[HIGHLIGHTS]

Sha Zhao *Department of Computer Science, Zhejiang University, China*
Julian Ramos *Human-Computer Interaction Institute, Carnegie Mellon University, US*
Jianrong Tao, Ziwen Jiang, Shijian Li, Zhaohui Wu, Gang Pan* *Department of Computer Science, Zhejiang University, China*
Anind K. Dey* *Human-Computer Interaction Institute, Carnegie Mellon University, USA* **corresponding authors*

Editors: Nic Lane and Xia Zhou

WHO ARE THE SMARTPHONE USERS?

Identifying user groups with apps usage behaviors

Photo, istockphoto.com

Understanding smartphone users is fundamental for creating better smartphones, improving the smartphone usage experience, and generating generalizable and reproducible research. However, smartphone manufacturers and most of the mobile computing research community make a simplifying assumption that all smartphone users are similar or, at best, constitute a small number of user types, based on their behaviors. Manufacturers design phones for the broadest audience and hope they work for all users. Researchers mostly analyze data from smartphone-based user studies and report results without accounting for the many different groups of people that make up the user base of smartphones. We challenge these elementary characterizations of smartphone users and show evidence of the existence of a much more diverse set of users. We analyzed one month of application usage from 106,762 Android users and discovered 382 distinct types of users based on their application usage behaviors, using our own two-step clustering and feature ranking selection approach, and gave a meaningful label to the users in each cluster, such as **Screen Checkers** and **Young Parents**. Our results have profound implications on the reproducibility and reliability of mobile computing studies, design and development of applications, determination of which apps should be pre-installed on a smartphone and, in general, on the smartphone usage experience for different types of users.

Excerpted from "Discovering different kinds of smartphone users through their application usage behaviors," from UbiComp 2016, *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, with permission. http://dl.acm.org/citation.cfm?id=2971696 © ACM 2016

The number and popularity of mobile applications is rising dramatically at the same time as there is an accelerating rate of adoption of smartphones. Meanwhile, a great number of research studies in recent years have sought to understand users' smartphone application usage behaviors, such as how individuals install apps [1], how many daily interactions they have with apps [2], how long the average application session lasts [3], which apps are frequently used together [4]. This past work analyzed application usage in the aggregate or explored the range across users. They have not explored any differences in application usage behavior between groups of users. They mostly treat all smartphone users as similar, which is a nice simplifying assumption for app developers, phone manufacturers and mobile carriers, except that it does not reflect reality.

The simplifying assumption that all smartphone users are similar or, at best, constitute a small number of user types, has led to a lack of reproducibility and generalizability in smartphone studies, as highlighted by Church *et al* [5], who suggest that this could be due to the existence of different user sub-populations among the larger smartphone user population. For example, Jones *et al.* [6] identified three distinct clusters of users based on their app revisitation patterns, by analyzing three months of application launch logs from 165 users. Banovic *et al.* [7] identified four types of users by analyzing 27 users' actions on emails displayed on their lock screen. Both studies are promising and very motivating for us in suggesting that there are at least 4 kinds of users, however their scope is limited by only examining a relatively simple behavior (revisitation for [6] and actions performed on emails for [7]) for relatively small populations. We challenge the simplifying assumption and attempt to show there are many more complex and diverse behaviors that make up the smartphone user population [8].

PROBLEM DEFINITION

Our goal is to discover different smartphone user groups through their app usage behaviors. A user group refers to a group in which the users are similar in a certain aspect. Users in different groups may have different needs and preferences, which makes it natural for the app usage from their smartphones to be different. Users with similar app usage behaviors may have similar needs or preferences, and form a group. Discovering smartphone user groups based on their app usage behaviors means to group users with similar app usage behaviors on smartphones. In the case of clustering, it groups a set of objects in such a way that objects in the same groups are most similar [9]. Particularly, it is appropriate to explore the similarity among the data points and to obtain an intuitive interpretation of each cluster embedded in the high-dimensional space. Thus, the task of discovering user groups is smoothly transformed into a clustering problem.

APPROACH TO IDENTIFY USER GROUPS

Users in different groups may have different needs and preferences, which makes it natural for the app usage from their smartphones to be different. Users with similar app usage behaviors may have similar needs or preferences and form a group. To solve this aggregation problem, we used unsupervised learning methods to cluster users. Please refer to our UbiComp'16 paper [8] for the details.

To cluster users based on their similarity of app usage behaviors, users were represented as vectors by their app usage in different time periods. With the very high number of distinct apps, it was impossible to directly compare application usage behaviors from the apps themselves. Instead, individual applications are transformed to their respective app category. This transformation allows us to compare application usage across different users interacting with different applications from the same category. Users can then be represented as vectors with the percentage of app category usage weight in different time periods, and these vectors can be used for clustering.

Our clustering method consists of two steps combining K-Means and MeanShift.

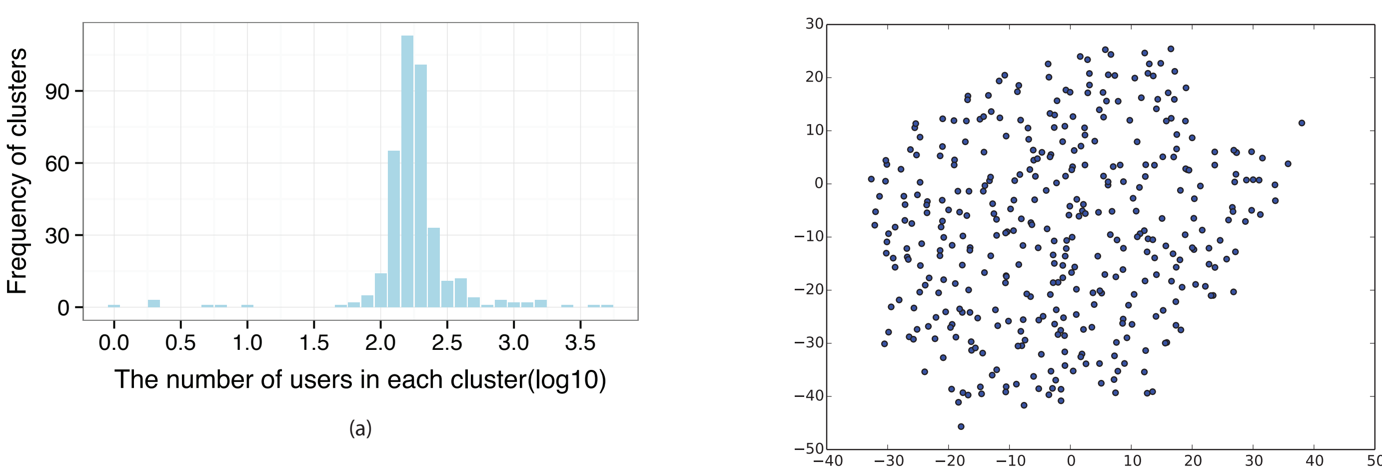

FIGURE 1. (a) Quantity of clusters with respect to number of users; (b) t-SNE representation of the cluster centroids.

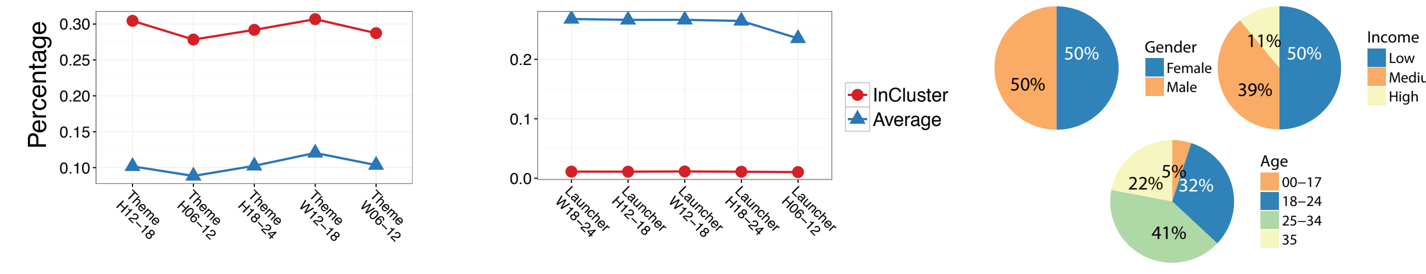

(a) Top 5 highest idiosyncratic features (b) Top 5 lowest idiosyncratic features (c) Proportion of users in each demographic attribute

FIGURE 2. The second-biggest cluster (3,814 users).

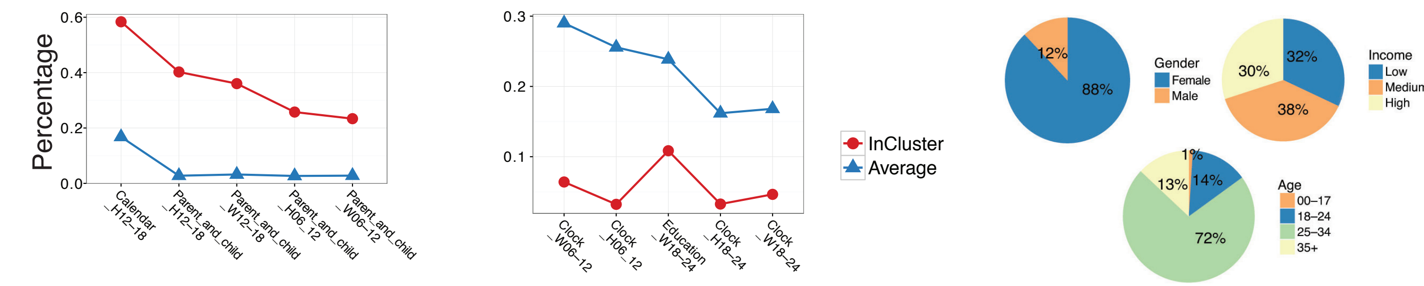

(a) Top 5 highest idiosyncratic features (b) Top 5 lowest idiosyncratic features (c) Proportion of users in each demographic attribute

FIGURE 3. Cluster219 (164 users).

In the first step, K-Means is used to reduce the available data points into a smaller representation and, in the second step, we used the centroids generated by K-Means as input for MeanShift, which automatically estimates the number of clusters in the data. The results from our two-step clustering were evaluated using a performance metric that considers penalties for complexity and non-uniform distribution of users across clusters.

We also proposed a novel *feature-ranking* scheme to identify meaningful *idiosyncratic features* for the selected clusters, which were those features that help distinguish each cluster from the "average user," a fictional user whose usage characteristics are equal to the average application usage of all users in the dataset. Then, we assigned a meaningful label to the users in each cluster based on their distinctive app usage behaviors.

FINDINGS

We identified different types of users from a dataset from the smartphones of 106,762 users from multiple provinces in China. For each smartphone, the dataset contains hourly updates on the 10 most recently used apps for the month of September 2015. All the apps were grouped into 29 semantic app categories. With our two-step clustering method, we obtained 382 clusters. By analyzing the clustering results, we found that:

1. Most clusters consist of 100-300 users. Figure 1 (a) shows the quantity of clusters with respect to the number of users. Among the 382 clusters, there are 85 percent of clusters consisting of 100-300 users, and only 9 clusters consisting of more than 1,000 users.

2. The centroids of clusters were nicely separated by our approach. For the visualization purpose, we used the t-SNE [10] transformation to represent the 382 cluster centroids in two dimensional space, shown in Figure 1 (b). The centroids appear to be nicely separated, giving a visual indication that our clustering approach was successful.

3. Most of the user groups have very distinctive app usage behaviors. With our feature-ranking scheme, the top 5 highest and lowest idiosyncratic features were selected to distinguish the app usage behaviors for the users in each cluster. For example, for the 3814 users in the second biggest cluster ranked by size, all of the top 5 highest idiosyncratic features are about theme apps, and the top 5 lowest idiosyncratic features are about launch apps, shown in Figure 2. It means that, compared with the users in other clusters, the users in this cluster use theme-related apps very often but rarely use launcher apps. They frequently wake up their smartphone but rarely unlock the screen and enter the main interface. We hypothesize that they are waking up their phones just to check the time or to see if there are any notifications. Thus, we tagged the users in this cluster as **Screen Checkers**. As shown in Figure 2 (c), this cluster has an even gender distribution, with the largest proportion of users being between 25-34 years old (41%) and half of the users having low income.

We also found some smaller yet interesting clusters composed of mainly male or female users. For example, **Cluster219** consists of

164 users, of which 88% are female users, and 72% are in age range from 25 to 34. It was found that the users in the cluster use the category of parent_and_child more frequently than others in the morning on both holidays and workdays, shown in Figure 3. More specifically, in the top 5 highest idiosyncratic features, there are 4 features related to the category of Parent_ and_child. The applications in this category are about how to raise a baby, how to help pregnant women, *etc*. Given the proportion of users that are 25-34 years old, we can label the users in this cluster as **Young Parents**.

4. Demographics have an important impact on how users use applications on their smartphones. It was found that gender, income level and age have a strong impact on the usage behaviors. For example, users with high-income levels use the categories of travel and health_and_fitness more frequently on holidays. Female users with ages between 25-34 years old use the category of parent_and_child more often in the daytime on both holidays and workdays, while female users between 0-17 use education-related apps more frequently in the evening on workdays [8].

CONCLUSIONS

By analyzing the app usage behaviors of more than 100,000 Android users, we discovered 382 distinct kinds of users using the proposed approach. With the selected general features and idiosyncratic features, we can roughly identify and semantically label each user group. We found that demographics have an important impact on how users use applications on their smartphones.

The several distinct groups of users we discovered prove that the assumption or simplification that all smartphone users are similar and that they can be treated as a uniform group, is not true. Smartphone manufacturers, mobile carriers, app developers and anyone who impacts the kinds of apps that are placed on phones, what apps are provided on phones and how people select apps to execute, can no longer treat users as if they all fit into one big group. The findings can be used in practice, such as the reproducibility and reliability of mobile computing studies, design and development of applications, determination of which apps should be pre-installed on a smartphone and, in general, on the smartphone usage experience for different types of users. ■

Sha Zhao is a PhD student at the College of Computer Science and Technology, Zhejiang University. Her research interests include ubiquitous computing, mobile sensing, and data mining. She received the Best Paper Award of ACM UbiComp '16. Zhao visited the Human-Computer Interaction Institute at Carnegie Mellon University as a visiting PhD student from September 2015 to September 2016.

Julian Ramos is a PhD student at the Human-Computer Interaction Institute at Carnegie Mellon University. His research interests lie at the intersection of machine learning and human-computer interaction for modeling human behavior. Ramos has a BS in mechatronics engineering from Militar Nueva Granada University. Contact him at ingenia@cs.cmu.edu.

Jianrong Tao is a Master's student at the College of Computer Science and Technology, Zhejiang University. His research interests include machine learning and data mining. Tao earned his BSc degree in computer science from Huazhong University of Science and Technology in 2014.

Ziwen Jiang is a Master's student at the College of Computer Science and Technology, Zhejiang University. His research interests include ubiquitous computing, machine learning and data mining. Jiang received his BSc degree in computer science from Zhejiang University in 2015.

Shijian Li is an associate professor with the College of Computer Science and Technology, Zhejiang University. His research interests include sensor networks, ubiquitous computing, and social computing. In 2010, he was a Visiting Scholar with the Institute Telecom SudParis, Evry, France. Li received his PhD degree from Zhejiang University, Hangzhou, China, in 2006.

Zhaohui Wu is a professor with the college of Computer Science and Technology, Zhejiang University. His research interests include intelligent systems, semantic grid, and ubiquitous embedded systems. He earned a PhD degree in computer science from Zhejiang University in 1993.

Gang Pan is a professor at the College of Computer Science and Technology at Zhejiang University. His interests include pervasive computing, computer vision, and artificial intelligence. He has co-authored more than 100 refereed papers, and has 25 patents granted. He has received three IEEE/ACM conference best paper awards (e.g. ACM UbiComp'16).

Anind K. Dey is a professor at Carnegie Mellon University's Human-Computer Interaction Institute. His research interests lie in the intersection of human-computer interaction and ubiquitous computing, including context-aware systems and more usable ubicomp systems. He has a PhD in computer science from the Georgia Institute of Technology. Contact him at anind@cs.cmu.edu.

REFERENCES

[1] Zhao, Sha, Gang Pan, Yifan Zhao, et al. Mining user attributes using large-scale app lists of smartphones. *IEEE Systems Journal*, 2017.11(1): p. 315-323.
[2] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, et al. Diversity in smartphone usage. in *Proceedings of the 8th International Conference on Mobile systems, Applications, and Services*. 2010. p. 179-194.
[3] Matthias Böhmer, Brent Hecht, Johannes Schöning, et al. Falling asleep with Angry Birds, Facebook and Kindle: A large scale study on mobile application usage. In *Proceedings of the 13th International Conference on Human computer interaction with mobile devices and services*. 2011. p. 47-56.
[4] Vijay Srinivasan, Saeed Moghaddam, Abhishek Mukherji, et al. Mobileminer: Mining your frequent patterns on your phone. in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 2014. p. 389-400.
[5] Karen Church, Denzil Ferreira, Nikola Banovic, et al. Understanding the challenges of mobile phone usage data. In *Proceedings of the 17th International Conference on Human Computer Interaction with Mobile Devices and Services*. 2015. p. 504-514.
[6] Simon L. Jones, Denzil Ferreira, Simo Hosio, et al. Revisitation analysis of smartphone app use. in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 2015. p. 1197-1208.
[7] Nikola Banovic, Christina Brant, Jennifer Mankoff, and Anind Dey. ProactiveTasks: the short of mobile device use sessions. In *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services*, 2014. p. 243-252.
[8] Zhao, Sha, Julian Ramos, Jianrong Tao, et al. Discovering different kinds of smartphone users through their application usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2016, p. 498-509.
[9] Jain, Anil K., M. Narasimha Murty, and Patrick J. Flynn. Data clustering: a review. *ACM computing surveys*, 1999(3): p. 264-323.
[10] Laurens Van der Maaten and Geoffrey Hinton, Visualizing data using t-sne. *Journal of Machine Learning Research*, 2008. 9(2579-2605): p. 85.