I'll Play on My Other Account: The Network and Behavioral Differences of Sybils

FRED MORSTATTER, USC Information Sciences Institute, USA
DO OWN (DONNA) KIM, USC Annenberg School for Communication and Journalism, USA
NATALIE JONCKHEERE, USC Annenberg School for Communication and Journalism, USA
CALVIN LIU, USC Annenberg School for Communication and Journalism, USA
MALIKA SETH, USC Information Sciences Institute, USA
DMITRI WILLIAMS, USC Annenberg School for Communication and Journalism, USA

This article studies the effects and implications of sybils (secondary accounts created by a person in an online platform) through the game *World of Tanks* from interdisciplinary, mixed-methods perspectives. Considering sybils allows us to access a "person based" network, instead of an "account based" network, revealing formerly undetected patterns. We move on to behavioral differences between the "parent" (the initial account) and "child" (those created afterwards) accounts in a sybil relationship. We explore the behavioral patterns of sybils using network, chat, and gameplay data. We find that sybils represent players experimenting with new roles or features without damaging their play record. We find that there are significant behavioral differences between different sybil accounts, and we leverage them to build a machine learning classifier to differentiate sybils. This classifier is able to identify sybil accounts with over 95% accuracy for sybil/non-sybil, 61% for parent/child. This study demonstrates the underexplored but rich potential for sybils to improve research and industry practitioners' understandings of user practices and experiences.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; *Empirical studies in collaborative and social computing*; • **Computing methodologies** → *Machine learning approaches.*

Additional Key Words and Phrases: sybils, social networks, topic modeling, machine learning, massively multiplayer online game (MMOG), self-presentation

ACM Reference Format:
Fred Morstatter, Do Own (Donna) Kim, Natalie Jonckheere, Calvin Liu, Malika Seth, and Dmitri Williams. 2021. I'll Play on My Other Account: The Network and Behavioral Differences of Sybils . *Proc. ACM Hum.-Comput. Interact.* 5, CHI PLAY, Article 277 (September 2021), 18 pages. https://doi.org/10.1145/3474704

277

1 INTRODUCTION

Sybil accounts are multiple accounts created by a single user on a platform. In the context of social networks, they are often created to disproportionately influence the impact that a single

Authors' addresses: Fred Morstatter, USC Information Sciences Institute, 4676 Admiralty Way Ste. 1001, Marina del Rey, CA, USA, 90292, fredmors@isi.edu; Do Own (Donna) Kim, USC Annenberg School for Communication and Journalism, 3502 Watt Way, Los Angeles, CA, USA, 90089, doownkim@usc.edu; Natalie Jonckheere, USC Annenberg School for Communication and Journalism, 3502 Watt Way, Los Angeles, CA, USA, 90089, njonckhe@usc.edu; Calvin Liu, USC Annenberg School for Communication and Journalism, 3502 Watt Way, Los Angeles, CA, USA, 90089, liucalvi@usc.edu; Malika Seth, USC Information Sciences Institute, 4676 Admiralty Way Ste. 1001, Marina del Rey, CA, USA, 90292, malikase@isi.edu; Dmitri Williams, USC Annenberg School for Communication and Journalism, 3502 Watt Way, Los Angeles, CA, USA, 90089, dcwillia@usc.edu.

2573-0142/2021/9-ART277 $15.00
https://doi.org/10.1145/3474704

person has on the site [37], often working in coordination to change users' perception towards or against a certain user or concept. On eBay, for example, sybils have been used to maliciously inflate sellers' ratings [41]. Another of the more negative uses of sybils is to allow an experienced player masquerading as a new player to be matched up with other newbies unfairly. The experienced player will generally beat up the new ones, inflating the sybil's ego, but discouraging the new players from continuing. In addition to breaking the general social norm of fair play [14], this also has obvious financial repercussions for a company that is more likely to lose a potential player or customer who is not enjoying themselves [17].

Sybils' purposes are not always purely nefarious, however. There are at least two harmless and even beneficial reasons that a user would want to create a second account, especially in online gaming. One is that most online games tend to keep a record of users' play history, such as win rates. A user may choose to create a new account to wash away their less-than-stellar win rate from the period in which they were learning how to play the game. The other is that a person may want to assume a new identity for a wide variety of reasons, ranging from compartmentalizing their relationships [34] to experimenting with their identity [35].

Each of these three reasons will have generalizability across many, and perhaps almost all, games. Any title that has competition and multiplayer functions introduces the incentive for experienced players to take advantage of inexperienced players. Even skill-based matchmaking systems designed to allow for fair play will struggle to properly classify a sybil with no record. Any title that has public win rates introduces the incentive to erase a record. Any title that has other players introduces the possible desire to manage multiple identities. Taken collectively, we should expect some degree of sybil account to appear in any multiplayer game, and especially in competitive ones. Multiplayer titles are increasingly ubiquitous; a recent estimate found that only 65% of players play with others, and 30% have met a good friend or significant other through play.[1] As a result, accounting for sybils is critical.

Understanding which accounts are sybils is important for both researchers and developers. For researchers, sybils create artifacts in networks. Unless sybils are identified, these users will be represented by several nodes, which will ultimately skew common network statistics. For game developers, understanding sybils is important because it will lead to a greater understanding of the desires and priorities of their games' players.

In this work, we study sybil accounts on *World of Tanks*, a massively multiplayer online vehicle battle simulator developed by Wargaming.net. Specifically, we explore the social dynamics of sybils as they are reflected through their gameplay and chat behavior. First, we investigate several motivations for sybil creation. We measure the extent to which these phenomena occur in the data. Next, we move on to behavioral differences between the "parent" (i.e. the initial account) and "child" (i.e. those accounts created afterwards) accounts in a sybil relationship. We demonstrate that there are statistically significant differences in behavioral characteristics of the users, in terms of the games they play and the languages they use. Finally, we show that these differences can be applied to a learning algorithm to differentiate parent and child accounts.

We investigate the following research questions:

(1) What are the behavioral distinctions between users who begin sybil accounts and those who do not? Furthermore, what behavioral differences exist between the parent and child accounts in a sybil relationship?
(2) Which behaviors are specific to different types of sybil accounts?
(3) Can we leverage these differences to identify sybils automatically?

[1] https://www.theesa.com/resource/2020-essential-facts/

Data Statement. The data used in this study were provided by Wargaming.net in an encrypted, anonymized format without any access to personally identifiable information. The Institutional Review Board of the researchers' university approved the study.

2 BACKGROUND

In this section, we discuss the motivations for creating sybil accounts, as well as studies that detect them. We introduce the game we study in this work, *World of Tanks*, and the data we used in our analyses.

2.1 Managing Sybil Accounts

Sybil accounts are defined as multiple accounts created by a single user, typically known as "smurf" or "alt" accounts in the gaming culture. We use the term sybil due to the implications that other similar terms may have. For example, while smurfing is a popular term in gaming culture, the term smurf tends to carry a negative connotation that is primarily associated with deception and the desire to win over less experienced players. Similarly, the term alt implies a secondary relationship to other accounts, whereas the relationships between accounts may be more nuanced. As such, we settle on the term sybil as a more holistic term for defining accounts shared by a user.

People may manage sybil accounts for a mix of diverse yet not mutually exclusive reasons, including performing different versions of self, accessing different groups, and gaining system benefits. The first reason why an individual may manage multiple accounts is impression management. Goffman [11] argues that there are two activities involved in impression management: one that they give and one that they give off (p. 2). According to Goffman, individuals will attempt to shape the impression others have of them by selectively presenting their performance of self, i.e. "front stage" behaviors. Their "true selves" are better reflected by their "back stage" behaviors, which may be available only to a narrower group of people with whom the individual feels relaxed and comfortable. Following Goffman, people may manage multiple accounts to perform on multiple "front stages" or in order to co-manage them with a "back stage" which they will share only with a selective group of people [3, 16, 18]. Privacy is another concern. In their study of Twitter, Marwick and boyd [21] explained that users may wish to use multiple accounts to negotiate with "context collapse," which refers to the "flattening of multiple relationships into one setting" (p. 122). In order to circumvent the limitations of existing privacy settings, these individuals create different accounts to assign to different privacy contexts [1, 18, 24, 27]. In the context of games, another reason could be the user's competitive motivation to dominate and to be perceived to have a dominating "front stage" (see [15, 39]). That is, the skilled players may create sybils to easily win against less skilled players, known colloquially within *World of Tanks* as "seal clubbing" [4], or to explore new features without negatively affecting their parent account's gameplay records, such as win rate.

Within computer science and social networks, sybils are often viewed as accounts used in addition to the primary account to bolster the opinions or social status of a malicious user. For instance, a "sybil attack" refers to when multiple accounts are created to manipulate results in a system [7, 24]. The majority of past work has focused on how to detect sybils, and many of the techniques have relied on friendship networks. For instance, Yang et al. [37] analyzed sybil accounts on the Chinese social networking site Renren and found that sybils tend not to connect to each other. In another study focusing on Renren [38], the authors found that sybils tend to generate many unsolicited social edges in the network in an attempt to bolster their credibility.

Sybil accounts are gaining limited but growing attention in games research. Studies show that players use sybils in a variety of contexts, including for smurfing [17] and in guilds [26]. Players may also use sybils out of a desire to freshly experience the game through a new account rather than dwelling on the "endgame" content [23]. Certain design elements may also affect why players create sybils; for example, fixed character jobs or designs, storage space, and additional fees for creating new characters can encourage or discourage the creation and management of sybil accounts in MMOs [6]. Sybils have been considered methodologically as well. In ethnographic studies, they have been used as the means for conducting field research (e.g. [23]). In survey research, some research has acknowledged sybils by considering only the user's most actively played, or "main" character (e.g. [35]).

In fact, not accounting for sybils in games research can lead to potentially problematic generalizations. Leavitt et al. [19] discussed the inconsistencies and the general lack of consideration regarding sybils in research practices as well as what issues they may incur, such as overlooked relationships and duplication of individuals in samples (see also [23, 35]). Moreover, as games and virtual world researcher Yee [40] argued, online contexts and behaviors can powerfully affect those offline, and vice versa. That is, missed linkages between the individual and their sybils can create discrepancies in accurately investigating the consequences in or from the user's multiple online engagements.

Sharing similar concerns, Leavitt et al. [19] investigated sybils by studying the server logs of *EVE Online*, where players could create upto three characters within an account. While they observed the management of sybil characters *within an account* not *across* sybil accounts, their contribution was a pioneering exploration of how people manage sybils *at scale*, expanding the emerging understandings on why. Regarding prevalence, for instance, they reported that although around 65% of users played with a single character, the total number of characters created by a user, including deleted characters in the past, followed a long-tail distribution. However, there are some contextual factors that limit their translatability to the *World of Tanks* context, some of the major ones being *EVE Online*'s genre (MMO*RPG*) and the ability to create multiple characters within an account. Yet, they nonetheless illustrated the richness and importance of adopting an user-based approach that collapses sybils into a user, as opposed to considering each sybil as a separate individual.

2.2 Detecting Sybil Accounts

There are several existing approaches to identify sybil accounts. Based on their findings, Yang et al. [38] developed a model to detect malicious sybils based upon their friend request patterns. Similarly, Gong, Frank, and Mittal [12] developed a belief propagation approach called "SybilBelief" that can predict sybil accounts based upon the neighborhood of each node. Others have attempted to take a more contextual approach. For example, Wang et al. [33] mined the clickstream log of users and clustered similar clickstream behavioral patterns to detect a preponderance of sybils in the clusters. Wang et al. [32] developed an algorithm based upon random walks to identify sybils in online social networks. In a similar vein, Gao et al. [10] uses loopy belief propagation to identify malicious sybils in social networks. A similar method is employed in Wang et al. [31] where posterior reputation scores are learned for a given node in the graph. These approaches resemble a similar approach by Yuan et al. [42] on WeChat where sybil malicious users are identified based upon features from their network and IP address. Similarly, Xia et al. [36] build a deep network-based approach to identify sybils in a social network using features from the friendship formation patterns and IP addresses.

Wargaming.net has developed its own approach to detect sybils, which serves as the ground truth for this study. This approach combines two internal approaches to detecting sybils. The first method is known as the session ID approach. Within the session ID approach, an account's user is identified through computer hardware. When an account logs in, Wargaming.net collects identifiers such as the MAC address and processor. Should a newer account share identifiers with an older account, then the newer account is marked as a sybil. The second method is the Trigram Distance approach, which inspects user email addresses. It compares a player's email address with all other registered email addresses before it and designates sybil status based on the level of similarity between email addresses. Sybil relationship data were provided by Wargaming.net in a hashed, encrypted format. We do not focus on the detection process itself in this project. Given that we were able to effectively skip the detection process, we instead focus on the implications and effects of sybils on the network and analysis. We acknowledge that there may be edge cases (e.g., users sharing devices) that add noise to the labels. Because of this, our analysis also considers the behaviors of each account.

2.3 *World of Tanks*

Our analysis centers on a network of players in *World of Tanks*. In the game, players select a tank and enter a battle with other players; the object of the game is to destroy all of the enemy tanks or to capture the enemy base. Within each team, players perform different functions, often based on the type of tank they choose to play.

World of Tanks features several different modes, ranging from random battles to skirmishes and other clan-based battles that allow players to select other players to be on their team. These battle modes enable different types of relationships among players, from random encounters to long-term teams. For example, players can choose to play with one or two others in "platoons" in random battles or form their own teams of seven, 10, or 15 in other battle modes.

World of Tanks also allows players to communicate with each other during battles. This can occur through text-based chat or voice, though many teams that prefer the latter tend to use external applications. Text-based chat can take on a range of functions, some of which are strategic and others more affective.

2.4 Data Used in this Study

In this study, we have three sets of data, each of which provides an independent view into the behavior on the network. Wargaming.net maintains many servers, each of which operate as their own universe for *World of Tanks*. Our data come from the North American server. The three sets of data we use are:

1. **Co-Play.** Co-play games are unranked, ad hoc games that occur between players of *World of Tanks*. Every time two users play in a co-play game, an edge is created between these two users. Co-play games can consist of multiple users, and a commensurate number of links are created for each game. This is a timestamped edge list of "co-plays" on *World of Tanks*. This dataset consists of 455,412 players, and 16,084,804 games that consist of 64,621,492 edges. The dataset covers the period from February 2016 to March 2019.
2. **Chat.** *World of Tanks*, like most online games, allows for in-game chat between the players. These data consists of all of the 63,962,134 chat messages. The primary languages are English, Russian, French, Spanish, and Portuguese.
3. **Warehouse Data.** In addition to networks, behavioral features have been shown to be important in gaming contexts [25]. Our analysis considers other behavioral features of *World of Tanks* player accounts using data from Wargaming.net's data warehouse, which included gameplay elements. Some warehouse features involve the tanks that users play; for context, it is important to understand how tanks are used in the game. A user can have multiple tanks in their inventory at a given time. They select a tank for each battle and use that tank for the duration of the battle. Every tank has a tier, which is an integer ranging from 1 to 10, inclusive. The lowest tier is 1 and has the least amount of power. The highest tier is 10 and has the most power. Progression in the game will unlock higher-powered tanks. Within each tier, there are different types of tanks that each perform a different function in a battle.

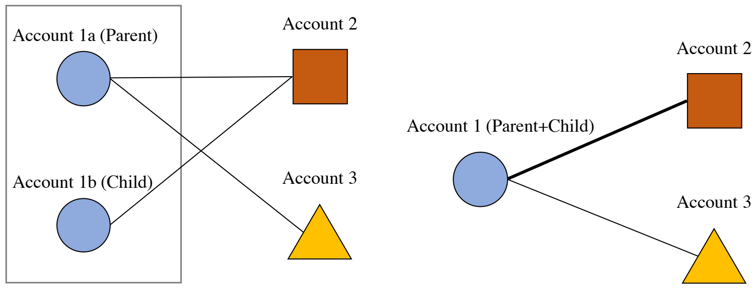

(a) Pre-Collapse. (b) Post-Collapse.

Fig. 1. A simple diagram demonstrating how we collapse networks to merge sybil accounts. Account 1a and Account 1b are owned by the same user. We combine them into one meta-node, and aggregate their edges in the collapsed network.

The warehouse data consist of the following features:

- The number of battles played.
- The number of battles won.
- The amount of experience points they earned across all of their battles.
- The maximum skill level of a tank they used in any battle.
- The tank skill diversity. We construct a distribution of tank skill levels from each battle the user plays. We then compute the Gini index [9] of this distribution.
- The number of days that the user engaged in at least one battle.
- The user's global rating on the server.

3 BEHAVIORAL ANALYSIS OF SYBIL USERS

This study focuses on the implications and effects of sybils on the network and behavioral analysis. We largely focus on two core differences: between sybil and non-sybil accounts, and between parent and child sybil accounts. These terms are defined as follows:

- **Non-sybil account:** An account that is the sole account of a user.
- **Sybil account:** An account that is part of a group of accounts used by a single user.
  - **Parent account:** The initial account in a group of accounts used by a single user.
  - **Child account:** The subsequent account(s) in a group of accounts used by a single user.

As a preliminary assessment of the effect of sybils on the network, we studied the effect of "collapsing" these users. In essence, we turned this from an "account-based" network to a "person-based" network by merging all of a person's sybil accounts into one meta-node. This helps us to understand the extent to which sybils distort network metrics. For instance, if an individual manages multiple accounts, it is possible that the individual is connected to nodes across multiple accounts. Alternatively, sybil accounts may have connections to nodes unique to each sybil. Collapsing sybil accounts would aggregate these patterns into one common meta-node: the individual managing the account. An example of how this collapsing might occur is shown in Figure 1. Aggregating these patterns provides communication and user experience research with key network insights, such as the diffusion of messages and behaviors, influences, and psychology of sociability.

Our exploration focused on the dynamics of sybils through three analyses. The first two employ network analysis, where (a) we studied the effect sybils have on the global network and its properties, and (b) we studied the effect of merging sybils at the ego network level. For the purposes of

Statistic	Pre-Merge	Post-Merge	Pct. Diff.
Nodes	169,149	159,787	-5%
Edges	4,233,190	4,202,663	-0.7%
Global Clustering Coeff.	6.89%	6.81%	-1.2%
Density	$2.9\text{x}10^{-4}$	$3.2\text{x}10^{-4}$	9.2%

Table 1. Network statistics pre- and post-merge of sybil accounts. The percent change ("Pct. Diff.") represents the change of the statistic as a result of the merge.

constructing the social network for this study, we used co-play data, which excluded random, "by-chance" team formation and only included instances where players made some kind of active choice to play with another specific player. Then, we analyzed topics based on chat data, which we discussed in connection with network patterns.

3.1 Global Network Analysis

For each account in the network, we looked up the sybils in the ground truth dataset that Wargaming.net provided to us and merged that account with all of its sybils into one meta-node. We computed some standard measures on the graph both pre- and post-merge. The impact this process had on the network can be seen in Table 1, where we conducted our collapsing approach on the graph. In particular, there is a disproportionate increase in density, which suggests a low overlap between the ego networks of sybil accounts shared by one person. This suggests that players with sybil accounts use their parent and child accounts to interact with different sets of players within the game.

Furthermore, this suggests that players use each account to address a different motivation or need, such as identity management or privacy concerns. Moreover, this means that an individual "person," not the "account," can reach a wider range of other users than an account-based network shows. This is important for system managers and researchers alike; in particular, it highlights how what has been previously considered a person's reach (which is that of an "account") does not reflect the person's actual reach (based on the "person"-based network). As a result, considering sybils can expand our understanding of how influence and information travel.

3.2 Ego Network Analysis

Ego networks have been shown to provide behavioral signal, especially in the context of gaming [22]. For each user, we constructed an ego network from the account's 1.5-hop network (the user's neighbors, as well as edges among those neighbors). We extracted this from the co-play dataset, which is a record of who battles with whom. This network is undirected. From this, we constructed an ego network for each user. We then extracted several features of the ego network.

(1) **Degree.** The degree of the user, calculated as the number of accounts the user has played with.
(2) **Weighted Degree.** The weighted degree of the user, where edge weight is determined by the frequency of battles. Essentially, this is the number of battles the user has played.
(3) **Weighted Degree Ratio.** The weighted degree divided by the degree, which is the average games per neighbor.
(4) **Density.** The density of the graph, which is the number of edges divided by the total possible edges.
(5) **Clustering Coefficient.** The number of triangles in the graph.

(6) **Components without User.** The number of connected components in the graph after removing the ego node.

Outside of this, we added two additional features from the two-hop graph of the user. The two-hop network is defined as the friends and friends-of-friends of the ego node. From this graph, we computed these features:

(1) **Two-Hop Size.** The number of vertices in the user's two-hop network.
(2) **Embeddedness.** The average Jaccard score [43] between the user's one-hop network and their neighbors' one-hop network.

We investigated the differences between two separations of users. First, we explored the differences between the ego networks of users in a sybil relationship and those not. Next, we explored differences between ego accounts of parent and child sybil accounts. We conducted a z-test between the distributions. Since we performed multiple tests, we compensated by performing a Bonferroni correction [30]. The results of these z-tests are shown in Table 2. These results are important because they show that nearly all features are strongly significantly different between sybils and non-sybils and between parent and child accounts. This implies that the network structures are vastly different across groups. We find that all of our features are significant between the two setups, except for density in the parent/child comparison. This indicates that there is a difference in network activity among sybil accounts, and between parent and child accounts.

To better understand these differences, we studied the difference in one particular metric: degree. The degree distribution is a common way to view the properties of a network [8], and we adopted it here to study the characteristics of these different subsets of users. The results of this are shown in Figure 2(a). The degree distributions of sybil users and non-sybils are very different, with users in sybil relationships having many more network connections than those without. That is, sybil accounts tend to connect with more accounts than non-sybils. They are possibly interested more in widening their network than in focusing on a concentrated, smaller network. However, this alone does not provide context for the differences in degree, which we explored through topics (see Section 3.3). Among those sybil users, parents are more prevalent. This is an artifact of our filtering approach, which requires that a user play at least one co-play game and post one chat message. Figure 2(b), which looks at parent and child accounts, exhibits a similar distribution to that of Figure 2(a). This means that parent accounts tend to have more connections in-game than child accounts and that child accounts tend to have a smaller number of co-players than parent accounts. This could be because child accounts were created at a later stage, but the presence of child accounts with high degrees suggest that this is likely not the only explanation. There are some cases in which child accounts exceed parent accounts at the higher degree range, which suggests the possibility that some child accounts have overtaken the dominance of parent accounts, despite being created later.

A user's ego network structure differs significantly based upon their sybil status. In a subsequent section, we use these network features to predict the user's sybil status.

3.3 Topics

Topic modeling is an unsupervised machine learning technique that identifies topics in a large corpus. Topics are abstract, with each topic represented as a distribution over the corpus vocabulary. High probability words are highly associated with the topics. This study used latent Dirichlet allocation (LDA) [2] to identify topics in the chats of *World of Tanks* players. Here, when we say "topic" we mean a statistical topic [2]. That is a probability distribution over all of the words in the corpus vocabulary. When we say a model has (e.g.,) 10 topics, that means that there are 10 different probability distributions over the vocabulary. Because the distributions differ, the most

Table 2. Bonferroni-corrected p-values for z-tests of two sets of distributions. "Sybil Relationship" is the set of users who are in a sybil relationship or not. "Parent/Child" is the set of users who are parents and children in a sybil relationship.

Feature	Sybil/Non-Sybil	Parent/Child
Degree	$< 10^{-10}$	$< 10^{-10}$
Weighted Degree	$< 10^{-10}$	$< 10^{-10}$
Weighted Degree Ratio	$< 10^{-10}$	$< 10^{-10}$
Density	$< 10^{-10}$	0.38
Clustering Coefficient	$2.7x10^{-7}$	$< 10^{-10}$
Components without User	$< 10^{-10}$	$< 10^{-10}$
Two-Hop Size	$< 10^{-10}$	$< 10^{-10}$
Embeddedness	$< 10^{-10}$	$< 10^{-10}$

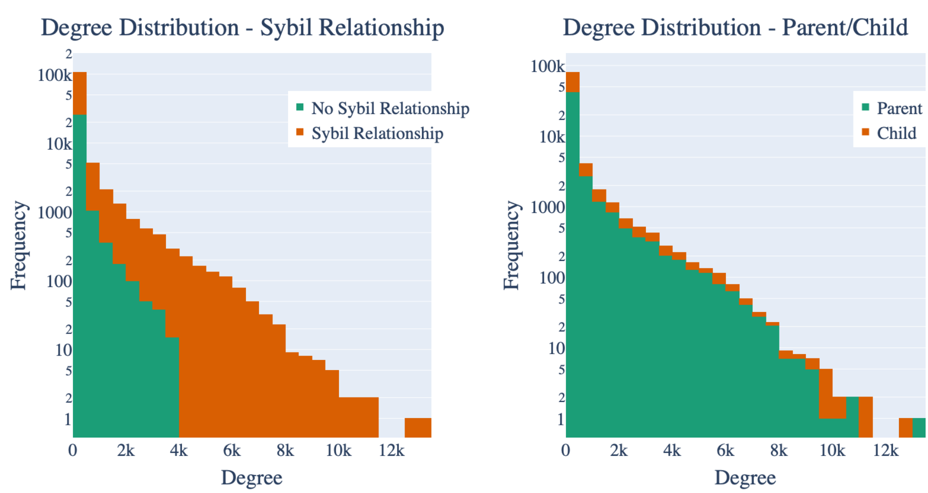

(a) Degree distribution of users in sybil relationships vs those not.

(b) Degree distribution of parents and children in sybil relationships.

Fig. 2. Degree distribution of different sets of users. On the left, the difference between users in a sybil relationship. On the right, the difference between parent and child sybil accounts. Frequencies are represented by the fraction of users in each pin. Note: there are more parents than children. This is because we filter out people who never posted a chat message, or never played a co-play game.

probable words from each topic will differ as well. As is commonly done [5], we investigate the top 10 words from a topic. The topics were learned through the LDA implementation provided in the "gensim" Python library [28]. Default α and β hyperparameters were used. In preprocessing, we converted the text to lowercase and we deleted any word containing fewer than three or more than 20 characters. We removed words that were a repeated sequence of a single character (words like "aaaaaa" were frequent in the raw data). Moreover, we removed any stopword that appears in the RANKS NL[2] stopword list for English, Russian, Spanish, Portuguese, or French. LDA models with 10, 25, 50, and 100 topics were assessed for interpretability by human coders. This was done

[2] http://www.ranks.nl/stopwords

by evaluating the top 10 words from each LDA topic. Based upon this assessment, 50 topics were determined to be the most interpretable. This set of 50 topics will be used throughout the rest of the paper. The mined topics were used as a lens into the experience of the players, and each user's affinity for each topic was used as a feature to predict their sybil status.

In this section, we analyze the set of 50 topics to gain further insight into user behaviors. For each language, we translated the meaning of the words into English. The authors either spoke or consulted speakers of some languages to confirm the translation of some words (e.g., Spanish, French, Arabic). Others were translated using both Google Translate and a series of contextual searches (e.g., including the original word, the translated meaning, and contexts such as game to confirm the accuracy of translation) to confirm their meaning. Words that were ambiguous, regardless of the accuracy of the translation, were marked as so. Three researchers then thematically coded the topics into six categories through a qualitative process [29] that involved three rounds of individual analysis followed by collective evaluation, discussion, and agreement. After the initial discussion of potential themes using randomly selected topics, the researchers engaged in an individual process of open-coding a portion of the data using the initially identified themes as guidance. During the collective session, the researchers discussed individual results segment-by-segment to evaluate the identified themes and to affirm confidence in the consistency of the analysis. For instance, the tone of the topics was dropped as an independent category during the iterative collaborative process because the researchers found it ambiguous and less representative of the larger data. Likewise, interpretative codes like "tank-related directions," "serious gameplay," and "commands" were unified under the "Utilitarian/Tactical" pattern code after the researchers agreed on their similarity. Ultimately, six categories, including a miscellaneous category for ambiguous topics, were identified and applied. More than one category was selected if deemed relevant. The six categories are (see also Table 3):

(1) **Utilitarian/Tactical:** Topics that contained mostly coordination of in-game actions or objectives.
(2) **Gameplay:** Topics that were focused on *World of Tanks* but were not meant for utilitarian/tactical coordination.
(3) **Conversational:** Topics that were focused on items outside of *World of Tanks*, such as general greetings, colloquial speech, and affective conversations.
(4) **Insults/Hate Speech:** Topics that focused on expressions of frustration or aggression towards other players.
(5) **Clan:** Topics that were focused on clan (organized player communities) aspects of the game.
(6) **Miscellaneous:** Topics that did not conform to the above categories, these included foreign language topics with unclear translations, topics with unclear themes, or esoteric topics.

3.4 Topic Interpretation

Topic analysis revealed differences in the type and quality of socialization that occur between sybil and non-sybil accounts and between parent and child accounts. Most notably, compared with non-sybils and parent accounts, sybil and child accounts tended to focus more on topics that directly concern the game itself than on those that imply a potential interest in forming personal, prolonged relationships. Additionally, parent accounts and non-sybil accounts did not display significant differences in the types of topics with which they engage, whereas parent accounts displayed different affinities for different topics than child accounts. Non-sybil accounts also displayed different topic affinities than child sybil accounts. Therefore, the creation of a child account may serve as a catalyst for changes in chat use patterns. Players may create child accounts with specific intentions that diverge from the goals of a parent account. These intentions may

Category	Languages	Sample Topic Words
Utilitarian/Tactical, Gameplay	English	arty (artillery), spot, push, move, kill
Utilitarian/Tactical, Gameplay	English	spotted, reloading, m53/m55, 44.11s
Utilitarian/Tactical, Conversational	Dutch, English	cheers, go*, ciao, hello*, play*
Insults/Hate Speech	English	team, idiot, noob, stupid, trash
Clan	English	clan, Discord, https, recruiting, join
Miscellaneous	Filipino, Arabic	here*, because*, one hundred one hundred*

Table 3. Exemplar topics that demonstrate the diversity of the data. Words translated to English are marked with an asterisk.

represent behaviors that would otherwise be too risky or difficult to perform on a parent account, such as playing against lower level players or trying out different roles or personas. For instance, sybil accounts may provide pseudo-anonymity to discuss topics with which players do not want to be personally associated, such as insults and hate speech, although the current topic patterns suggest that conversational topics, including political topics, tend to appear primarily among non-sybils and parent accounts. The different types of socialization between sybils and non-sybils also supports this idea.

In the following section, drawing on our topic analyses, we discuss (a) the similarities between non-sybil and parent account behaviors and between sybil and children accounts, (b) the types of socialization they engage in, and (c) their language patterns, which we then (d) compare with the network results.

3.4.1 Types of socialization. Topic analysis found that insults and utilitarian speech were common across all groups. Utilitarian speech was likely encouraged by the game itself. As a team game, success in *World of Tanks* incentivizes effective tactical communication between players. Alongside utilitarian speech, insults may have also constituted a way for players to overcome language barriers.

Non-sybil accounts were more likely to engage in conversational and clan-based topics compared to sybil accounts. Topics closely affiliated with sybil accounts, identified through affinity scores, did not include any topics categorized as "Conversational" through words such as "cheers" and "ciao" nor those categorized as "Clan" through words such as "clan" or "Discord" (see Table 3 for examples). This points to non-sybil accounts being more invested in building social relationships between other players. Non-sybil accounts may possess more time and resources to devote to forging relationships, as they do not split time between different accounts.

Similarly, parent sybil accounts are more likely to engage in conversational and clan-based topics compared to child sybil accounts, whereas child sybil accounts are focused more on conversations about gameplay. Alike the aforementioned pattern observed with sybil accounts, topics closely affiliated with child sybil accounts did not include any topics categorized as "Conversational" nor those categorized as "Clan" (see Table 3 for examples). Because parent accounts are older, they have had more time to invest in social relationships with other players, which may explain why parent accounts display more affinity for topics that foster social relationships.

The similarities between parent accounts and non-sybil accounts suggest that sybils act as a means for trying out new behaviors. It may be appropriate to consider non-sybils as pre-sybils, in other words, accounts that have not yet encountered a need or desire to diverge from their current behaviors.

3.4.2 Languages. Topic analyses reveal that players on the North American server of *World of Tanks* communicate in a number of languages, including but not limited to English, Russian, French, Spanish, and Portuguese. In addition, topics commonly mix several of these languages, which could be a result of several factors, including multilinguality of users and differences in parent/child accounts and in sybil/non-sybil accounts.

To begin, the top five topics ranked by both the difference between parent and child accounts and between all accounts and sybil accounts were solely in English. In general, this should not be surprising, as the data are from the North American server, in which a number of players are from the United States or Canada. These topics reveal a mix of players communicating about the gameplay (as in giving directions to teammates) and lobbing insults at teammates. The lack of affective or conversational chat suggests that these English-speaking players do not use the chat for socialization; rather, they are focused on winning.

Because these top five topics are in English, it follows that these English-speaking players are more likely to create child accounts and thus become sybil accounts. Patterns for speakers who use or mix other languages are not as clear. However, in general, English-only topics have greater child than parent affinity, but mixed-language topics have greater parent than child affinity. This could suggest that non-native English speakers use their child accounts to play in English and their parent accounts to play in their native language.

3.5 Combining Networks and Topics

Thus far, this analysis has considered network results and topic analyses separately. The network results reveal the number of co-players a player has; this could suggest the quantity of socialization, i.e. someone who plays with more co-players has more connections in-game and could socialize more. By contrast, the topic analyses reveal the type of socialization. Considering these two analyses together reveals differences in behavior between sybil and non-sybil accounts and between parent and child accounts.

In general, sybils have more in-game interactions, but these interactions are more utilitarian. Network analyses reveal that sybil accounts play with more co-players (see Figure 2), but topic analyses reveal that sybils are less likely than non-sybils to engage in friendly chat. This suggests a difference between the quantity and quality of communication within the game; players with only one account (non-sybils) may not interact with as many players as those with multiple accounts, but they are more likely to develop friendly relationships with some players with whom they play.

Parent and child accounts also reveal different patterns; the parent accounts not only have more connections in-game but are also more likely to participate in affective or conversational chat. In general, parent accounts also play with more co-players than child accounts. One reason for this could be because these accounts were created first and therefore have a longer gameplay history. Another reason could be social; because parents are also more likely to engage in friendly chat, they may also be inclined to play more and thus interact with more players. Child accounts, by contrast, have fewer co-players but a higher degree; in these cases, it is possible that players moved from parent to a child account, so the child account becomes the dominant account. There are several potential motivations for this. Some players may want to erase poor performance history and create a child account to start fresh. Others may feel inclined to use their child account to masquerade as an inexperienced player [4].

4 PREDICTING SYBIL ACCOUNTS

So far, we have discovered that there are statistically significantly different patterns in the behavior of parent and child sybil accounts. For example, the way that they interact with topics and the network structures formed by their interactions with other users are different. Based on these findings, we hypothesize that machine learning techniques will be able to differentiate users in sybil relationships from those who are not and to further differentiate the child role from the parent role. In this section, we investigate this hypothesis by comparing the predictions of our supervised machine learning classifier with the list of sybil accounts that Wargaming.net provided, operationalizing our observations as features to a machine learning classifier.

4.1 Representation of User Accounts

Driven by our analysis, we developed a feature representation of the different accounts. We only consider users who played at least one co-play game *and* posted at least one chat message. Based upon their network interactions and behavior in the game, we computed several features that may be predictive of sybil behavior. This representation is as follows:

(1) **Network Features.** To obtain a higher-level representation of the graph structure, we leveraged recent techniques in graph embedding. We used node2vec [13] to represent the network. We used the PecanPy [20] implementation, which scales well to large graphs. This yielded 128 network embedding features from the co-play network data.
Additionally, we used all of the ego network features discussed in the "Ego Network Analysis" section, yielding eight additional features.
(2) **Behavioral Features.** These data consist of the warehouse features. An exhaustive list of warehouse features is enumerated in Section 2.4.
(3) **Text Features.** We incorporated the topics as features. From LDA, we used the user-topic matrix as a proxy for the affinities each user has for each topic. This is represented as 50 features.

In the next section, we investigate the predictive power of these features at two tasks related to the use of sybils. We begin by incorporating *all* of these features into our classifier. In other words, we have the 128 network embedding features, the eight ego network features, the seven warehouse data features, and the 50 LDA topic features for a total of 193 features. In our ablation tests we test combinations of feature groups, where groups correspond to the groups named in the previous sentence. Experiments were conducted on a labeled dataset of 89,880 users who engaged in co-play and chat.

4.2 Sybil Relationship Prediction

Understanding whether or not the account is in a sybil relationship helps contextualize the behavior of the user and in absence of ground truth can determine who is participating in a sybil relationship. This is a binary classification task where the positive class is "in a sybil relationship" (precision, recall, and $F1$ are computed with respect to the positive class). We did not consider whether the user is the "parent" or "child" in the sybil relationship. This echoes the tasks pursued in [36, 42], however with the important caveat that we do not use the same information. For example, IP addresses were used in both of the aforementioned approaches but are not available to the researchers in this study. We trained varying feature combinations to study which are the most effective. In all cases, we used a random forest classifier with 100 trees. To prevent overfitting, we limited the max tree depth to 2, and required that there be at least three instances on each leaf. We employ cross-validation, randomly withholding 30% of the data for testing. All values reported are computed from this held-out set.

The results of this experiment are shown in the top row of Table 4. The baseline column shown in this table represents the class distribution, or the performance of a classifier that guesses the most frequent class. The set of features that is most predictive of this class is the warehouse features,

Task	Feature Set	Baseline	Accuracy	Precision	Recall	F_1
Sybil Relationship	Warehouse + Ego + Topic	0.819	0.951	0.953	0.951	0.952
Parent Child	Warehouse + Ego + Embedding	0.560	0.611	0.606	0.611	0.608

Table 4. Predictive performance results of all tasks. The baseline is a majority-class baseline.

the ego network, and the set of topic features. The majority class is "not in a sybil relationship," a label assigned to 81.9% of the accounts. Our classifier is able to greatly surpass that, achieving 95.1% accuracy at this task. This indicates that the user's behavior (warehouse features), topics of discussion (topics), and social interactions (ego network features) are all important in understanding whether an account participates in a sybil relationship. Next, we dive into an analysis of which of these features are most important.

We have discovered that our set of features lend to good predictive performance, enabling us to achieve over 95% accuracy on the task of identifying which accounts are in a sybil relationship. However, which of our feature sets are the most predictive at this task? To answer this question, we performed a feature ablation test. The results of this are shown in Table 5.

For the sybil/non-sybil task, the most predictive features are the weighted degree ratio, degree, and weighted degree of the user. This corresponds with what we see in Figure 2, where those participating in sybil relationships have a different degree distribution from those that do not. Beyond this, we see that the two-hop size is predictive, which may be a further reflection of the degree of the nodes. The clustering coefficient was also predictive, with sybil users generally having a lower clustering coefficient. This shows that the users with which sybils play games do not tend to play games with each other. Embeddedness reflects this pattern at the two-hop network. These results strongly suggest that sybil accounts are more likely to play games with random users than non-sybil accounts. Combined with the insights from the topic analysis, this may mean that sybils' higher degree results from weak ties rather than strong ties. In other words, in *World of Tanks*, sybil accounts tend to play with a wider pool of people, but may not deeply engage with them.

4.3 Parent or Child

Within a sybil relationship, there is one parent and one or more child accounts. Knowing which is which can help us to understand the motivations of the account and to better understand the dynamics of a game. For instance, if a child account sees vastly more play time than the parent, this is a strong indication that the user has moved on from their parent account. This is useful for game developers to understand the experience of their players.

Using the same sets of features as above, we investigated the efficacy of these features at predicting whether the account is a parent or a child. Here, we conditioned the data on those accounts that are in a sybil relationship and only sought to predict their role. This became a binary prediction task. Using the same hyperparameters for the random forest classifier as before, we explored which feature combinations are the most predictive.

The results of this experiment are shown in the bottom row of Table 4. The performance is notably worse than the first task. This is understandable; as the accounts are controlled by the same person, they will exhibit similar behavior patterns. However, the fact that we were able to achieve above-baseline performance in this task shows that there are some behavioral differences between parent and child accounts, which are explored through topic analysis. Due to the lack of predictive performance, we omitted a feature comparison.

Feature	Importance Score
Coplay Count	0.044
Embeddedness	0.045
Clustering Coefficient	0.048
Two Hop Size	0.086
Weighted Degree	0.100
Degree	0.148
Degree Ratio	0.185

Table 5. Top features as determined by ablation test. The importance score is the accuracy with the feature - the accuracy without the feature. A higher score indicates a more predictive, or important feature.

5 CONCLUSIONS AND FUTURE DIRECTIONS

In this work, we studied the behavior of sybil accounts on *World of Tanks*, a massively multiplayer online game. Sybils are a component of most systems, including *World of Tanks*. As such, the nature and purpose of these alternative accounts are matters of increasing concern in both academic and industry spaces. We conducted a large-scale study of one server, containing hundreds of thousands of user accounts. We examined these users from the perspective of what they say in aggregate through topic modeling algorithms. Then, we used machine learning techniques to detect sybils and to see which behavioral patterns are most indicative of sybil behavior.

Our analysis of the topics reveals a rich, multilingual environment. We find that sybils and non-sybils tend to socialize differently. For example, non-sybils engage in more conversational topics than those in sybil accounts. Within sybil accounts, the child accounts are more likely to focus on topics that concern gameplay. Despite these differences, there are many similarities between parent sybil accounts and those not engaged in sybil accounts.

Next, we attempt to classify users based upon their sybil status. This comprises two separate but related classification tasks. The first is to predict if a user is a sybil of any type or not. In this task, we were able to achieve over 95% accuracy using features from the user's gameplay, topics, and social network. The second task we undertook was to predict who are the parent and child accounts within a sybil relationship. Because these accounts are controlled by the same people, this task proved to be much more difficult, only surpassing a majority-class baseline by about 5%.

This work provides insight into the behavioral differences between different types of users in the sybil space. Factoring in sybils would mean that researchers or system managers will be able to accurately consider the *peoples'* experiences and motivations, not individual *accounts'*. We can more closely approach user needs and experiences—whether it be identity management or exploration [3, 11], privacy concerns [16, 18, 21], influencing public opinions [38], or gaining visibility or dominance [4, 38]—as well as the system features that facilitate their negotiations or manipulations. For instance, if users tend to create an account for the wider public and a separate one for a smaller circle [16], introducing features that assist privacy management could help. If users create sybils to dominate less skilled players, games can re-adjust the matching algorithm so players with a competitive motivation or an intent to explore new features can do so without negatively affecting their parent account's gameplay records, including win rate. A natural extension of this project is to understand why users create sybils to begin with. Our analysis suggests a handful of trends that future researchers could investigate.

For example, one explanation for the creation of sybil accounts is that the user is tired of the identity they have created and wishes to adopt another account. Another could be their competitive motivation [15] to dominate the game against lower-skilled players, colloquially known as "seal clubbing" in the *World of Tanks* context [4]. However, our analyses show that these players still maintain and use their original (parent) accounts. For example, about two-thirds of child accounts play fewer games than their parent accounts, suggesting that most players with sybil accounts spend more time playing with their parent account.

Furthermore, we find that the child accounts are less skilled than the parent accounts. Slightly less than half (47.8%) of child accounts have a higher win rate than their parents, and only 14.8% of child accounts use a higher tank tier than their parent account. Thus, it is possible that sybil users in *World of Tanks* are interested in exploring roles that differ from their existing experiences [11].

Our analysis suggests that users may be inclined to create sybils to try out different roles towards different groups of users, particularly in regards to gameplay rather than socialization for conversational or clan-related purposes. In this game, we cannot conclude that users create sybils for malign purposes. This diverges from the common view of sybils as solely bad actors (i.e. "sybil attacks" [38]). *World of Tanks* is a game, so the primary context for creating sybil accounts is likely to surround gameplay. However, insults and hate speech appeared across sybils, non-sybils, parents, and children alike, so there is a possibility that some users intend to try on a more toxic role. Game practitioners could consider the person-based network to afford the system a more comprehensive view of problematic user behaviors.

Identifying sybils within games would allow developers treat multiple accounts as one person; holistically viewing sybils as a unified player experience affords developers greater insights into catering services, customer support, and other incentives towards players as a whole rather than individual accounts. After assessing the context of their game and its player needs, the developers can implement features such as customizable privacy or visibility settings, low-risk trials of various types of roles and character types, support for management of multiple accounts through the primary account, and resetting or renaming the original account to various degrees.

Another application of the sybil approach would be to help the "true" new users to acclimate to the game, especially in games incorporating player versus player elements. Identifying problematic user behaviors related to sybils, like seal clubbing, would help game developers improve user experience and retention. Such uses of sybils may subvert the game experience for the benefit of an individual at the expense of others, so regulating those accounts would have implications for both developers and users. Developers may be able to introduce new features to address issues caused by sybils. For example, if a "seal clubbing" account is detected, developers may build in some functionality that satisfies player desires for empowerment, but not at the expense of others. For the victims of such behaviors, removing or redirecting the offenders will likely improve their experience, especially when they are at the start of their learning curves. This benefits the players, as well as the companies, who are more likely to retain them and their potential spending.

This study also presents a means for detecting sybils from in-game user behaviors. Applications of these insights into other contexts, such as social media, can further the understanding of the behavior and psychology of sybils. For instance, it can contribute to detecting, understanding, and preventing the dissemination of mis- and disinformation, potentially both within and across "echo chambers." Delving into user motivations of creating and using sybils may help further our understanding of privacy management and group identification and help create tools that support these motivations, such as features for group differentiation. However, we would like to note that such implementations should be in respectful consideration of users' understanding of the platform and in favor of protecting personal privacy. Moreover, we strongly recommend that qualitative and contextual components be considered when applied to regulatory application.

As this project was in collaboration with Wargaming.net, we have a unique level of access to sybils and their characteristics. We acknowledge that our insights draw on data that may not be readily accessible outside of such collaborations. However, our access was based on limited and anonymized data sets to respect user privacy and maintain our relationship with Wargaming.net. Institutions with more comprehensive access to similar kinds of data should be cognizant of the privacy risks to users. Nevertheless, as our study has shown, further research into sybil accounts in games and other media has rich potential.

To better understand the motivations behind creating sybil accounts, future work could involve more detailed analysis of player warehouse data in relation to sybils and non-sybils. Due to the variety of features available and the ongoing collection of warehouse data, there is a breadth of behavioral features that are unexplored. These traces of player data could provide insights into the reasoning and use patterns behind sybil accounts. We acknowledge that these measures have a limited capacity to infer player intentions and recommend that methods such as surveys, interviews, and ethnographic studies be applied to provide further support for these inferences. Nevertheless, this study demonstrates the underexplored but rich potential for sybils to improve understandings of user practices and experiences for researchers and industry practitioners alike.

ACKNOWLEDGMENTS

The authors thank Wargaming for providing access to their anonymized data, and Eugene Kislyi, Sergey Laptenok and Jeremy Ballenger for their feedback.
This material is based upon work supported by the National Science Foundation under Grant No. 2035064. It was also supported in part by the USC Viterbi School of Engineering through the Center for Knowledge-Powered Interdisciplinary Data Science (CKIDS).

REFERENCES

[1] Erfan Aghasian, Saurabh Garg, Longxiang Gao, Shui Yu, and James Montgomery. 2017. Scoring users' privacy disclosure across multiple online social networks. *IEEE access* 5 (2017), 13118–13130.
[2] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *JMLR* 3 (2003), 993–1022.
[3] Liam Bullingham and Ana C Vasconcelos. 2013. 'The presentation of self in the online world': Goffman and the study of online identities. *Journal of information science* 39, 1 (2013), 101–112.
[4] Sam Chandler. 2016. Why Seal Clubbing Is Good. https://bit.ly/seal-clubbing-is-good.
[5] Jonathan Chang, Jordan Boyd-Graber, Chong Wang, Sean Gerrish, and David M Blei. 2009. Reading tea leaves: How humans interpret topic models. In *Neural information processing systems*, Vol. 22. Citeseer, 288–296.
[6] Mia Consalvo. 2009. *Cheating: Gaining advantage in videogames.* Mit Press.
[7] John R Douceur. 2002. The sybil attack. In *Int'l workshop on peer-to-peer systems.* Springer, 251–260.
[8] David Easley, Jon Kleinberg, et al. 2010. *Networks, crowds, and markets.* Vol. 8. Cambridge university press Cambridge.
[9] Frank A Farris. 2010. The Gini index and measures of inequality. *The American Mathematical Monthly* 117, 10 (2010), 851–864.
[10] Peng Gao, Binghui Wang, Neil Zhenqiang Gong, Sanjeev R Kulkarni, Kurt Thomas, and Prateek Mittal. 2018. Sybilfuse: Combining local attributes with global structure to perform robust sybil detection. In *IEEE CNS.* IEEE, 1–9.
[11] Erving Goffman et al. 1978. *The presentation of self in everyday life.* Harmondsworth London.
[12] Neil Zhenqiang Gong, Mario Frank, and Prateek Mittal. 2014. Sybilbelief: A semi-supervised learning approach for structure-based sybil detection. *IEEE Transactions on Information Forensics and Security* 9, 6 (2014), 976–987.
[13] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *KDD.* 855–864.
[14] Jessie Cameron Herz. 1997. *Joystick nation: how videogames gobbled our money, won our hearts, and rewired our minds.* Abacus.
[15] Adam S Kahn, Cuihua Shen, Li Lu, Rabindra A Ratan, Sean Coary, Jinghui Hou, Jingbo Meng, Joseph Osborn, and Dmitri Williams. 2015. The Trojan Player Typology: A cross-genre, cross-cultural, behaviorally validated scale of video game play motivations. *Computers in Human Behavior* 49 (2015), 354–361.
[16] Jin Kang and Lewen Wei. 2020. Let me be at my funniest: Instagram users' motivations for using Finsta (aka, fake Instagram). *The Social Science Journal* 57, 1 (2020), 58–71.
[17] Yubo Kou. 2020. Toxic Behaviors in Team-Based Competitive Gaming: The Case of League of Legends. In *Proceedings of the Annual Symp. on Computer-Human Interaction in Play.* 81–92.
[18] Alex Leavitt. 2015. "This is a Throwaway Account" Temporary Technical Identities and Perceptions of Anonymity in a Massive Online Community. In *CSCW.* 317–327.

[19] Alex Leavitt, Joshua Clark, and Dennis Wixon. 2016. Uses of multiple characters in online games and their implications for social network methods. In *CSCW*. 648–663.

[20] Renming Liu and Arjun Krishnan. 2020. PecanPy: a fast, efficient, and parallelized Python implementation of node2vec. *BioRxiv* (2020).

[21] Alice E Marwick and danah boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society* 13, 1 (2011), 114–133.

[22] Marçal Mora-Cantallops and Miguel-Ángel Sicilia. 2018. Player-centric networks in League of Legends. *Social Networks* 55 (2018), 149–159.

[23] Bonnie Nardi and Justin Harris. 2006. Strangers and friends: Collaborative play in World of Warcraft. In *CSCW*. 149–158.

[24] Constantinos Patsakis, Athanasios Zigomitros, Achilleas Papageorgiou, and Edgar Galván-López. 2014. Distributing privacy policies over multimedia content across multiple online social networks. *Computer Networks* 75 (2014), 531–543.

[25] Johanna Pirker, André Rattinger, Anders Drachen, and Rafet Sifa. 2018. Analyzing player networks in Destiny. *Entertainment Computing* 25 (2018), 71–83.

[26] Nathaniel Poor and Marko M Skoric. 2014. Death of a guild, birth of a network: Online community ties within and beyond code. *Games and Culture* 9, 3 (2014), 182–202.

[27] Kate Raynes-Goldie. 2010. Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook. *First Monday* (2010).

[28] Radim Rehurek and Petr Sojka. 2011. Gensim–python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic* 3, 2 (2011).

[29] Johnny Saldana. 2015. *The Coding Manual for Qualitative Researchers*. SAGE.

[30] Tyler J VanderWeele and Maya B Mathur. 2019. Some desirable properties of the Bonferroni correction: is the Bonferroni correction really so bad? *American journal of epidemiology* 188, 3 (2019), 617–618.

[31] Binghui Wang, Jinyuan Jia, and Neil Zhenqiang Gong. 2019. Graph-based Security and Privacy Analytics via Collective Classification with Joint Weight Learning and Propagation. In *ISOC Network and Distributed System Security Symp.*

[32] Binghui Wang, Le Zhang, and Neil Zhenqiang Gong. 2017. SybilSCAR: Sybil detection in online social networks via local rule based propagation. In *IEEE INFOCOM 2017-IEEE Conf. on Computer Communications*. IEEE, 1–9.

[33] Gang Wang, Tristan Konolige, Christo Wilson, Xiao Wang, Haitao Zheng, and Ben Y Zhao. 2013. You are how you click: Clickstream analysis for sybil detection. In *22nd {USENIX} Security Symp. ({USENIX} Security 13)*. 241–256.

[34] Dmitri Williams, Nicolas Ducheneaut, Li Xiong, Yuanyuan Zhang, Nick Yee, and Eric Nickell. 2006. From tree house to barracks: The social life of guilds in World of Warcraft. *Games and culture* 1, 4 (2006), 338–361.

[35] Dmitri Williams, Tracy LM Kennedy, and Robert J Moore. 2011. Behind the avatar: The patterns, practices, and functions of role playing in MMOs. *Games and Culture* 6, 2 (2011), 171–200.

[36] Zenghua Xia, Chang Liu, Neil Zhenqiang Gong, Qi Li, Yong Cui, and Dawn Song. 2019. Characterizing and detecting malicious accounts in privacy-centric mobile social networks: A case study. In *KDD*. 2012–2022.

[37] Zhi Yang, Christo Wilson, Xiao Wang, Tingting Gao, Ben Y Zhao, and Yafei Dai. 2014. Uncovering social network sybils in the wild. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 8, 1 (2014), 1–29.

[38] Zhi Yang, Yusi Zhang, and Yafei Dai. 2018. Defending against Social Network Sybils with Interaction Graph Embedding. In *2018 IEEE CNS*. IEEE, 1–9.

[39] Nick Yee. 2006. Motivations for play in online games. *CyberPsychology & behavior* 9, 6 (2006), 772–775.

[40] Nick Yee. 2014. *The Proteus paradox: How online games and virtual worlds change us-and how they don't*. Yale University Press.

[41] Haifeng Yu, Phillip B Gibbons, Michael Kaminsky, and Feng Xiao. 2008. Sybillimit: A near-optimal social network defense against sybil attacks. In *2008 IEEE Symp. on Security and Privacy*. IEEE, 3–17.

[42] Dong Yuan, Yuanli Miao, Neil Zhenqiang Gong, Zheng Yang, Qi Li, Dawn Song, Qian Wang, and Xiao Liang. 2019. Detecting fake accounts in online social networks at the time of registrations. In *SIGSAC*. 1423–1438.

[43] Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. 2014. *Social media mining: an introduction*. Cambridge University Press.

Received February 2021 ; revised June 2021 ; accepted July 2021