

Fake it till you make it: Fishing for Catfishes

Walid Magdy*, Yehia Elkhatib[§], Gareth Tyson[‡], Sagar Joglekar[†], and Nishanth Sastry[†]

*School of Informatics, The University of Edinburgh, UK

[§]School of Computing and Communications, Lancaster University, UK

[‡]School of Electronic Engineering and Computer Science, The Queen Mary University of London, UK

[†]Department of Informatics, King's College London, UK

wmagdy@inf.ed.ac.uk, y.elkhatib@lancaster.ac.uk, gareth.tyson@eecs.qmul.ac.uk,
sagar.joglekar@kcl.ac.uk, nishanth.sastry@kcl.ac.uk

Abstract—Many adult content websites incorporate social networking features. Although these are popular, they raise significant challenges, including the potential for users to “catfish”, *i.e.*, to create fake profiles to deceive other users. This paper takes an initial step towards automated catfish detection. We explore the characteristics of the different age and gender groups, identifying a number of distinctions. Through this, we train models based on user profiles and comments, via the ground truth of specially verified profiles. Applying our models for age and gender estimation of unverified profiles, we identify 38% of profiles who are likely lying about their age, and 25% who are likely lying about their gender. We find that women have a greater propensity to catfish than men. Further, whereas women catfish select from a wide age range, men consistently lie about being younger. Our work has notable implications on operators of such online social networks, as well as users who may worry about interacting with catfishes.

I. INTRODUCTION

Adult content has been a long standing innovator in technology. Whereas, originally, online adult content was primarily distributed via pay-per-view sites and peer-to-peer networks, we have recently witnessed a radical shift termed “*Porn 2.0*” with the integration of diverse features into popular portals, *e.g.*, videos, images, webcams and chat functionality. For example, many now support user-generated content (UGC), as well as video commenting and rating. These services have exploded in popularity, yet research has not kept pace with their advancement.

One of the most powerful features introduced is that of online social networking (OSNs). Adult OSNs, much like Facebook, allow users to create and interconnect social profiles. Anecdotally, these profiles have led to a plethora of fake accounts created by users in an attempt to deceive others regarding their true identity — so called “catfishing” [1]. In this paper, we restrict ourselves to two forms of identity deception: users lying about either their gender or age (or both). Although we do not take a moral position about catfishing and perhaps in many cases it does not harm users, there is also a notable subset of situations in which such fake accounts could directly damage others (or themselves). For example, catfishes may pretend to be younger than their true age in an attempt to attract partners; similarly, younger users may pretend to be

older to avoid legal age filters. Alternatively, men may pretend to be (homosexual) women in an attempt to garner more female friends [2]. Indeed, our exploration of one such Porn 2.0 OSN userbase reveals that certain demographics do excel in terms of popularity metrics (§IV), potentially motivating deceit. Thus, any mechanism to detect such deception would have significant value to both the OSN operators and users concerned about being deceived.

In this paper we ask *to what extent does catfishing occur*, and also *is it possible to automatically detect them?* To answer this, we study the Pornhub adult OSN, one of the largest and most popular Porn 2.0 portals in the world. We have crawled the Pornhub website recording all data from 99,727 OSN profiles (§III). 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”. This makes it ideal as a case-study for analysis.

A particularly novel feature of Pornhub is the ability to create “verified” profiles that are manually checked by Pornhub employees. This gives a unique ground-truth of accurate profiles that is useful for our purpose. Using this data, we first characterise the activities of these different genders and age groups to observe notable differences (§IV). With these analyses, we move on to explore several algorithms to detect users lying about their age and/or gender (§V). As an input we utilise profile features and the comments left by users. Using well recognised state-of-the-art techniques [3], [4], and with additional set of features we propose, we build models to predict true user demographics (gender and age) based on these inputs.

With this subset of catfish profiles, we then begin to explore their characteristics (§VI). Although our focus is *not* to discover the unique reasoning behind catfishing, our results do shed light on some of the benefits. We compute popularity metrics (*e.g.*, number of friends) for catfishes to find that they tend to gain larger friendship and subscription groups. Although this metric may be skewed by the possibility of catfishes being more proactive in befriending others, we also note that catfishes gain more profile views as well (a metric

which is more difficult to game as it depends on others choosing to visit your page). We surmise that the increasing integration of our offline and online personas will result in this becoming even more important (§VII). Of course, although catfish detection may be more explicitly required in adult portals, our methods will also have benefits for various other services, *e.g.*, online dating and mentorship.

The contributions and roadmap of this work is as follows:

- We present the first work to explore and quantify the presence of catfishes on adult OSNs (§IV), making our dataset available for public use¹ (§III).
- We build classifiers for age estimation and gender prediction for adult OSN users, and compare various features for achieving the best performance (§V).
- We profile the demographics and activities of catfish users, highlighting how they differ from truthful users, and benefit from their deception (§VI).

II. RELATED WORK

Pornography is anecdotally the most searched for content on the web. Whereas work has gone into understanding the sociological aspects of sexual activities [5], [6], little is known about the online engines that enable its distribution, especially the expanding “Porn 2.0” phenomenon. Numerous YouTube-like websites have emerged (*e.g.*, PornHub, YouPorn), with in-built social features. Recent work [7] 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 [8].

A novel requirement that stands out in this domain is the need for robust demographics verification. This is particularly driven by the potential incentives for people to deceive others about their true gender and age. There have been a number of recent studies looking at the automated verification of demographics [9], [10], [11]. For instance, [3], [12] utilised textual features including text content, part-of-speech (POS) tags, and discourse styles to train a model using linear regression for age prediction. They tested their approach on three different collections including online forums. They managed to achieve an age prediction with correlation 0.535 and mean-absolute-error (MAE) in age of 6.5 years. In an extension to their work [4], they applied their methods on Twitter data, and achieved a MAE of 4.1, which was shown to be better than human estimation of age. In another study, [13] utilised profile and network information for age and gender prediction. The results show that network information leads to significant improvement in prediction, and they explain it as a reason of the principle of homophily of Twitter users. Unlike the above studies, We do not strive to make a contribution to the algorithmic detection of users lying about demographics. Instead, we exploit these existing state-of-the-art methods and apply them to this new and important domain. This is a topic of increasing importance, with countries such as the UK proposing mandatory age-checks for such websites [14].

To the best of our knowledge, there has been no prior work focusing on demographic verification within the adult domain, although various works have focused on related components of Porn 2.0, such as pornographic practices, communities and subcultures [15]; interest recommendations [16]; dating services [17], [18]; user commenting [19]; content popularity [20] and illegal content dissemination [21]. The closest to our work is a recent study of the PornHub social network [2], although it did not touch upon age verification.

A number of studies have also looked at online gender swapping more generally [22]. For example, [23], [24] found that gender swapping in online gaming is commonplace. The reasons for such activities have also been explored via surveys [25]. Reasons given include: (*i*) curiosity and the desire to experiment; (*ii*) the perception that the opposite gender is treated better; and (*iii*) the belief that playing the opposite gender will allow new forms of behaviour, or gain advantages. This complements our own work, although we do not focus on finding the reasons for users to catfish.

III. METHODOLOGY & DATA

A. Data Collection

PornHub is a video sharing website that allows users to upload and view adult content, much like YouTube. It also has a built-in OSN, allowing 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 have scraped the associated OSN profiles attached to PornHub. To initiate the crawl, we used the search facilities on PornHub to retrieve all users from the 60 largest cities within the US, as well as 48 capital cities selected around the world, *e.g.*, Beijing, Berlin, London, New York, Tokyo, *etc.* Starting from this seed set, we performed a breadth-first crawl (traversing friendship links). If not already available, we also collected the account information for any users who had left a comment on another user’s wall. In total, we collected 99,727 profiles. In each profile, where available, we recorded user age, gender, sexual preference, their “wall” of comments, number of profile views, number of videos that they have watched, and their social relationships. Figure 1 presents a sample account from our data collection showing the information gathered for each account. User IDs in our collection were anonymised to protect users’ privacy.

We collect two profile types. The most populated is what we term “unverified”, which any user can set up without any pre-requisites (*e.g.*, verification, payment). These accounts constitute 94.5% 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 in profile details (particularly gender and age) than typical. To increase our number of verified profiles, we specifically searched for all verified profiles in the above cities, and scraped their

¹Link to be provided in final version

```

<USER>
  <ID>8423012381</ID>
  <TYPE>0</TYPE> <!-- 1 for verified accounts -->
  <AGE>27</AGE>
  <PROFILE>Girl,Single,Girls,united%states</PROFILE> <!-- gender, status, interested in,
  country -->
  <ACTIVITY>7611,202</ACTIVITY> <!-- #videos watched, #posts -->
  <FRIENDS>1350,594,76,680</FRIENDS> <!-- #total friends, males, females, others -->
  <SUBSCRIBERS>605,246,25,334</SUBSCRIBERS> <!-- same as above -->
  <SUBSCRIBEDTO>106,40,7,59</SUBSCRIBEDTO> <!-- same as above -->
  <NCOMMENTS>11,10,87,54</NCOMMENTS> <!-- #comments, #unique comments, #words, #unique
  words -->
  <COMMENTS>
    <comment count=2>thanx sugar</comment>
    <comment count=1>I can definitely say the same! You are damn sexy!</comment>
    <comment count=1>Awwwww you guys are so sweet to this poor horny girl!!!</comment>
    <comment count=1>Thank you all are real sweet</comment>
    <comment count=1>Mmmmmmm PLEASE!!!!</comment>
    <comment count=1>Well thank you sugar!</comment>
    <comment count=1>Ohhh yeah I think I could</comment>
    <comment count=1>Hehehehehe thanx sugar!</comment>
    <comment count=1>No problem...my pleasure in fact!</comment>
    <comment count=1>Damn you are one fine woman!!!! *kiss kiss*</comment>
    <comment count=1>GRRRR I NEED SOMEONE XXXXXX MY XXXX AND XXX RIGHT NOW!!!</comment>
  </COMMENTS>
</USER>

```

Fig. 1: A sample account in the collected data. Pornographic terms are replaced by X's.

accounts. Overall, 5.5% of the accounts crawled were verified (5,484 accounts in total). Despite being a smaller sample size, verified accounts offer a more reliable view of dedicated users, providing a form of (partial) ground truth against which statistics about unverified accounts can be compared. We later exploit these “groundtruth” verified accounts to train models that can be used to classify the remaining unverified accounts.

B. Limitations & Ethical Considerations

Before continuing, it is important to recognise the limitations of our data. First, we acknowledge that the PornHub userbase is not necessarily representative of the wider population. Hence, our results are specific to the PornHub social network. That said, we posit that building catfish detection methods *must* be specific to individual platforms, and therefore we do not view this as a concern.

Second, we later utilise verified profiles to train age and gender prediction models (for use on unverified profiles). An obvious possibility is that the verified profiles and unverified profiles are sufficiently different to make comparability difficult. Although this is a theoretical possibility, our extensive manual inspections have not suggested this is the case. On more practical grounds, we also note that this is the *only* viable near-groundtruth dataset available in the domain. Hence, although the data likely contains noise, we are confident it is not sufficiently significant to undermine our overall results and trends. Further exploration and validation is a major theme within our future work.

Finally, it is important to note that we are collecting personal information about a large number of users. Hence, we took a number of steps to address ethical concerns. We avoided analysing personally identifiable information, *e.g.*, names. We also ensured that our methodology did not involve any communications with the accounts under study. Critically, PornHub

allows users to set their profiles to either public or private. Our data is exclusively made-up of users who have set their profiles to public (*i.e.*, anyone can view them). We did not make any attempt to access private profiles, which would have involved befriending others.

IV. CHARACTERISING PORNHUB DEMOGRAPHICS

Before exploring the presence of catfishes, we briefly characterise the demographic make-up of PornHub accounts. We include both verified and unverified accounts in this analysis, but we ignore accounts labelled as “couple”, “company”, “transgender”, or “not specified” (these collectively make just 4.4% of all accounts) to focus on catfishing in the context of male and female accounts. Note that we interchangeably use the terms ‘girl’ / ‘guy’ and female / male as these are the terms used on the portal.

Figure 2 presents a histogram revealing the number of accounts of each age. 64% of accounts fall into the 18–30 bracket, confirming a young demographic. Perhaps more noteworthy is the distinction between male and female accounts, with female users marginally younger (average of 28 vs. 31). This is driven by a larger population of older male profiles. This could perhaps be explained by that fact that younger women have shown greater affinity to pornographic content (compared to older), whilst men continue to maintain an interest into older years [26].

As well as being able to stipulate their own gender, profiles can also list which other gender(s) they are interested in. Figure 3 presents the number of accounts that stipulate an interest in each gender (male, female, both). We find that 68% of accounts are interested in women (bottom middle plot) of which 90% are heterosexual males, yet only 4.5% of accounts (top right plot) are females interested in males who could satisfy this need. In fact, female accounts are *more* likely to




Fig. 2: Distribution of male and female profiles across age ranges




Fig. 3: Distribution of male and female users interested in men, women and both across age ranges

stipulate an interest in both, rather than men alone. In other words, PornHub suffers from a significant supply-and-demand mismatch: Female profiles are in high demand, yet only a small proportion of profiles can satisfy this demand.

To explore this further, we look at the number of friends garnered per user for both female and male accounts (Figure 4). Profiles labelled as female consistently obtain larger social groups than male ones: an average of 396 vs. 185. A downward trend can also be observed across the age ranges with younger users gaining more friends, which is again more pronounced for female accounts. For example, the average 21 year old female accounts gets 433 friends compared to just 176 for a 50 year old female.




Fig. 4: Number of friends per profile vs. reported age

One explanation for the above could be that younger users are more proactive in befriending others. To explore this, we inspect the number of subscribers for each account type (these are equivalent to followers on Twitter, *i.e.*, directed relationships). Again, we see a strong bias towards female profiles who, on average, have 309 subscribers vs. just 97 for men. Similarly, the number of friends, on average, for female accounts is 403 vs. 188 for men. Clearly, the subscriber numbers are primarily driven by the behaviour of *other* users (as a user cannot subscribe to themselves). Hence, this adds evidence against the claim that young female profiles may get large social groups simply through proactive befriending other. In fact, the average female account is subscribed to just 106 channels compared to 142 for male accounts.

In summary, the above confirms a significant difference between male and female profiles in terms of social popularity (friendship groups). Our aim is not to speculate on *why* this is the case. However, it is clear that these demographic pressures might create incentives for people to lie about profile details. Due to this, we spend the rest of the paper exploring the propensity for users to lie on this social network.

V. DEMOGRAPHIC PREDICTION

To identify deceptive users (catfishes), it is first necessary to extract and predict profile features that indicate true demographics. To do this, we utilise PornHub’s *verified profiles*; these are profiles that have been independently verified by PornHub. Although an assumption, we are confident that these are significantly more trustworthy than non-verified accounts, since they must pass a process of validation by the website.

We therefore use the verified profile data as a training set for building models for age estimation and gender prediction. Models are then used to estimate the demographics of the unverified accounts and compare them to the claimed ones. It is also worth noting that, theoretically, verified users may exhibit different characteristics (*e.g.*, activity level, sexual preferences, education level) to their unverified equivalents, thereby making their use within the classifier less accurate. To mitigate this, we *only* identify users as catfishes when they exhibit a *significant* disparity to their predicted demographics. Regardless, we emphasise that the lack of a groundtruth means that our results only offer a lower-bound.

A. Feature Extraction

First, a set of features is extracted for training a regression model for age estimation and a classifier for gender prediction. Many features are inspired by the state-of-the-art [3], [4], [13]; however, we also introduce a set of novel features that are specific to our data. The set of feature groups extracted from the data set are as follows:

Content features: The comments collected for each user’s profile are used for generating a set of features. Some of the content features are inspired by [3], while we introduce an additional one to measure text formality based on the comments nature on such website:

- **Comments-content:** a set of features representing the bag-of-words (BOW) by each user, where each term represents a binary feature that is set to 1 if appeared in the user’s comments and zero otherwise.
- **Linguistic Inquiry and Word Count (LIWC):** a set of features representing the percentage of terms used in user’s comments that have a given LIWC category [27]. This includes: emotions, questions, self reference, family reference *etc.* This set of features should show the nature of terms used by a given user, where it might be a good indication of the user’s age and/or gender.
- **Comments count:** This includes, (i) number of comments by the user; (ii) percentage of unique comments, as we noticed many users repeats their exact comments on different posts; and (iii) vocabulary variety, which is the number of unique terms used by the user divided by the total length of comments.
- **Comments-formality:** One behaviour that young generations are commonly characterised by is the usage of slang, including shortcuts and abbreviations (*e.g.*, “gr8”, “thx”). We introduced a feature to measure the percentage of informal language in users’ comments. To detect informal text, we used the Xerox part of speech tagger (POS)² to label POS tags of terms in comments. Terms with an undefined tag were considered informal language. We found this approach to be accurate in many cases. However, we noticed that some users use non-English terms. Thus, we applied the same POS tagger for three languages: English, French, and German. Those terms that obtained undefined tags with all three languages were considered informal. Finally, the feature value is the number of informal terms divided by the comments length.

Network features: Another set of features were extracted to represent non-textual information. Most of these features are dependent on the network characteristics of Pornhub, which contains more information on individuals than other general social websites. These features are:

- **Profile:** including a user’s country, status (single / in a relationship), and interested in (men / women / both).
- **Activity:** numbers of videos watched and posted (in log scale).
- **Network:** Numbers of friends, subscribers, and subscribed to (in log scale). In addition, percentage of males and females in each of these lists.

Using the above features, we perform gender and age prediction to identify discrepancies, where users may be exhibiting anomalous behaviour for their stipulated demographics.

B. Gender Prediction

For gender prediction, we train a support vector machine (SVM) binary classifier [28] using the above features of the verified accounts. The created model allows us to then classify the remaining unverified accounts.

To measure the performance of the built classifier, we used 10-fold cross validation on the verified set. We only focused on the verified accounts with at least 10 comments in their profile to avoid training the classifier with samples of sparse content features. The number of verified accounts with 10+ comments are 1,231 out of which only 1,119 accounts had their gender identified. Numbers of males and females are 820 and 299 respectively. Precision, recall, and F-measure were calculated for each gender separately; then macro-F-measure and accuracy were calculated for the overall performance.

Different combinations of the feature groups were tested to find the most effective features. Table I reports the results for gender prediction on the verified accounts. As shown, network features are more effective than content features in predicting user’s gender. This is driven by the differences seen in the popularity of the two genders (§IV). However, when combining content and network features, significantly better accuracy (92%) is achieved. We note that the classifier is more effective in predicting males than females, which is expected since male training samples are almost three times than of female samples. Nevertheless, this result is acceptable for the prediction of unverified accounts for our analysis.

C. Age Estimation

The same set of features has been used for age estimation. We applied two regression techniques for age estimation; namely, SVM regression³ and deep neural networks (DNNs). We have examined different models for SVM, including linear and polynomial regressions. As for the neural network approach, we designed a multi-layered perceptron (MLP) trained to do regression over all the possible age values. The neural network had 3 hidden layers of rectified linear unit (ReLU) activated neurons each. This means the activation function for each neuron was of the form $f(x) = \max(0, x)$

Our experimentation showed that SVM linear regression achieved the most effective results when compared to the other non-linear SVM regression models and DNNs. Actually, DNNs showed the poorest performance, which could be explained by over-fitting. Unlike SVMs that are resilient to over-fitting because of their regularization techniques, DNNs are highly susceptible to over-fitting, especially in situations where feature dimensionality exceeds training dataset size. This is because of exponentially higher number of free parameters. We therefore utilise SVM.

Similar to what we performed in gender prediction, 10-fold cross-validation was applied on the verified accounts with 10+ comments to evaluate the effectiveness of the features and regression techniques. All the 1,231 verified accounts with 10+ comments were used, where the input for training the regression models was the extracted features set, and the label is the declared age by the verified account. We normalized ages over 60 to “60+” in order to reduce the sparsity for these ages. For measuring the performance, we followed the same method applied in literature [3], [12], [4], where Pearson’s correlation

²<https://open.xerox.com/Services/fst-nlp-tools/>

³<http://svmlight.joachims.org/>

Features	Female			Male			Macro F1	Accuracy
	R	P	F1	R	P	F1		
Content	0.542	0.885	0.672	0.974	0.854	0.910	0.791	0.859
Network	0.702	0.847	0.768	0.954	0.898	0.925	0.846	0.887
All	0.769	0.920	0.838	0.976	0.921	0.947	0.893	0.920

TABLE I: Gender prediction performance

and mean absolute error (MAE) between estimated age and actual age are measured.

Features	MAE	Correlation
Content	5.581	0.509
Network	6.119	0.234
All	5.783	0.440

TABLE II: Age estimation performance (SVM)

The best result we achieved with DNNs was a correlation of 0.28 and MAE of 6.8 years. This is much lower than the best results achieved with SVMs. Table II shows the obtained MAE and correlation using different features set when applying SVM linear regression model for age estimation. Unlike gender detection, network features are not effective for age estimation. Actually, it degrades the performance of content features when added to it. Content features from users' comments are the most effective as shown in the table. The model can estimate the age of an account holder from his/her comments with an error of ± 5 years on average. The correlation between estimated ages and actual ones is 0.51. This result compares well with the state-of-the-art methods on other data sets [3]⁴

VI. RESULTS AND DISCUSSION

Next, we utilise our classifiers to explore the unverified users who are classified as catfishes. We consider a catfish to be either: (i) a user who is classified as having a gender different to the one reported; or (ii) a user with a predicted age that is more than 5.581 years⁵ different than their claimed age. Of course, this definition would also identify some verified accounts as catfishes in the case of classification errors. Consequently, we emphasise that our analysis highlights users who are likely to be catfishes — it is not a groundtruth. Hence, our classifiers can be used for raising flags to service provider for applying additional vetting for users with suspicious behaviour.

For consistency, only unverified accounts with more than 10 comments were considered in our analysis, since our models were trained in the same way. Thus, the number of unverified accounts with 10+ comments included in our upcoming analysis is 11,182.

A. Who catfishes?

We start by inspecting the number of accounts that are classified as lying about their gender. Figure 5 shows the

⁴The best result for age estimation achieved in [3] for medical blog users was a correlation of 0.535 and 6.537 MAE

⁵Our model achieves a MAE of 5.581. Thus, we cannot consider catfishes who only marginally change their age.




Fig. 5: The age breakdown for users who lie about their gender (left) and those who don't (right).




Fig. 6: The probability density function of difference between estimated and reported user age, broken down for female and male accounts.

number of catfishes of either gender and in each age range (compared to the number of users classified as honest). Note that in Figure 5, the gender refers to the *reported* gender listed on the profile (not the one that was predicted by our classifier). A key fact emerges from Figure 5 — surprisingly, we find that a large proportion of *female* users pretend to be men. Just 35.1% of male accounts are classified as catfishes compared to 60.7% for women. Although this sounds high, it is roughly equivalent to prior questionnaire-based results (48% [1]). It is also worth noting that a similar observation has been made

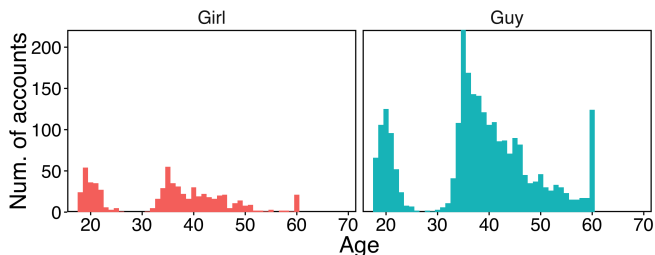


Fig. 7: Reported age of accounts lying about their age, broken down by gender.

in online gaming [23], [24]. A major reason for this was the perception that men are treated differently; we conjecture a similar motive could be at play here.

We can also inspect accounts that are classified as lying about their age. Figure 6 present the difference between reported and predicted ages. It can be seen that catfish accounts reported as male (right) tend to select a marginally younger age, but also experiment with far older age ranges. In contrast, catfish accounts reported as female more consistently select younger ages. The average age of a male catfish is 38.15, but just 34.87 for female catfish. This can clearly be seen in Figure 7, which presents a histogram of the reported ages for the catfish accounts: Male catfish profiles are more likely to be older than female catfish profiles. This tendency is conducive with common social theory; studies have shown that women are generally interested in same-aged to somewhat older men, whereas men exhibit preferences towards younger women [29]. Hence, it seems likely that men pretending to be women would select young ages, whilst women pretending to be men would select from a wider range of ages.

B. What are the benefits of catfishing?

An obvious question is what drives users to catfish. Although there are many potential reasons, we have previously hypothesised that frustrated (unpopular) users may catfish in an attempt to gain more attention (§IV). This can be measured using the profile view counter listed on each profile, as well as the number of friends a profile garners. To inspect this, Figure 8 plots the number of views *vs.* the number of friends. Profiles are separated into catfish (false) and honest (true) users. It can be seen that *honest* (true) female accounts gain by far the most profile views and friend relationships. As the number of views increase, so does the number of friends. On average, these female users have 32,150 views and 444 friends. Honest male users, on the other hand, gain just 6,098 views and 195 friends. This is in-line with our earlier analysis. More interesting, however, is comparing these against users classified as catfishes. Male catfishes, *i.e.*, male users with a female account (red line in the right plot), do substantially better and, in fact, mirror the popularity of women (16590 views and 314 friends on average). In contrast, the average view count for truthful male profiles is just 6098, while they gain only 195 friends. Such observations shed light on the motivation of a male catfish, although we are reticent to say this is the reason. The situation is very different for female catfish, *i.e.*, women pretending to be men (red line in the left plot), who gain far fewer views and friends than their honest counterparts (8719 and 151 respectively). This is because, to other users, they will appear to be men. This makes it unlikely that such users would catfish simply for more attention.

C. What do catfishes want?

We can also decompose the profiles into accounts interested in each option (men, women, both). This allows us to see what profiles classified as catfish are typically seeking. Figure 9 plots the number of friends each account accumulated based




Fig. 8: Number of Profile Views for profiles with false and true genders (smoothed with 95% confidence interval). We use the classifier to separate accounts into male and female (rather than using the reported gender).




Fig. 9: The number of friends per account as a proxy gain for lying about gender, broken down by gender and interest.

on the user’s age. It also breaks users down based on what they stipulate they are interested in. It can be seen that very few male profiles register an interest in both genders (8.0%), whilst this is significantly higher for female accounts (60.5%). We conjecture this might be because females choose to catfish so that they can experiment with other women whilst also interacting with men, whilst men choose to catfish primarily so they can overcome the difficulty of befriending women. Hence, male catfishes tend to lie about being homosexual women; these are actually the users who gain the most (in terms of increased friendship sizes). This is especially the case for men aged 30-55 years (red line in top middle plot), who tend to do quite poorly when not lying.

VII. CONCLUSION

In this paper, we have explored the behaviour and characteristics of users in a major Porn 2.0 social network. Our focus has been on understanding the roles of age and gender, exploiting state-of-the-art algorithms to detect users who lie about these features: so called “catfishes”. To achieve our goal, we built classifiers for gender and age estimation. Using this, we predict that 25% and 38% of unverified users on PornHub are likely lying about their gender and age respectively. We found that women actually have a greater propensity to lie than men. The two genders tend to select different ages when

catfishing, though. Whereas men pretending to be women usually select younger ages, women pretending to be men select from a wider range. Our analysis also reveals that male users frequently benefit from such lies, and obtain larger numbers of views and friends. The opposite is the case for women pretending to be men, confirming that their motives go beyond the acquisition of attention.

This work is just the first step towards understanding the nature of online catfishing. Here, we have only explored deception related to age and gender. Our future work will expand this to understand how some users lie about other attributes such as location and interests. We also plan to use machine learning to better model the similarities between some male and female accounts (e.g., based on viewing histories and browsing patterns [30]). A key priority is corroborating our findings via ground truth information from surveys and questionnaires. This is, of course, because it is currently impossible to gain definitive data beyond the verified profiles – our method is, thus, underpinned by the assumption that there is limited catfishing related to gender and age in verified profiles. Although not ideal, we argue this is the only significant dataset to allow such exploration; further, although noise may exist, the Pornhub verification process has proven highly effective overall. We also assume that the true underlying attributes of unverified profiles are broadly similar to that of verified profiles. This, too, needs to be tested using other methods such as surveys. That said, we remain confident that the state-of-the-art techniques employed here are accurate enough to underpin our findings. Finally, we wish to understand *why* users choose to lie, and how individual users strategise this deception. Again, we plan to execute this via surveys that can garner deeper insight into such reasoning.

REFERENCES

- [1] A. Cooper, D. L. Delmonico, and R. Burg, "Cybersex users, abusers, and compulsives: New findings and implications," *Sexual Addiction and Compulsivity*, vol. 7, no. 1-2, pp. 5–29, 2000.
- [2] G. Tyson, Y. Elkhatib, N. Sastry, and S. Uhlig, "Are people really social on porn 2.0?" in *ICWSM*, 2015.
- [3] D. Nguyen, N. A. Smith, and C. P. Rosé, "Author age prediction from text using linear regression," in *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*. Association for Computational Linguistics, 2011, pp. 115–123.
- [4] D.-P. Nguyen, R. Trieschnigg, A. Doğruöz, R. Gravel, M. Theune, T. Meder, and F. de Jong, "Why gender and age prediction from tweets is hard: Lessons from a crowdsourcing experiment." Association for Computational Linguistics, 2014.
- [5] J. S. Carroll, L. M. Padilla-Walker, L. J. Nelson, C. D. Olson, C. M. Barry, and S. D. Madsen, "Generation XXX pornography acceptance and use among emerging adults," *Journal of adolescent research*, vol. 23, no. 1, 2008.
- [6] K. Daneback, A. Sevcikova, S.-A. Månsson, and M. W. Ross, "Outcomes of using the internet for sexual purposes: fulfilment of sexual desires," *Sexual Health*, vol. 10, no. 1, 2012.
- [7] G. Tyson, Y. Elkhatib, N. Sastry, and S. Uhlig, "Measurements and analysis of a major adult video portal," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 12, no. 2, 2016.
- [8] P. Fiadino, A. Bär, and P. Casas, "HTTPTag: A flexible on-line HTTP classification system for operational 3G networks," in *INFOCOM*, 2013.
- [9] T. Buzzell, "Demographic characteristics of persons using pornography in three technological contexts," *Sexuality & Culture*, vol. 9, no. 1, pp. 28–48, 2005.
- [10] J. Marquardt, G. Farnadi, G. Vasudevan, M.-F. Moens, S. Davalos, A. Teredesai, and M. De Cock, "Age and gender identification in social media," in *Proc. of CLEF 2014 Evaluation Labs*, 2014.
- [11] F. Rangel and P. Rosso, "Use of language and author profiling: Identification of gender and age," *Natural Language Processing and Cognitive Science*, vol. 177, 2013.
- [12] D.-P. Nguyen, R. Gravel, R. Trieschnigg, and T. Meder, "'How old do you think I am?' A study of language and age in Twitter," 2013.
- [13] F. Al Zamal, W. Liu, and D. Ruths, "Homophily and latent attribute inference: Inferring latent attributes of Twitter users from neighbors," *ICWSM*, vol. 270, 2012.
- [14] J. Killock, "A database of the uk's porn habits. what could possibly go wrong?" *Open Right Group*, Oct. 2016. [Online]. Available: <https://www.openrightsgroup.org/blog/2016/a-database-of-the-uks-porn-habits-what-could-possibly-go-wrong>
- [15] F. Attwood, *Porn.com: Making sense of online pornography*. Peter Lang, 2010, vol. 48.
- [16] M. Schuhmacher, C. Zirn, and J. Völker, "Exploring YouPorn categories, tags, and nicknames for pleasant recommendations," in *Proc. Workshop on Search and Exploration of X-Rated Information*, 2013.
- [17] K. Jacobs, "Is there life on Adult FriendFinder? Sex and logic with the happy dictator," in *Proceedings of the Digital Arts and Culture Conference*, 2009.
- [18] G. Tyson, V. C. Perta, H. Haddadi, and M. C. Seto, "A first look at user activity on tinder," in *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016.
- [19] I. Trestian, C. Xiao, and A. Kuzmanovic, "A glance at an overlooked part of the world wide web," in *Proc. WWW Conference*, 2013.
- [20] G. Tyson, Y. Elkhatib, N. Sastry, and S. Uhlig, "Demystifying porn 2.0: a look into a major adult video streaming website," in *Proc. IMC*, 2013.
- [21] R. Hurley, S. Prusty, H. Soroush, R. J. Walls, J. Albrecht, E. Cecchet, B. N. Levine, M. Liberatore, B. Lynn, and J. Wolak, "Measurement and Analysis of Child Pornography Trafficking on P2P Networks," in *Proc. WWW Conference*, 2013.
- [22] A. Bruckman, "Gender swapping on the internet," *High noon on the electronic frontier: Conceptual issues in cyberspace*, pp. 317–326, 1996.
- [23] R. M. Martey, J. Stromer-Galley, J. Banks, J. Wu, and M. Consalvo, "The strategic female: gender-switching and player behavior in online games," *Information, Communication & Society*, vol. 17, no. 3, 2014.
- [24] J.-K. Lou, K. Park, M. Cha, J. Park, C.-L. Lei, and K.-T. Chen, "Gender swapping and user behaviors in online social games," in *Proc. WWW Conference*, 2013.
- [25] Z. Hussain and M. D. Griffiths, "Gender swapping and socializing in cyberspace: An exploratory study," *CyberPsychology & Behavior*, vol. 11, no. 1, pp. 47–53, 2008.
- [26] M. Ferree, "Women and the web: Cybersex activity and implications," *Sexual and Relationship Therapy*, vol. 18, no. 3, pp. 385–393, 2003.
- [27] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count (LIWC): A computerized text analysis program," *Mahwah (NJ)*, vol. 7, 2001.
- [28] T. Joachims, *Learning to classify text using support vector machines: Methods, theory and algorithms*. Kluwer Academic Publishers, 2002.
- [29] J. Antfolk, B. Salo, K. Alanko, E. Bergen, J. Corander, N. K. Sandnabba, and P. Santtila, "Women's and men's sexual preferences and activities with respect to the partner's age: evidence for female choice," *Evolution and Human Behavior*, vol. 36, no. 1, pp. 73 – 79, 2015.
- [30] Y. Elkhatib, R. Killick, M. Mu, and N. Race, "Just browsing? Understanding user journeys in online TV," in *International Conference on Multimedia*. ACM, 2014, pp. 965–968.