

Are People Really Social on Porn 2.0?

Gareth Tyson

Queen Mary
University of London
gareth.tyson@qmul.ac.uk

Yehia Elkhatib

Lancaster University
y.elkhatib@lancs.ac.uk

Nishanth Sastry

King's College London
nishanth.sastry@kcl.ac.uk

Steve Uhlig

Queen Mary
University of London
steve.uhlig@qmul.ac.uk

Abstract

Social Web 2.0 features have become a vital component in a variety of multimedia systems, e.g., YouTube and Last.fm. Interestingly, adult video websites are also starting to adopt these Web 2.0 principles, giving rise to the term “Porn 2.0”. This paper examines a large Porn 2.0 social network, through data covering 563k users. We explore a number of unusual behavioural aspects that set this apart from more traditional multimedia social networks. We particularly focus on the role of gender and sexuality, to understand how these different groups behave. A number of key differences are discovered relating to social demographics, modalities of interaction and content consumption habits, shedding light on this understudied area of online activity.

1 Introduction

Social Web 2.0 features have become a core element of many content-sharing websites. Prominent examples include music sharing and curation in Last.fm (Aiello et al. 2012); image sharing and curation in Flickr (Mislove et al. 2008) and Pinterest (Zhong et al. 2013); and user generated video sharing in YouTube (Cheng, Dale, and Liu 2008; Wattenhofer et al. 2010). Whereas these topics have been investigated in a range of online multimedia repositories, there is a domain which, as of yet, has received next to no attention. Studies have shown that a significant portion of Internet traffic is *adult video content* (Schulze and Mochalski 2009), often far exceeding that of traditional streaming services (Fiadino, Bär, and Casas 2013). In line with their burgeoning popularity, many adult video portals have begun to introduce social Web 2.0 principles; a propensity that has given rise to the term “Porn 2.0”. For example, many now support user-generated content (UGC), as well as video commenting and rating. These services have exploded in popularity, yet have seen little research. We therefore know next to nothing about the nature of Porn 2.0, nor the behavioural traits of its users.

A prominent emerging feature of Porn 2.0 is that of *online social networks* (OSNs), which are increasingly being integrated into well known adult video portals such as RedTube and xHamster. Whereas this is an interesting concept in itself, it also offers the potential to shed light on a topic

which has long been a largely untouched taboo. It raises several particularly interesting questions for the field of user behaviour analysis: Do users actually behave socially on such websites? What is the impact of gender and sexuality? (How) is social content shared? Who do users interact with? What kinds of practices and behaviours make users popular?

To answer these questions, we study the PornHub adult video repository, one of the largest Porn 2.0 websites, consistently featuring amongst the Alexa Top 100. We crawled the PornHub website, recording information from their public social network. PornHub incorporates some of the most sophisticated social networking features seen in the domain, including the ability to form friendships, upload and share content, send messages and post on each others’ “walls”. Our crawl collected 563k user profiles, with publicly available information about demographics, interactions, social groups and multimedia provision/consumption. We also discovered that over 3k accounts have been verified by PornHub through a manual vetting process, giving us a non-trivial sample of ground truth information on demographics.

In this paper, we offer the first study of an adult multimedia OSN, using our unique dataset to explore various aspects of the domain, highlighting unusual observations when compared to more traditional OSNs. A common theme of exploration throughout the remainder of the paper is that of *gender and sexuality*. Sexual differences between the genders have long been known in this domain. These include differing sexual attitudes (Geer and Robertson 2005), differing content preferences (Hald 2006), and differing rates of consumption (Kvalem et al. 2014). As of yet, however, these have focussed on traditional (passive) pornography consumption, rather than active social environments. This study offers the opportunity to give larger-scale insights, building on prior studies compiled by methods such as interviews and questionnaires (Short et al. 2012).

We begin by analysing the demographic make-up of the users. Confirming smaller studies (Kvalem et al. 2014), we observe a prominence of male heterosexual users (§ 4). We show that this group struggles to accumulate friends, whilst women succeed in gaining large social groups (§ 5). This property leads us to ask the question of how such demographics impact the wider activity of the system and why users might make friends in the OSN. We show that demographics have a significant impact on users’ social capital,

with women outperforming men in all modalities of interaction (§ 6). In reaction to these occurrences, we find preliminary evidence suggesting that some male users might even be posing as homosexual women (§ 4).

Inspired by these findings, we correlate various profile features with popularity, finding that male profiles can best excel by being proactive, measured, for example, by the number of subscriptions they make (§ 5.2). This indicates that such users must make their own success and actively find others, employing strategies to stand out from the crowd (Griskevicius et al. 2006). In contrast, female profiles find it much easier to gain large friendship groups without such effort; this is therefore representative of more traditional mating environments, where males often have had to compete more (Kruglanski and Stroebe 2012). That said, in all cases, we observe a positive correlation between the amount of content uploaded by a profile and the number of relationships created (§ 6.3). Thus, we confirm that media exchange plays a key role in the OSN. Our findings have strong implications for both sociologists wishing to understand the role of gender and sexuality in this emerging area, as well as computer scientists, who could use our findings to inform the design of systems.

2 Background and Related Work

Pornography is anecdotally the most searched for content on the web. Whereas work has gone into understanding the sociological aspects of online sexual activities (Carroll et al. 2008; Daneback et al. 2012), little is known about the engines that enable its distribution, especially the expanding “Porn 2.0” phenomenon. Recently, numerous YouTube-like websites have emerged (e.g., PornHub), allowing users to freely upload, view, rate and comment on adult videos. The scale of these services is huge; recent work (Tyson et al. 2013) recorded over 60 billion views on one such Porn 2.0 website, whilst another study found that some adult video sites can even exceed the traffic footprint of traditional video sites (Fiadino, Bär, and Casas 2013).

Although much effort has been invested in exploring user behaviour in traditional multimedia social networks, the adult domain remains largely untouched. That said, various communities have focussed on a number of specific sub-components of Porn 2.0, such as pornographic practices, communities and subcultures (Attwood 2010); interest recommendations (Schuhmacher, Zirn, and Völker 2013); online adult dating (Jacobs 2009); user commenting (Trestian, Xiao, and Kuzmanovic 2013); and illegal content dissemination (Hurley et al. 2013). Interestingly, our data is tightly coupled to the multimedia corpus itself, offering further insights into the behaviour of users in the wider video service.

Related to our own work are the variety of studies into more traditional multimedia social networks. These include music sharing and curation (Aiello et al. 2012); image sharing and curation (Mislove et al. 2008); and user-generated video sharing (Cheng, Dale, and Liu 2008; Wattenhofer et al. 2010). These studies have highlighted a number of interesting findings in terms of how users operate in these social environments. For example, it was found that 25% of all YouTube views are driven by social networks (Wattenhofer

et al. 2010). There have also been a number of studies into more general purpose OSNs, e.g., Google+ (Schiöberg et al. 2012) and Facebook (Backstrom et al. 2011; Dow, Adamic, and Friggeri 2013; Farahbakhsh et al. 2013). A recurrent theme of our work is the impact of gender and sexuality on behaviour. Differences relating to this have been identified in several social networks, including Facebook and Twitter (Taylor 2012). Women in Pinterest, for example, gain more repins than men (Otoni et al. 2013), whilst they also tend to have slightly more friends in MySpace (Thelwall 2008) and pay greater attention to their closest friends in Facebook (Backstrom et al. 2011). Sexuality, however, is a topic that has received far less attention in mainstream social networks. There is a wealth of sociological literature on gender roles in pornography usage and sexual practices. These have highlighted different content preferences (Hald 2006), reactions to visual stimuli (Carvalho et al. 2013) and frequency of use (Kvalem et al. 2014). However, prior studies have focussed on small-scale exploration (e.g., via interviews). Further, we are currently unaware how these observations relate to online social environment, rather than the traditional model of passive pornography consumption. We are particularly interested to see how gender and sexuality impact activities within the social network other than content consumption.

3 PornHub Dataset

PornHub is a video sharing website that allows users to upload and view adult content, much like YouTube. It also has an expanding base of features, with a built-in OSN. This allows users to create profiles and form social connections (friendships and subscriptions) with each other. They can exchange private messages, upload/share videos and pictures, and post on each others’ “walls”.

We scraped the associated OSN profiles attached to PornHub during October and November 2013. To initiate the crawl, we used the search facilities on PornHub to retrieve all users from the 60 largest cities within and the 48 largest cities outside of the USA (based on population), giving us a seed set of 102k users. Starting from this seed set, we performed a breadth-first crawl (traversing friendship links) aiming to discover the largest connected component of the social graph. In total, we collected 563k profiles. There are 26.7 million directed friendships and 3.9 million subscriptions. In each profile, where available, we recorded user demographics, their “wall” of comments, number of profile views, number of videos that they have watched, videos they have favourited, all videos and picture galleries they have uploaded, their social relationships, and their public video viewing history.

We witness four profile types. The most populated is what we term “normal”, which any user can set up without any prerequisites (e.g., verification, payment). These accounts constitute 99.3% of all users. There is, however, a special form of user account, called “verified”. These accounts have been manually validated by PornHub; this involves taking a picture of oneself with some message that links the individual to the account (e.g., next to their username). Although not foolproof, it does lend a far higher level of user trust

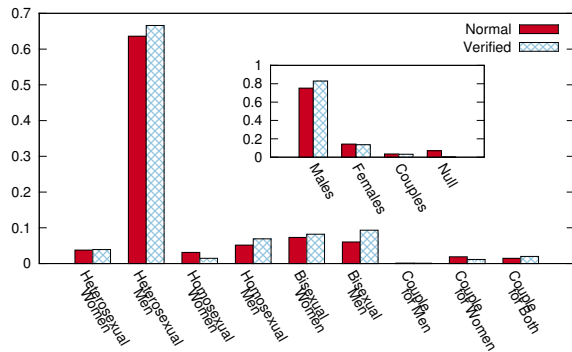


Figure 1: Distribution of reported gender and sexuality

in profile details (e.g., gender) than typical. Only 0.005% (3169) of the accounts crawled have this status. Despite being a small sample size, verified accounts offer a more reliable view of dedicated users, providing a form of (partial) ground truth against which statistics about normal accounts can be compared. There are also 209 company profiles and 8 PornHub “staff” profiles. Unless otherwise stated, we filter out staff and producer accounts to get a more accurate portrayal of typical users.

4 Truth and Lies in Profile Demographics

A key question that underlies any behavioural study is who is participating. Social studies have highlighted a prominence of pornography usage by male users: 38.8% of men between the ages of 18–30 view it in excess of twice a week, compared to just 6.9% of women (Hald 2006). An open question is therefore whether or not this male dominance maps into adult social environments too. Whilst self reported data is, of course, susceptible to manipulation by the account owner, it offers an accurate portrayal of how people wish to present themselves (as well as how others perceive them).

Figure 1 provides an overview of the distribution of reported genders and sexual preferences. The overwhelming majority of users classify themselves as male (75%), with a notable portion as female (14%). This outstrips reports on websites such as Adult Friend Finder, where, for example, only 8% of accounts were female in Hong Kong (Jacobs 2009). It is also very different to OSNs such as Twitter and Facebook, where females dominate (Taylor 2012). That said, it is broadly similar to other studies that, for example, identified 21% of cybersex addicts to be female (Cooper, Delmonico, and Burg 2000). We also discover a small fraction (3%) of accounts registered by pairs of individuals sharing a single account (represented as a gender termed “couple”). This is a novelty, as shared profiles are a rarity in social media. The remainder (8%) do not specify a gender.

We also inspect reported ages. Figure 2 compares the age distribution in PornHub against that of averages taken from 24 OSNs (Pingdon 2012). First, it can be seen that users of PornHub are far younger than in other OSNs; everybody is registered as above 18, although the majority of users

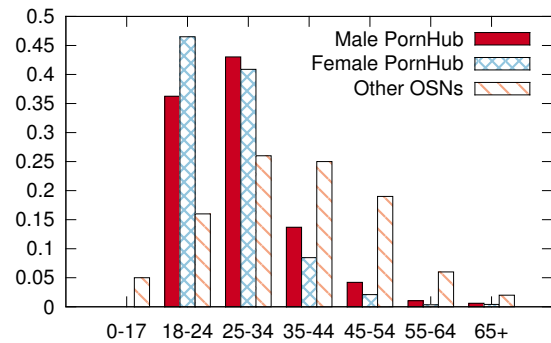


Figure 2: Distribution of reported age in PornHub compared against other OSNs (Pingdon 2012)

report their age as under 40 (90%). We can also compare the reported ages of men and women against each other, shown in Figure 2. The patterns are quite different, with 47% of women reporting themselves as under 25 compared to just 36% of men. This could be explained by the fact that younger women have shown greater affinity to explicit visual content (compared to older), whilst men continue to maintain an interest into older years (Ferree 2003). An alternative explanation might be that these ages are simply misrepresented; studies have shown that women are generally interested in same-aged to somewhat older men, whereas men exhibit preferences towards younger women (Antfolk et al. 2015). It is therefore possible that the observed ages could simply be a manifestation of these preferences, with users reporting what they anticipate to be attractive. Lying in online sexual interactions is commonplace: 48% of Internet users report misrepresenting personal facts, at least in terms their age (Cooper, Delmonico, and Burg 2000). Unfortunately, there is no ground truth to compare these theories against, but we note that the distribution of ages across both normal and verified accounts is very similar.

We next inspect the respective sexual preferences of each user, also shown in Figure 1. Intuitively, one would expect declaration of sexuality to be more accurate than gender and age classifications, as users would likely want to collect friends who match their true preferences. Although the majority of users are heterosexual men, there is a range of other user groups. Bisexuality amongst both men and women is quite common in the dataset. This could be an example of users trying to garner as much attention as possible by not limiting their scope for friends. We can again turn to the verified dataset to compare. Broadly speaking, classifications are quite similar. However, we observe that the number of heterosexual men is higher in the verified set than in the normal one. When looking at the verified group, this shift is manifested in an increased number of homosexual women. We conjecture that these could be, in fact, men posing as women. Unlike Adult Friend Finder, we posit that very few people intend to actually meet each other (c.f. § 6.1), making it conducive for providing false data regarding gender. Sociological work has explored a variety of reasons why peo-

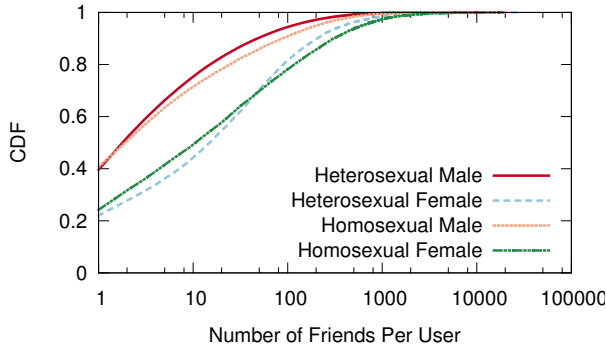


Figure 3: CDF of number of friends per user

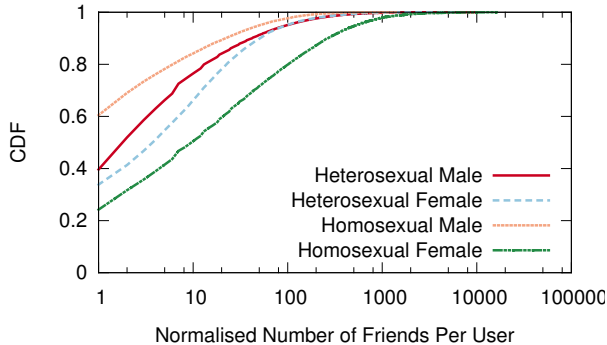


Figure 4: CDF of number of friends per user normalised by availability of gender

ple “gender swap” online (e.g., fun and exploration (Hussain and Griffiths 2008)), and it is worth noting that the saturation of male heterosexuals would make befriending women difficult (c.f. § 5.1). By claiming to be women interested in other women, it is possible that men get more attention. Although not as prominent in more “real-world” OSNs, similar inclinations have been observed in other domains such as online gaming (Lou et al. 2013). We are not able to confirm this conjecture with our dataset, but we consider it an interesting future line of work.

5 How to Win Friends and Influence People

Online sexual gender dynamics can be significantly different to the real world (Ferree 2003). As such, past gender work in other social networks does not translate to PornHub well (Ottoni et al. 2013; Backstrom et al. 2011). We thus explore how gender and sexuality impact friendship groups.

5.1 Relationship Creation

Social environments naturally lead to an uneven distribution of friendship sizes (Backstrom et al. 2011). This is driven by many factors, including accessibility, perceptions and commonality of interest. We ask how *popularity*, as measured by the number of friends, arises in PornHub. 70% of users have fewer than 10 friends, and 90% have fewer than 10

Gender	Friends	Subscribers	Subscriptions
Hetero. M	27	2	7
Homo. M	41	11	11
Bi. M	55	9	11
Hetero. F	98	20	4
Homo. F	136	25	8
Bi. F	179	32	6

Table 1: Breakdown of average number of relationships based on reported gender and sexual preferences

subscribers. Remarkably, 10% of all users have no friends whatsoever.

Here, we investigate how this diversity in relationship sizes is impacted by reported gender and sexual preference. For example, as 69% of users claim that they are interested in women, it would be logical that women can accumulate most attention, particularly as they only constitute 14% of the population (according to declared gender). Table 1 presents the average friendship group size for each gender. It can be seen that high diversity exists, with distinct trends. Most notably, it can be seen that female profiles consistently outperform male ones. This is quite different to more general-purpose OSNs that show broadly even friendship group sizes across genders (Volkovich et al. 2014). To explore this further, Figure 3 presents the distribution of friends across different genders and sexualities. It can be seen here that female profiles of both sexualities far outperform male profiles, and that there is a strong skew towards an upper percentile of “heavy hitters”. The oversupply of male profiles therefore generates unique dynamics, that do not exist in most typical real-world environments (Kruglanski and Stroebe 2012) or other OSNs (Thelwall 2008; Ottoni et al. 2013). To address the differences in gender populations, Figure 4 introduces a normalised friendship size (ϕ) given as,

$$\phi_i = |F_i| \cdot \left(1 - \frac{|U_g|}{|U|}\right) \quad (1)$$

where U is the set of all users, and U_g is the set of users that belong to a gender g (male or female); F_i is the friend set of user $i \in U$, who is interested in people from U_g . Thus, the metric scales up friendship sizes for those interested in women. It can be seen that, through this normalisation, the difference between male and female heterosexual friendship sizes narrows dramatically. The exception is that of homosexual female profiles, which still achieve larger friendship groups, despite the paucity of women in the system. These specific users become popular regardless of competition. This is quite unique, as most multimedia and general purpose social networks have little concept of gender preference (e.g., Facebook, Pinterest, Spotify).

A possible explanation is that homosexual women might be very flexible in who they befriend. To explore this, for each profile, we compute the percentage of friends who match their sexual preference. For example, a heterosexual male who only has female friends would be allocated a score of 100%. Due to the ability of users to befriend others who are outside of their sexual preference, we treat friendships

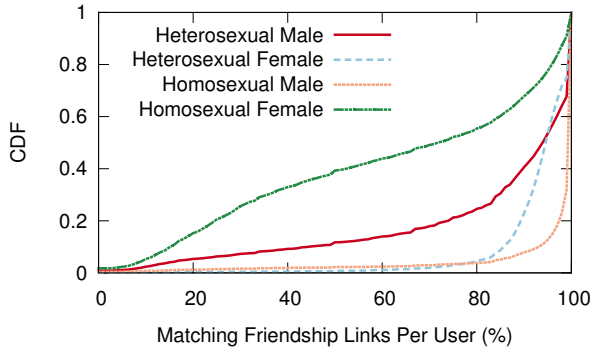


Figure 5: CDF of per user percentage of sexual preference matching friendship links

as directed. Figure 5 presents the CDFs across each sexuality group. We include only users with over 5 friends. A significant divergence can be seen, with users who are interested in women possessing far fewer preference matching links. This could be for one of two reasons: (i) the paucity of women in the system may be *forcing* these social users to befriend men as well; or (ii) these users are simply less discerning about who they befriend. Perhaps more noteworthy is the fact that homosexual female profiles are by far the least discerning, possessing significant numbers of male friends. It is unclear why exactly this occurs, however, we conjecture that the bulk of male profiles are far more likely to accept (or request) friendship from homosexual women than other men (thereby allowing women to easily form friendships with the large population of men). This explains why homosexual women can consistently achieve large social groups.

5.2 Popularity Factors

The above disparity across genders and sexualities raises the obvious question of why some users are more popular than others. To explore this, we have separated profiles into sexuality groups and performed correlation analysis to detect characteristics related to friendship size. Due to space constraints we focus on four sexuality groups, shown in Figure 6, for the 10 features most correlated (by Spearman rank) with friendship group size: profile views received, users subscribed to, wall comments sent, videos watched, wall comments received, private and public albums, videos favoured, self wall comments, and video views received from other users.

Some intuitive findings can be seen; for instance, the number of profile views and the number of comments sent/received correlate quite well with friendship sizes. There are, however, some surprising discoveries. First, it can be seen that the relative importance of these factors varies heavily across the different user groups. For example, the number of comments received by male profiles is far less correlated with popularity than when compared to women: Whereas popular women gain many comments, popular men do not (c.f. § 6.2). In other words, even men who gain large friendship groups, struggle to gain attention

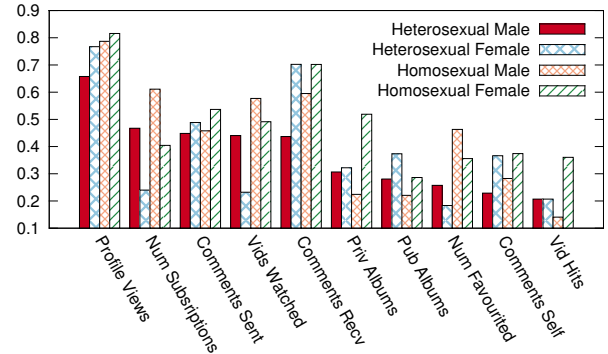


Figure 6: Spearman's rank correlation coefficient for user friendship sizes

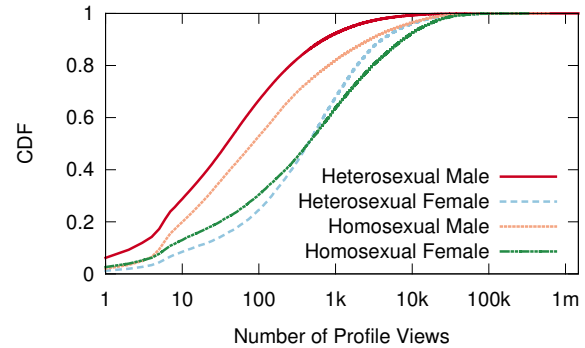


Figure 7: CDF of number of profiles views per account

from those friends. Also, the number of video views received is only correlated with popularity for homosexual women; i.e., other groups might upload popular content but it does not necessarily equate to large social groups.

Another key observation is that user activity impacts popularity. Both the number of subscriptions made and the number of videos watched by a user correlates (weakly) with popularity. This suggests that highly active users gain more friends in return. However, this does not hold true for heterosexual female profiles, which do well regardless of their activity. These principles are also evident in other areas of social theory, where men may have to be more active in mating (Griskevicius et al. 2006). The remainder of the paper explores these features to understand their relationship with gender, sexuality and the impact on popularity.

6 Exploring User Activity

The previous section has identified a number of key profile features that correlate with user friendship sizes. Many of these features relate to specific modalities of user activity, e.g., viewing profiles or uploading content. Next, we explore these activities to understand how they vary across reported sexualities.

6.1 Profile Browsing

Joinson *et al.* state that “virtual people watching”, where users passively observe others, plays an important role in OSNs, even if there is no interest in meeting offline (Joinson 2008). Unlike OSNs such as Facebook, we anticipate few users to know each other offline. To investigate this, we compute the distance between friends using their reported city/country locations. On average, friends are 4796 km away from each other, with only 22% of friends within 1000 km (compared to 50% for Google+ (Schiöberg *et al.* 2012)). Manual inspection confirms that few openly elicit real-world meet-ups. This is quite similar to other “interest-based” social networks such as Twitter, which has an average link distance of 5117 km (Scellato *et al.* 2010).

It therefore seems intuitive that many users may people watch without an interest in offline meeting. To explore this, we turn our attention to the number of profile views each user receives. Overall, we record approximately half a billion views, far higher than the 26.7 million directed friendships. This confirms the propensity of people to browse profiles. Again, we find a strong skew with the top 10% of accounts collecting 85% of the views. We therefore ask whether gender and sexuality influences this skew.

Figure 7 presents the distribution of profile views across different user types. It can be seen that female users gain far more attention. Another interesting observation is that sexuality has a significant impact on the attention a profile receives. Homosexual male profiles do far better than heterosexual ones. Further, it can be seen that homosexual female profiles have a greater tendency towards a few heavy hitters: The top 1% of homosexual female profiles constitute 20.7% of their views, compared to 26.4% for heterosexual females. This indicates that these small user groups effectively manage to stand out from the large mass of heterosexual male users. We argue that the ability to differentiate yourself from others is a key requisite for forming large friendship groups in a social networks that lacks existing real-world context.

6.2 Posting Comments

Similar to other OSNs, users in PornHub can receive and post comments on each others’ walls. This allows us to see how truly interactive users are. Some users (3%) make their comment wall private. 24% of users with private walls are female, possessing a mean of 152 friends compared to just 44 for all others. They also upload more content, suggesting that privacy is more important for the socially active (note females in other OSNs have exhibited higher privacy concerns (Thelwall 2008)). As we cannot see their comments, we filter private walls.

In total, we observe 1.3 million postings. The majority of walls (76%) lack any comments. However, there is a notable set of very active users, with several thousand profiles possessing hundreds of comments. For the most part, the comments are text-based (96.7%), but there is a non-negligible set of multimedia posts: 2.3% are videos, whilst 1% are pictures. The next question is how does sexuality and gender impact this commenting behaviour?

Unsurprisingly, women get the most attention, receiving an average of 9.6 comments compared to just 0.8 for men.

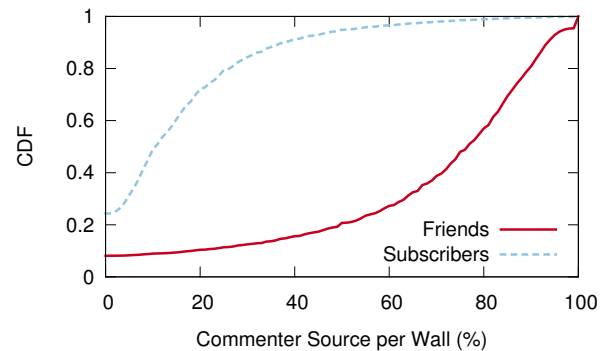


Figure 8: CDF of sources (friends or subscribers) of posts per wall (walls with over 10 comments)

When removing users with empty walls, these averages increase to 15.4 and 4.4 respectively (median 4 and 1). Much of this attention, however, is not directly responded to, with only 0.76 self-posts for every 10 received. Overall, 66% of all comments are on female profiles (despite the small number of women in the overall population). To further explore this, we study the nature of interactions between the different genders. Female profiles tend to receive notably more multimedia posts than males. 3.7% of comments sent from men to women are multimedia, whilst only 1.9% are sent from women to men. Inspection shows that this is often user-generated content.

A further feature of PornHub is the ability to post on strangers’ walls. Strangers are users who are neither friends nor subscribers. Figure 8 presents the distribution of “commenters” per wall separated into friends and subscribers. 67% of comments have been left by friends, and only 16% by subscribers: A surprising number of strangers choose to post on people’s walls. The lack of (utilised) privacy controls seem to encourage users to initiate such interactions. A notable number of females also post on strangers’ walls (18% are from females). This indicates that, proportionally, women are just as proactive in posting comments on strangers’ walls as men are. Initially, we anticipated many of these might be spam accounts. However, manual inspection revealed that the bulk of heavy commenters are simply highly active users. The activities on these social networks therefore are quite unusual compared to other OSNs. Either way, it is clear that there is a considerable amount of social interaction occurring, and that female profiles continue to exhibit dominance.

6.3 Content Uploading

The third type of interaction modality is that of content uploading. Related work has shown that OSNs can have a significant impact on content access, with 45% of all YouTube views being driven by social networks in a video’s first week (Wattenhofer *et al.* 2010). We are therefore curious to see how this translates into an adult environment. Both picture albums and videos can be uploaded as either public or private (requiring friendship before the content becomes

	Public V	Private V	Public P	Private P
Hetero. M	0.147	0.195	0.28	0.306
Hetero. F	0.091	0.114	0.373	0.322
Homo. M	0.265	0.330	0.220	0.224
Homo. F	0.115	0.173	0.286	0.518

Table 2: Spearman’s rank correlation coefficients for sexuality group’s popularity based on content provision (V is video, P is picture albums)

accessible in the latter option).

We have recorded 692k content uploads from social profiles. 80% of users provide no objects (neither public nor private), indicating most users are mere consumers. Male profiles, generally, view more videos: 44% of males view over 100 objects, whereas this is only 31% of female profiles. The tendency to share is also impacted heavily by reported gender. 33% of females upload content, whilst it is only 19% for male profiles. Thus, it seems that supply-and-demand is somewhat being addressed, whereby the demographics in greatest demand, proportionally, offer the greatest supply.

Our first goal is to understand how content uploading impacts popularity. Figure 9 presents the median friendship sizes correlated with various types of upload. An upward trend can be observed, with heavy uploaders achieving larger social groups. Interestingly, it does not take a large number of objects to attract attention, with increases in friendship sizes levelling out after the first few uploads. The exception here is private picture albums, which have a dramatic impact on popularity. We argue such provisions are another example of users wishing to stand out from the crowd (Griskevicius et al. 2006). This suggests that availability of private content may be a key motivating factor for befriending people. On the one hand, private content increases popularity because friendship is required to view the content, but we also posit that the amount of private content is a proxy for how engaged a user is in the OSN (inherently, private content depends on friendship). As such, this would drive greater friendship sizes as those who offer content probably also pursue social interactions.

To explore this further, we compute the correlation between friendship sizes and content provision for each of the sexuality groups. Table 2 shows that the correlations vary between the different groups. Although the correlations are not strong in all cases, private content provision is the more prominent indicator of popularity. Most notably, private photo albums are well correlated with homosexual female popularity. Clearly, these users have private content that is highly desirable. The significance of user contributions is also clearly shown in absolute values. Male profiles that offer content get 7 times as many friends as male profiles that do not, whilst female profiles offering content get only 5.7 times as many as female profiles that do not. This is because female profiles can achieve high friendship counts regardless of content provision, whereas male profiles need to “fight” for their social status.

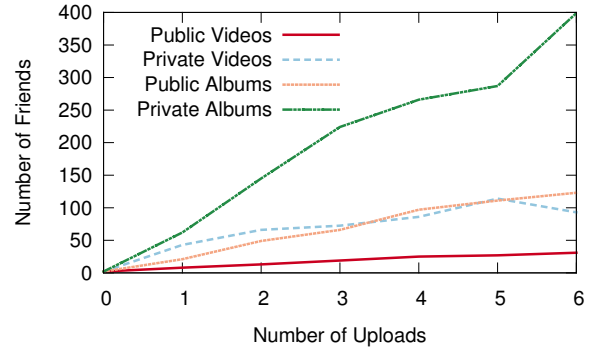


Figure 9: Median number of friends vs. user uploads

6.4 Content Category Preferences

Last, we shed light on the content consumption preferences for the different groups. In choices of pornographic content, the genders have long been understood to differ. For instance, women tend to be more positive towards less explicit content, whilst men often prefer the opposite (Hald 2006). However, related findings have also suggested that women react in a more sexual way to more explicit imagery in the same way as men (Carvalho et al. 2013). To explore this, we inspect public viewing histories listed in users’ profiles. Each video can be tagged with one or more category labels, taken from a controlled set of 61 defined by the website. Although, as previously stated, there is the potential for users to misreport their gender/sexuality, we argue that our analysis provides valuable insight into general trends.

We begin by compiling the set of videos watched at least once by members of each of the four sexuality groups. We then divide this set into subsets of views per category. We manually removed the “Pornstar” category because this applied to most videos. For each category subset, we compute the average number of views per user. This presents a locally normalised popularity of video tags within that sexuality group. Figure 10 presents the rankings for each group in a stacked bar (note that each stack is not cumulative).

It can be seen that a ranking is followed, with some categories far more popular than others. There is, however, significant diversity across the different user groups. Perhaps unsurprisingly, the second most popular category for profiles reported as homosexual females is “Lesbian”, a category that does not feature in the top 10 of any other group. Male-focussed categories similarly are more prominent for profiles interested in men; for instance, “Big Dick” is the second most highly ranked category for homosexual males.

There are also many less intuitive findings. In line with (Carvalho et al. 2013), various particularly explicit categories are present in female profile viewing histories. Surprisingly, the “Female Friendly” category is also disproportionately popular with profiles reported as homosexual males, despite often featuring women. Further, it can be seen that each group consumes content (to some extent) from all categories. This is caused partly by the shared category tags, which introduces noise. However, clear trends show

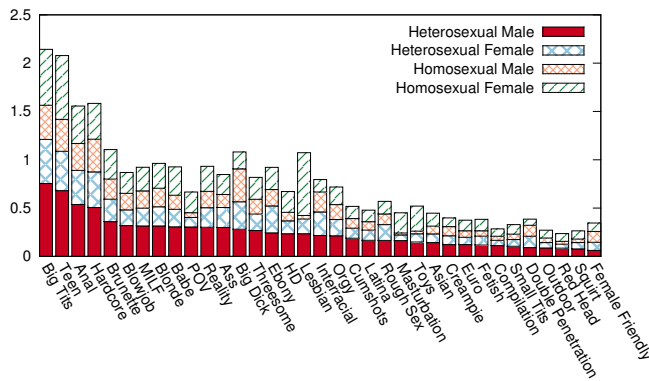


Figure 10: Category rankings based on per user average view counts for each category (ordered by heterosexual male ranking); stacked bars are not cumulative

that many sexual groups extensively experiment with content that intuitively one would imagine a lack of interest in. Most notably, it can be seen that profiles reported as homosexual men consume large amounts of heterosexual content: Their highest ranked category heavily features women.

7 Summary and Implications

This paper has presented the first study of the social features of “Porn 2.0”. Intentionally, we have given a broad overview of the platform. However, a key thread linking our analyses is the importance of gender and sexuality. The scale of these adult OSNs is significant, and their characteristics are quite unique within the wider domain. Unlike past work that has analysed gender in OSNs, we study a system that has an explicit purpose for gender, captured by sexual preference. Although we have focussed on a single adult OSN, it is likely that our findings are generalisable across others. We believe that more “2.0” services (of all varieties) will continue to emerge and, as such, it is important to understand and extract both technological and sociological insight.

First, we explored the demographic make-up of the system. We showed that these adult OSNs are primarily populated with young (<40) males in stark contrast to general purpose OSNs (Pingdon 2012). Although past work (Kvalem et al. 2014; Hald 2006) has highlighted an increased likelihood of males consuming pornography, we believe this to be the first work to shed light on OSN demographics. We show that these males often struggle to compete in the ecosystem, and do significantly worse than female profiles in all modalities of interaction (e.g., comments). This is exacerbated by the tendency for users to select other users who match their gender preferences, thereby limiting the capacity of users interested in women to form large groups. We argue that to be successful such OSNs *must* therefore expand their female population dramatically. Of course, algorithms that can effectively pair users would also be highly attractive. More generally, this selective type of behaviour also shows that many users are actively engaged in online social interaction. We find that males can take steps to improve their social standing, just as in the real-world:

Males who are more active (e.g., subscribing more) gain larger social groups. In contrast, high activity is less correlated with increased popularity for women, who do well regardless. Related observation are made elsewhere; for example, in mating, men have often had to compete more than women (Kruglanski and Stroebe 2012).

We also confirmed that content is a prominent component of this system, augmenting other forms of interaction. Users with content (both male and female) significantly increase their popularity. That said, the proportional benefit is actually greater for male profiles. This is the classic example of a “standing out from the crowd” strategy (Griskevicius et al. 2006). We further inspected user content preferences. Interestingly, there were notable similarities across the content accessed by different reported genders and sexualities. However, clear distinguishing features could be extracted with, for example, homosexual female profiles being attracted to categories not featuring men, such as “Toys”. This creates important implications for any technological services underpinned by this data. Most notably, recommendation engines could exploit this demographic information when suggesting content. Similarly, automated friendship suggestions could be driven by content interest similarities.

This is just the first step towards fully understanding the nature and implications of adult OSNs. Such services could offer a further wealth of sociological insights. Here, we have only explored a subset, focussing on the roles of gender and sexuality. Our future work will be to expand this to cover also bisexual users and couples, who have very different dynamics. Another fascinating future area is exploring more deeply the open questions regarding gender swapping. We have found preliminary evidence that this occurs, but have not investigated in detail. Clearly, it is critical to understand how this property affects our findings. We plan to use machine learning to better model the similarities between some male and female accounts (e.g., based on viewing histories). Finally, so far, we have largely relied on quantitative data; these analyses will also be extended to cover more qualitative issues. Most important is the question of *intent*, exploring what the goals of individual users are.

References

- Aiello, L. M.; Barrat, A.; Schifanella, R.; Cattuto, C.; Markines, B.; and Menczer, F. 2012. Friendship prediction and homophily in social media. *ACM Transactions on the Web* 6(2).
- Antfolk, J.; Salo, B.; Alanko, K.; Bergen, E.; Corander, J.; Sandnabba, N. K.; and Santtila, P. 2015. Women’s and men’s sexual preferences and activities with respect to the partner’s age: evidence for female choice. *Evolution and Human Behavior* 36(1):73 – 79.
- Attwood, F. 2010. *Porn.com: Making sense of online pornography*, volume 48. Peter Lang.
- Backstrom, L.; Bakshy, E.; Kleinberg, J. M.; Lento, T. M.; and Rosenn, I. 2011. Center of attention: How facebook users allocate attention across friends. In *ICWSM*.
- Carroll, J. S.; Padilla-Walker, L. M.; Nelson, L. J.; Olson, C. D.; Barry, C. M.; and Madsen, S. D. 2008. Genera-

- tion XXX pornography acceptance and use among emerging adults. *Journal of adolescent research* 23(1):6–30.
- Carvalho, J.; Gomes, A. Q.; Laja, P.; Oliveira, C.; Vilarinho, S.; Janssen, E.; and Nobre, P. 2013. Gender differences in sexual arousal and affective responses to erotica: The effects of type of film and fantasy instructions. *Archives of sexual behavior* 42(6):1011–1019.
- Cheng, X.; Dale, C.; and Liu, J. 2008. Statistics and social network of youtube videos. In *IWQoS*.
- Cooper, A.; Delmonico, D. L.; and Burg, R. 2000. Cybersex users, abusers, and compulsives: New findings and implications. *Sexual Addiction and Compulsivity* 7(1-2):5–29.
- Daneback, K.; Sevcikova, A.; Månsson, S.-A.; and Ross, M. W. 2012. Outcomes of using the internet for sexual purposes: fulfilment of sexual desires. *Sexual Health* 10(1).
- Dow, P. A.; Adamic, L. A.; and Friggeri, A. 2013. The anatomy of large facebook cascades. In *ICWSM*.
- Farahbakhsh, R.; Han, X.; Cuevas, A.; and Crespi, N. 2013. Analysis of publicly disclosed information in facebook profiles. In *IEEE/ACM ASONAM*.
- Ferree, M. 2003. Women and the web: Cybersex activity and implications. *Sexual and Relationship Therapy* 18(3):385–393.
- Fiadino, P.; Bär, A.; and Casas, P. 2013. HTTPTag: A flexible on-line HTTP classification system for operational 3G networks. In *INFOCOM Poster Session*.
- Geer, J. H., and Robertson, G. G. 2005. Implicit attitudes in sexuality: Gender differences. *Archives of Sexual Behavior* 34(6):671–677.
- Griskevicius, V.; Goldstein, N. J.; Mortensen, C. R.; Cialdini, R. B.; and Kenrick, D. T. 2006. Going along versus going alone: when fundamental motives facilitate strategic (non) conformity. *Journal of personality and social psychology* 91(2):281.
- Hald, G. M. 2006. Gender differences in pornography consumption among young heterosexual danish adults. *Archives of sexual behavior* 35(5):577–585.
- Hurley, R.; Prusty, S.; Soroush, H.; Walls, R. J.; Albrecht, J.; Cecchet, E.; Levine, B. N.; Liberatore, M.; Lynn, B.; and Wolak, J. 2013. Measurement and Analysis of Child Pornography Trafficking on P2P Networks. In *WWW Conference*.
- Hussain, Z., and Griffiths, M. D. 2008. Gender swapping and socializing in cyberspace: An exploratory study. *CyberPsychology & Behavior* 11(1):47–53.
- Jacobs, K. 2009. Is there life on Adult FriendFinder? Sex and logic with the happy dictator. In *Digital Arts and Culture Conference*.
- Joinson, A. N. 2008. Looking at, looking up or keeping up with people?: Motives and use of facebook. In *CHI*.
- Kruglanski, A. W., and Stroebe, W. 2012. *Handbook of the history of social psychology*. Psychology Press.
- Kvalem, I. L.; Træen, B.; Lewin, B.; and Štulhofer, A. 2014. Self-perceived effects of internet pornography use, genital appearance satisfaction, and sexual self-esteem among young scandinavian adults. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*.
- Lou, J.-K.; Park, K.; Cha, M.; Park, J.; Lei, C.-L.; and Chen, K.-T. 2013. Gender swapping and user behaviors in online social games. In *WWW Conference*.
- Mislove, A.; Koppula, H. S.; Gummadi, K. P.; Druschel, P.; and Bhattacharjee, B. 2008. Growth of the flickr social network. In *Workshop on Online social networks*.
- Otoni, R.; Pesce, J. P.; Las Casas, D. B.; Franciscani Jr, G.; Meira Jr, W.; Kumaraguru, P.; and Almeida, V. 2013. Ladies first: Analyzing gender roles and behaviors in pinterest. In *ICWSM*.
- Pingdon, R. 2012. Report: Social network demographics in 2012. <http://royal.pingdom.com/2012/08/21/report-social-network-demographics-in-2012/>.
- Scellato, S.; Mascolo, C.; Musolesi, M.; and Latora, V. 2010. Distance matters: geo-social metrics for online social networks. *Workshop on Online Social networks*.
- Schiöberg, D.; Schneider, F.; Schiöberg, H.; Schmid, S.; Uhlig, S.; and Feldmann, A. 2012. Tracing the birth of an OSN: Social graph and profile analysis in google+. In *Web Science Conference*.
- Schuhmacher, M.; Zirn, C.; and Völker, J. 2013. Exploring youporn categories, tags, and nicknames for pleasant recommendations. In *Workshop on Search and Exploration of X-Rated Information*.
- Schulze, H., and Mochalski, K. 2009. Internet study 2007–2009.
- Short, M. B.; Black, L.; Smith, A. H.; Wetterneck, C. T.; and Wells, D. E. 2012. A review of internet pornography use research: Methodology and content from the past 10 years. *Cyberpsychology, Behavior, and Social Networking* 15(1):13–23.
- Taylor, C. 2012. Women win facebook, twitter, zynga; men get linkedin, reddit. <http://mashable.com/2012/07/04/men-women-social-media/>.
- Thelwall, M. 2008. Social networks, gender, and friending: An analysis of myspace member profiles. *Journal of the American Society for Information Science and Technology* 59(8):1321–1330.
- Trestian, I.; Xiao, C.; and Kuzmanovic, A. 2013. A glance at an overlooked part of the world wide web. In *WWW Conference*.
- Tyson, G.; Elkhatib, Y.; Sastry, N.; and Uhlig, S. 2013. Demystifying porn 2.0: A look into a major adult video streaming website. In *ACM IMC*.
- Volkovich, Y.; Laniado, D.; Kappler, K. E.; and Kaltenbrunner, A. 2014. Gender patterns in a large online social network. In *Social Informatics*. Springer. 139–150.
- Wattenhofer, M.; Interian, Y.; Vaver, J.; and Broxton, T. 2010. Catching a viral video. In *Workshop on Social Interaction Analysis and Service Providers*.
- Zhong, C.; Shah, S.; Sundaravadivelan, K.; and Sastry, N. 2013. Sharing the loves: Understanding the how and why of online content curation. In *Proc. Conference on Weblogs and Social Media*.