

Your Installed Apps Reveal Your Gender and More!

Suranga Seneviratne^{a,b}

suranga.seneviratne@nicta.com.au

Prasant Mohapatra^c

prasant@cs.ucdavis.edu

Aruna Seneviratne^{a,b}

aruna.seneviratne@nicta.com.au

Anirban Mahanti^b

anirban.mahanti@nicta.com.au

^aSchool of EET, University of New South Wales, Australia

^bNICTA, Australia

^cDepartment of Computer Science, University of California, Davis

In this paper, we highlight a potential privacy threat in the current smartphone platforms, which allows any third party to collect a snapshot of installed applications without the user's consent. This can be exploited by third parties to infer various user attributes similar to what is done through tracking. We show that using only installed apps, user's gender, a demographic attribute that is frequently used in targeted advertising, can be instantly predicted with an accuracy around 70%, by training a classifier using established supervised learning techniques.

I. Introduction

Smartphone usage and the associated app market ecosystem are expanding rapidly. The current 25% – 30% smartphone penetration in the mobile device market is expected to reach 60% by the year 2019 [18]. The two most popular smartphone operating systems, Android and iOS, which have a combined market share of 91% [8], reached the one million mark in terms of number of apps in their official markets in 2013 [28, 10]. Advertising accounts for 32% of the monetisation from apps [19] and in year 2014 mobile advertising revenue is expected to reach \$31.45 billion, showing 75% increment [17] compared to previous year.

Increased app monetisation through advertising raises multiple privacy and security issues as developers try to collect user data through *over-permission* (i.e., asking for permissions which are not required for the functionality of the app) and share with third parties. This information is used to provide personalised advertisements, which have a higher likelihood of catching users' interest. For example, according to Yan et al. [29], behavioral targeted advertisements can significantly improve advertisement *click-through-rates*.

Information leakage through *over-permission* can be controlled by the users. Users can either control permissions an app has access to or uninstall the apps which ask for *over-permission*. However, in the platforms Android and iOS, the list of installed apps can be collected by any app developer or an embedded

tracking library without the knowledge of user [12, 4]. This allows third parties to learn a range of information about the user instantly, which can be equivalent to inferences made after a period of data collection through tracking.

In this paper, we show how a third party can infer smartphone users' gender by mining a corpus of apps installed by users. First, we provide a basic characterisation of the differences in app installation patterns between male and female users, using a dataset collected from over 200 smartphone users. Then, we convert these differences into features and build a *linear SVM classifier* to predict a smartphone user's gender, given the list of installed applications and show that accuracy around 70% can be achieved. Finally, we compare our results with other tracking-based data sources and show that the snapshot of installed apps performs comparable given the predictions are instant.

Our results highlight a privacy threat in the current smartphone ecosystems, as some users are not comfortable with disclosing their gender online. For example, according to Chaabne et al. [7], 20% of the *Facebook* users did not reveal their gender. Moreover, this type of instant inferences can be exploited by the service providers to quickly verify or assign a confidence value to user-entered information in smartphone apps. A recent study [16], showed that in the popular online game *World of Warcrafts*, 23% of the male users used female avatars, while 7% of the female used male avatars. There can be multiple reasons for users to give fake information about gender, with privacy concerns being one of them. Neverthe-

less, unrestricted access to other information such as installed app lists allows the service providers to identify users who entered fake information, which may be undesirable for some users.

The remainder of this paper is organised as follows. Section II discusses the related work. Section III describes our dataset and the features used for predictions. Section IV presents the prediction methodology and performance of the predictions. Section V compares our results with prior work and Section VI concludes the paper.

II. Related Work

The possibility of demographic inference through various online footprints of the users has been discussed multiple times [24, 15, 11, 14, 23, 21, 20, 5, 31]. One such data source is the content and style of writing. Using a corpus of content obtained from 71,000 blogs at *blogger.com*, Schler et al. [23] showed that *writing style related features* such as parts-of-speech, function words, and hyper-links and *content-based features* such as simple content words and special classes of words taken from the handcrafted LIWC (Linguistic Inquiry and Word Count) [22] categories can be used to predict the writer’s gender and age.

Similarly, Schwartz et al. [24] showed how the language in *Facebook* status update messages can be exploited to predict user demographics, age and gender as well as user’s personality using *differential language analysis*. Otterbacher et al. [21] predicted the gender of the authors of IMDB movie reviews based on stylistic and content features.

Another source of data used for demographic prediction is the web browsing patterns. Hu et al. [14] used data on page clicks from a major website to predict age and gender of the users, through a *Bayesian framework*. Goel et al. [11] used one-year client-side browsing history of 250,000 users to predict user’s age, gender, race, education, and income through *linear Support Vector Machines*.

Bi et al. [5] predicted age, gender, religion, and political views of the users based on search queries using models trained from *Facebook* likes. Ying et al. [31] used *behavioral features* in the likes of application usage, SMS usage, voice call usage, and *environment features* in the likes of number of Bluetooth and WiFi devices detected per day on mobile phones to predict demographics such as gender, age, and relationship status.

In prior work [26], we developed a framework to identify user traits such as languages spoken and re-

lationship status through the presence of individual apps. However, this method does not work for gender as it cannot be directly associated to the presence of individual apps. To the best of our knowledge, this paper presents the first attempt to predict smartphone users’ gender, solely based on installed apps.

III. Dataset & Feature Engineering

We collected a dataset of installed apps and the corresponding user demographic attributes through an Android app [25]. When installed, this app uploads the list of *user installed apps* to a server and generates a random identifier. *User installed apps* are differentiated from *system installed apps* using the *system* flag provided in Android developer API which identifies the apps came with the system image [2].

The app was distributed among volunteers and users recruited via Amazon Mechanical Turk [1]. These users installed the app and answered a brief questionnaire about demographic attributes such as gender and age together with the random identifiers so that the questionnaire output could be correlated with the collected app list.

Out of 218 users who provided an answer to gender, 131 ($\approx 60\%$) users were male and 87 users were female. In total, there were 4167 unique apps and on average each user had approximately 40 *user installed apps* in their smartphones. For each application, we queried official Android app market, Google Play Store and collected the metadata of the apps such as price, category and the app description. In the next subsections, we define and characterise various features that can be used to predict user’s gender using list of installed apps.

III.A. Numeric Features

Figure 1 shows the *cumulative distribution function (CDF)* of the number of *user installed apps* for each gender. No significant differences can be found between the two genders as the two graphs closely follow each other.

Figure 2 shows the total cost of the *user installed apps*. We observe that male users tend to have more paid apps than female users. For instance, approximately 70% of the female users had no paid apps installed while the corresponding value is only 50% for male users. However, when the money spent increases the difference in the gender distribution decreases. For example, the users who spent more than AUD\$11.00, are equally likely to be male or female.




Figure 1: Cumulative distribution of the number of apps installed by the users




Figure 2: Cumulative distribution of the total cost of the apps installed by the users

III.B. Category-based Features

Google Play Store categorises apps into 30 categories. For each app in our dataset, we queried Google Play Store and obtained the assigned app category. Then we tried to identify the categories that show a high difference in popularity between the two genders. We calculated the *discrete entropy* for each app category with respect to gender. Since we have more male users compared to female users, for the entropy analysis and for the succeeding subsections, we *under-sampled* the male users to match the number of female users.

Figure 3 shows the 5 app categories with lowest entropy for gender (i.e. highest difference in popularity between genders) and Figure 4 shows the 5 app categories with highest entropy. As can be seen from Figure 3, categories *Libraries & Demos* and *Sports Games* are more popular among male users while the category *Casual* is more popular among female users. Categories such as *Personalisation*, *Racing* and *Social* are equally popular between both the genders. To con-

vert these observations to features related to individual users, we selected the percentage of apps in each category as the features representing individual users.




Figure 3: High dispersion in gender




Figure 4: Low dispersion in gender

III.C. Item-based Features

Individual apps showing lowest entropies according to gender are shown in Figure 5. Since there were a large number of apps in our dataset, for this analysis we selected only the apps, that have been installed by at least 10% of the users.

Notice that certain apps show a strong tendency towards one gender. For example, apps such as *Google Translate* and *Reddit is fun* are more popular among male users while the apps such as *Pinterest* and *Ibotta-Cash* are more popular among female users. Some of these observations corroborate market reports published on app popularity among the genders. For example, according to [27] 83% of the users of the photo sharing app *Pinterest* are female, and 97% of the fans

Table 2: Performance of the classifiers
 NB*: Naive Bayes, SVM*: Support Vector Machines

Feature type	Classi.	Accuracy	AUC
Numeric	NB	0.493	0.499
	SVM	0.602	0.500
Category	NB	0.618	0.689
	SVM	0.715	0.693
Item	NB	0.635	0.695
	SVM	0.679	0.640
Content (100)	NB	0.663	0.722
	SVM	0.641	0.646
Content (1K)	NB	0.687	0.740
	SVM	0.680	0.686
Category+Item	NB	0.650	0.698
	SVM	0.683	0.667
Category+Content (1K)	NB	0.698	0.740
	SVM	0.681	0.689
Category+Item+Content (1K)	NB	0.698	0.737
	SVM	0.685	0.690

among users in one gender. For example, terms *network* and *rss* are popular among male users, whereas terms *decor* and *pregnanc* are more popular among female users.

IV. Methodology & Results

From the total dataset, we randomly selected 50% of user data for training and 50% for testing. To measure an average classifier performance, we built 5 training and testing set pairs by sampling the original dataset. Then, we represented each user as a feature vector using the features mentioned in Sections III.A to III.D. To address class imbalance, we randomly under sampled the majority male user class to match the size of the female user class in the training set.

We trained naive Bayes and linear kernel SVM (*Support Vector Machine*) classifiers [9] using the machine learning tool Weka [13] using individual categories of features as well as feature categories in combination. Table 2 summarises the performance of the classifiers observed over the 5 samples.

The results show that numeric features alone do not perform better than the random classification, as the AUC (*Area Under the Curve*) values are approximately around 0.5. Category-based features show higher performance under the SVM model achieving 72% accuracy and 69% AUC values.

Item-based features perform better compared to random prediction; however, they do not outperform category-based features. From the content-based features, selection of more terms (1K) performed better compared to the selection of lower number of terms. Combining features did not provide a significant performance improvement. For example, combining other feature categories with content-based features

Table 3: Classifier Performance
Category + Content(1K)

Sample	Gender	Precision	Recall	AUC	Accuracy
1	Male	0.817	0.754	0.784	0.750
	Female	0.667	0.744	0.76	
2	Male	0.727	0.615	0.681	0.630
	Female	0.528	0.651	0.666	
3	Male	0.849	0.692	0.802	0.741
	Female	0.636	0.814	0.756	
4	Male	0.733	0.704	0.776	0.704
	Female	0.596	0.791	0.728	
5	Male	0.704	0.769	0.674	0.667
	Female	0.595	0.512	0.709	

(1K) did not improve the performance significantly; indicating content-based features are the most influential features deciding the classifier performance.

To explore further, we selected the naive Bayes classifier built using *Content (1K)+Category* features as the best performing classifier as it has the highest classifier performance and accuracy and analysed its performance in detail. Table 3, shows the accuracy, precision, recall, and the AUC of the selected classifier over the 5 samples.

According to Table 3, 3 samples showed over 70% accuracy. Across all samples a precision over 70% could be achieved for male user prediction with a recall varying between 62%–77%. Precision for female users varied between 53%–67% with a recall varying between 51%–81%. The performance of the classifiers can be considered as moderate, as they perform better compared to the random classification, and do not show very high accuracy.

V. Discussion

We showed using a single snapshot of installed apps; gender of smartphone users can be predicted with an accuracy around 70%. In this Section, we compare our results with other online data sources that can be used for predicting the gender of the users.

In Table 4, we compare our results with the work of Hu et al. [14] in terms of precision and recall. Performance of our classifier was lower than the classifier built based on users’ web page clicks. However, prediction of male users showed a close precision value only at a 10% drop in recall. Hu et al. used a one-week long dataset, whereas we used only a single snapshot of installed apps.

Table 5, compares our work with prior research, such as study of *Facebook* likes [15], language style in *Facebook* status updates [24], search queries [5], client side web browsing logs [11] and mobile phone usage monitoring [31]. Our method’s accuracy is sim-

Table 4: Performance comparison
Apps (*Category + Content - 1K*) vs. web page
clicks

Related Work	Gender	Precision	Recall
Wep page click logs [14]	Male	0.791	0.810
	Female	0.805	0.782
Apps - <i>Category+Content (1K)</i>	Male	0.767	0.701
	Female	0.604	0.702

Table 5: Performance comparison
Apps (*Category + Content - 1K*) vs. other data
(NA* - Result of the performance metric not
available)

Related Work	Accuracy	AUC
<i>Facebook</i> likes [15]	NA*	0.93
<i>Facebook</i> status updates [24]	91.9%	NA*
Search queries [5]	NA*	0.803
Client side browser history [11]	$\approx 75\%$	≈ 0.85
Mobile phone usage & environment [31]	$\approx 82-85\%$	NA*
Apps - <i>Category+Content (1K)</i>	70%	0.74

ilar to techniques using client side browser history. Further the performance is close to the performance of search query and mobile usage monitoring. A difference of approximately 12%–15% in accuracy with the mobile usage monitoring method is interesting as monitoring usage activities on a mobile phone is a resource intensive task and considered as violating user privacy.

Comparisons show that in the smartphone context, the use of a snapshot of installed apps to predict gender provides performance comparable to other data sources and methods that require much richer datasets collected by tracking user activities over a period of time or collecting more intrusive personal data such as *Facebook* likes or status updates. We believe access to a much larger dataset of the users’ installed apps and the corresponding ground truth of gender will enable higher classifier performance as more advanced features can be generated. Moreover, it might be possible to infer other user attributes such as age, ethnicity, and income level following a similar method.

VI. Conclusion

This paper provided insights on various features related to installed apps on a smartphone, which can be used to predict the gender of the user. Using a dataset of over 200 smartphone users, we showed that by only observing a single snapshot of installed apps, it is possible to predict gender of the smartphone user with an average accuracy of approximately 70%. We compared our results with other data sources used for gender prediction, which involve data collection for

longer periods and/or require more personal data from users such as *Facebook* likes or their social media status updates. We showed that through a single snapshot of apps we could achieve comparable results to longer and more intrusive inference techniques.

References

- [1] Amazon Mechanical Turk. <https://www.mturk.com/mturk>.
- [2] Applicationinfo Android Developers. <http://developer.android.com>, 2013.
- [3] Language Detection API. <http://detectlanguage.com>, 2013.
- [4] D. Amitay. iOS App Detection. <http://www.ihasapp.com>, 2012.
- [5] B. Bi, M. Shokouhi, M. Kosinski, and T. Graepel. Inferring the demographics of search users: Social data meets search queries. In *WWW 2013 - Proceedings of the 22nd International Conference on World Wide Web*, pages 131–140, 2013.
- [6] D. Bolton. Why I stopped playing Candy Crush Saga. <http://news.dice.com>, 2013.
- [7] A. Chaabane, G. Acs, M. A. Kaafar, et al. You are what you like! information leakage through users’ interests. In *NDSS 2012 - Proceedings of the 19th Annual Network & Distributed System Security Symposium*, 2012.
- [8] A. Cocotas. Android grabs a record share of the global smartphone market. <http://au.businessinsider.com>, 2013.
- [9] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [10] S. Fiegerman. Apple’s App Store tops 1 million apps. <http://mashable.com>, 2013.
- [11] S. Goel, J. Hofman, and M. I. Siner. Who does what on the web: Studying web browsing behavior at scale. In *ICWSM 2012 - Proceedings of the 6th International Conference on Weblogs and Social Media*, 2012.
- [12] M. C. Grace, W. Zhou, X. Jiang, and A. Sadeghi. Unsafe exposure analysis of mobile in-app advertisements. In *WiSec 2012 - Proceedings of the 5th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 101–112, 2012.

- [13] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.
- [14] J. Hu, H. Zeng, H. Li, C. Niu, and Z. Chen. Demographic prediction based on user’s browsing behavior. In *WWW 2007 - Proceedings of the 16th International World Wide Web Conference*, pages 151–160, 2007.
- [15] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 110(15):5802–5805, 2013.
- [16] 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 and Society*, 17(3):286–300, 2014.
- [17] eMarketer Inc. Driven by Facebook and Google, mobile ad market soars 105% in 2013. <http://www.emarketer.com>, 2014.
- [18] Ericsson Inc. Ericsson Mobility Report. <http://www.ericsson.com>, 2013.
- [19] M. Meeker. Internet Trends 2014. <http://www.kpcb.com/internet-trends>, 2014.
- [20] A. Mukherjee and B. Liu. Improving gender classification of blog authors. In *EMNLP 2010 - Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 207–217, 2010.
- [21] J. Otterbacher. Inferring gender of movie reviewers: exploiting writing style, content and metadata. In *CIKM 2010 - Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 369–378, 2010.
- [22] J. W. Pennebaker, M. E. Francis, and R. J. Booth. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, 71:2001, 2001.
- [23] J. Schler, M. Koppel, S. Argamon, and J. Pennebaker. Effects of age and gender on blogging. In *AAAI Spring Symposium - Technical Report*, volume SS-06-03, pages 191–197, 2006.
- [24] H. A. Schwartz, J. C. Eichstaedt, M. L. Kern, L. Dziurzynski, S. M. Ramones, M. Agrawal, A. Shah, M. Kosinski, D. Stillwell, M. E. P. Seligman, and L. H. Ungar. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS ONE*, 8(9), 2013.
- [25] S. Seneviratne. Apptronomy: You are what your apps are. <https://play.google.com/store/apps>, 2013.
- [26] S. Seneviratne, A. Seneviratne, P. Mohapatra, and A. Mahanti. Predicting user traits from a snapshot of apps installed on a smartphone. *ACM SIGMOBILE Mobile Computing and Communications Review*, 18(2):1–8, 2014.
- [27] K. Wagstaff. Men are from Google+, women are from Pinterest. <http://techland.time.com>, 2012.
- [28] C. Warren. Google Play hits 1 million apps. <http://mashable.com>, 2013.
- [29] J. Yan, N. Liu, G. Wang, W. Zhang, Y. Jiang, and Z. Chen. How much can behavioral targeting help online advertising? In *WWW’09 - Proceedings of the 18th International World Wide Web Conference*, pages 261–270, 2009.
- [30] Y. Yang and J. O. Pedersen. Icm1 1997 - a comparative study on feature selection in text categorization. In *Proceedings of the Fourteenth International Conference on Machine Learning*, pages 412–420, 1997.
- [31] J. J. Ying, Y. Chang, C. Huang, and V. S. Tseng. Demographic prediction based on users mobile behaviors. *Mobile Data Challenge*, 2012.