedia recommender. *Information Processing & Management* 52, 3 (2016), 459–477.
- [57] Jae Kyeong Kim, Hyea Kyeong Kim, Hee Young Oh, and Young U Ryu. 2010. A group recommendation system for online communities. *International journal of information management* 30, 3 (2010), 212–219.
- [58] Tor Lattimore and Csaba Szepesvári. 2018. Bandit algorithms. *preprint* (2018), 28.
- [59] Tricia M Leahey, Melissa M Crane, Angela Marinilli Pinto, Brad Weinberg, Rajiv Kumar, and Rena R Wing. 2010. Effect of teammates on changes in physical activity in a statewide campaign. *Preventive medicine* 51, 1 (2010), 45–49.
- [60] Fengjiao Li, Jia Liu, and Bo Ji. 2019. Combinatorial Sleeping Bandits with Fairness Constraints. In *Proceedings of the IEEE Conference on Computer Communications*. IEEE, 1702–1710.
- [61] James J Lin, Lena Mamykina, Silvia Lindtner, Gregory Delajoux, and Henry B Strub. 2006. Fish'n'Steps: Encouraging physical activity with an interactive computer game. In *International conference on ubiquitous computing*. Springer, 261–278.
- [62] Mu Lin, Nicholas D Lane, Mashfiqui Mohammad, Xiaochao Yang, Hong Lu, Giuseppe Cardone, Shahid Ali, Afsaneh Doryab, Ethan Berke, Andrew T Campbell, et al. 2012. BeWell+ multi-dimensional wellbeing monitoring with community-guided user feedback and energy optimization. In *Proceedings of the conference on Wireless Health*. 1–8.
- [63] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalaya Mandal, and David Parkes. 2017. Calibrated Fairness in Bandits. *CoRR abs/1707.01875* (2017).
- [64] Sisi Ma and Roshan Tourani. 2020. Predictive and causal implications of using shapley value for model interpretation. In *Proceedings of the 2020 KDD Workshop on Causal Discovery*. PMLR, 23–38.
- [65] Regan Mandryk, Kathrin Gerling, Rohit Ashok Khot, Florian Mueller, et al. 2016. Exertion games.
- [66] Judith Masthoff. 2004. Group modeling: Selecting a sequence of television items to suit a group of viewers. In *Personalized digital television*. Springer, 93–141.
- [67] Emery J Max, Stephen Samendinger, Brian Winn, Norbert L Kerr, Karin A Pfeiffer, and Deborah L Feltz. 2016. Enhancing aerobic exercise with a novel virtual exercise buddy based on the Köhler effect. *Games for health journal* 5, 4 (2016), 252–257.
- [68] Joseph F McCarthy. 2002. Pocket restaurantfinder: A situated recommender system for groups. In *Workshop on mobile ad-hoc communication at the 2002 ACM conference on human factors in computer systems*, Vol. 8.
- [69] Joseph F McCarthy and Theodore D Anagnost. 1998. MusicFX: an arbiter of group preferences for computer supported collaborative workouts. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*. 363–372.
- [70] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
- [71] Gina Merchant, Nadir Weibel, Laura Pina, William G Griswold, James H Fowler, Guadalupe X Ayala, Linda C Gallo, James Hollan, and Kevin Patrick. 2017. Face-to-face and online networks: College students' experiences in a weight-loss trial. *Journal of health communication* 22, 1 (2017), 75–83.

- [72] Andreas Messalas, Yiannis Kanellopoulos, and Christos Makris. 2019. Model-agnostic interpretability with shapley values. In *2019 10th International Conference on Information, Intelligence, Systems and Applications (IISA)*. IEEE, 1–7.
- [73] Harvey H Millar and Mona Kiragu. 1998. Cyclic and non-cyclic scheduling of 12 h shift nurses by network programming. *European journal of operational research* 104, 3 (1998), 582–592.
- [74] Hervé Moulin. 1992. An application of the Shapley value to fair division with money. *Econometrica: Journal of the Econometric Society* (1992), 1331–1349.
- [75] Yuri Nakao, Simone Stumpf, Subeida Ahmed, Aisha Naseer, and Lorenzo Strapelli. 2022. Towards Involving End-users in Interactive Human-in-the-loop AI Fairness. *ACM Transactions on Interactive Intelligent Systems (TiiS)* (2022).
- [76] Yehuda Naveh, Yossi Richter, Yaniv Altschuler, Donna L Gresh, and Daniel P Connors. 2007. Workforce optimization: Identification and assignment of professional workers using constraint programming. *IBM Journal of Research and Development* 51, 3.4 (2007), 263–279.
- [77] Ellinor K Olander, Helen Fletcher, Stefanie Williams, Lou Atkinson, Andrew Turner, and David P French. 2013. What are the most effective techniques in changing obese individuals' physical activity self-efficacy and behaviour: a systematic review and meta-analysis. *International Journal of Behavioral Nutrition and Physical Activity* 10, 1 (2013), 29.
- [78] Mark O'Connor, Dan Cosley, Joseph A Konstan, and John Riedl. 2001. PolyLens: A recommender system for groups of users. In *ECSCW 2001*. Springer, 199–218.
- [79] Arjun Panesar. 2019. *Machine learning and AI for healthcare*. Springer.
- [80] Pablo Paredes, Ran Gilad-Bachrach, Mary Czerwinski, Asta Roseway, Kael Rowan, and Javier Hernandez. 2014. PopTherapy: Coping with stress through pop-culture. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*. ICST (Institute for Computer Sciences, Social-Informatics and ...), 109–117.
- [81] Juha Pärkkä, Luc Cluitmans, and Miikka Ermes. 2010. Personalization algorithm for real-time activity recognition using PDA, wireless motion bands, and binary decision tree. *IEEE Transactions on Information Technology in Biomedicine* 14, 5 (2010), 1211–1215.
- [82] Raymond Paternoster, Robert Brame, Paul Mazerolle, and Alex Piquero. 1998. Using the correct statistical test for the equality of regression coefficients. *Criminology* 36, 4 (1998), 859–866.
- [83] Mashfiqui Rabbi, Min Hane Aung, Mi Zhang, and Tanzeem Choudhury. 2015. MyBehavior: automatic personalized health feedback from user behaviors and preferences using smartphones. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 707–718.
- [84] Mashfiqui Rabbi, Angela Pfammatter, Mi Zhang, Bonnie Spring, and Tanzeem Choudhury. 2015. Automated personalized feedback for physical activity and dietary behavior change with mobile phones: a randomized controlled trial on adults. *JMIR mHealth and uHealth* 3, 2 (2015), e42.
- [85] Aida Rahmattalabi, Shahin Jabbari, Himabindu Lakkaraju, Phebe Vayanos, Max Izenberg, Ryan Brown, Eric Rice, and Milind Tambe. 2021. Fair Influence Maximization: a Welfare Optimization Approach. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*. 11630–11638.
- [86] Brianna Richardson, Jean Garcia-Gathright, Samuel F Way, Jennifer Thom, and Henriette Cramer. 2021. Towards Fairness in Practice: A Practitioner-Oriented Rubric for Evaluating Fair ML Toolkits. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [87] Herbert Robbins. 1952. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* 58, 5 (1952), 527–535.
- [88] Lionel P Robert, Casey Pierce, Liz Marquis, Sangmi Kim, and Rasha Alahmad. 2020. Designing fair AI for managing employees in organizations: a review, critique, and design agenda. *Human-Computer Interaction* 35, 5-6 (2020), 545–575.
- [89] Alvin E Roth. 1988. *The Shapley value: essays in honor of Lloyd S. Shapley*. Cambridge University Press.
- [90] Maria Salamó, Kevin McCarthy, and Barry Smyth. 2012. Generating recommendations for consensus negotiation in group personalization services. *Personal and Ubiquitous Computing* 16, 5 (2012), 597–610.
- [91] Christophe Senot, Dimitre Kostadinov, Makram Bouzid, Jérôme Picault, Armen Aghasaryan, and Cédric Bernier. 2010. Analysis of strategies for building group profiles. In *International Conference on User Modeling, Adaptation, and Personalization*. Springer, 40–51.
- [92] Lloyd S. Shapley. 1953. A value for n-person games. *Contributions to the Theory of Games* 2, 28 (1953), 307–317.
- [93] Petr Skobelev. 2018. Towards autonomous AI systems for resource management: applications in industry and lessons learned. In *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer, 12–25.
- [94] Paulina Sockolow, Seran Schug, Jichen Zhu, TJ Smith, Yalini Senathirajah, and Sandra Bloom. 2017. At-risk adolescents as experts in a new requirements elicitation procedure for the development of a smart phone psychoeducational trauma-informed care application. *Informatics for Health and Social Care* 42, 1 (2017), 77–96.
- [95] Constantine D Spyropoulos. 2000. AI planning and scheduling in the medical hospital environment. , 101–111 pages.
- [96] Olov Ståhl, Björn Gambäck, Preben Hansen, Markku Turunen, and Jaakko Hakulinen. 2008. A mobile fitness companion. (2008).
- [97] William R Thompson. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3/4 (1933), 285–294.
- [98] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [99] Jennifer Villareale, Robert C Gray, Anushay Furquan, Thomas Fox, and Jichen Zhu. 2019. Enhancing Social Exergames through Idle Game Design. In *Proceedings of the 14th International Conference on the Foundations of Digital Games*.
- [100] Thomas A Wills. 1981. Downward comparison principles in social psychology. *Psychological bulletin* 90, 2 (1981), 245.
- [101] Eyal Winter. 2002. The shapley value. *Handbook of game theory with economic applications* 3 (2002), 2025–2054.
- [102] Joanne V Wood. 1989. Theory and research concerning social comparisons of personal attributes. *Psychol. Bull.* 106, 2 (1989), 231.
- [103] Elad Yom-Tov, Guy Feraru, Mark Kozdoba, Shie Mannor, Moshe Tennenholtz, and Irit Hochberg. 2017. Encouraging physical activity in patients with diabetes: intervention using a reinforcement learning system. *Journal of medical Internet research* 19, 10 (2017), e338.
- [104] Zhiwen Yu, Xingshe Zhou, Yanbin Hao, and Jianhua Gu. 2006. TV program recommendation for multiple viewers based on user profile merging. *User modeling and user-adapted interaction* 16, 1 (2006), 63–82.
- [105] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna Gummadi. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*. 1171–1180.
- [106] Mo Zhou, Yoshimi Fukuoka, Yonatan Mintz, Ken Goldberg, Philip Kaminsky, Elena Flowers, and Anil Aswani. 2018. Evaluating machine learning-based automated personalized daily step goals delivered through a mobile phone app: Randomized controlled trial. *JMIR mHealth and uHealth* 6, 1 (2018), e28.
- [107] Jichen Zhu, Diane H Dallal, Robert C Gray, Jennifer Villareale, Santiago Ontañón, Evan M Forman, and Danielle Arigo. 2021. Personalization Paradox in Behavior Change Apps: Lessons from a Social Comparison-Based Personalized App for Physical Activity. In *ACM PACM on Human Computer Interaction*.
- [108] Jichen Zhu, Yuanyuan Feng, Anushay Furquan, Robert C Gray, Timothy Day, Jessica Nebolsky, and Karina Caro. 2018. Towards Extending Social Exergame Engagement with Agents. In *Proceedings of the 21st ACM Conference on Computer-Supported Cooperative Work and Social Computing*.
- [109] Jichen Zhu and Santiago Ontañón. 2019. Experience management in multi-player games. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–6.
- [110] Jichen Zhu and Santiago Ontañón. 2020. Player-centered AI for automatic game personalization: Open problems. In *International Conference on the Foundations of Digital Games*. 1–8.
- [111] Jichen Zhu, Jennifer Villareale, Nithesh Javvaji, Sebastian Risi, Mathias Löwe, Rush Weigelt, and Casper Hartevelde. 2021. Player-AI interaction: What neural network games reveal about AI as play. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17.