Using Opportunistic Face Logging from Smartphone to Infer Mental Health: Challenges and Future Directions

Rui Wang

Dartmouth College rui.wang@cs.dartmouth.edu

Andrew T. Campbell Dartmouth College campbell@cs.dartmouth.edu

Xia Zhou Dartmouth College xia@cs.dartmouth.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

Ubicomp/ISWC'15 Adjunct, September 07-11, 2015, Osaka, Japan © 2015 ACM. ISBN 978-1-4503-3575-1/15/09...\$15.00 DOI: http://dx.doi.org/10.1145/2800835.2804391

Abstract

We discuss the opportunistic face logging project that collected 5811 opportunistic photos using the front-facing camera from the smartphones of 37 participants over a 10-week period. We present our experiences using computer vision and human labeling approaches on the photos to assess mental health of the participants. Finally, we discuss our insights, challenges, and future directions.

Introduction

Mental illness has a profound impact on people's functioning, health, and guality of life [15]. Detecting early warnings of depression or other mental illness is challenging. Recently, there have been a number of studies [5, 7, 14, 26, 31] using sensor data from smartphones to assess mental health by looking at changes in social interaction, speech patterns, sleep and physical activity. In this paper, we ask the question: can a sequence of opportunistic photos taken of a user by the front-camera of their phones infer depression, stress, affect, and flourishing? People particularly the younger generation – interact with their phones throughout the day - from the moment they wake up to the point they go to sleep. If every time you pick up your phone, an inference system analyzes a naturalistic photo of your face, could it detect your

mental state? This is the question that drives this initial study presented in this paper. These opportunistic photos represent glimpses or data points of a person's experience and could be potentially revealing if mental health could be inferred.

The opportunistic face logging study uses 5811 opportunistic photos collected from 37 smartphone users over 10 weeks using the front-facing camera of the phone. The data was collected as a part of StudentLife study [31] but is not a part of the publicly released dataset because of privacy reasons. The goal of the StudentLife project is assessing mental health, academic performance and behavioral trends of college students using smartphones [31]. Photos of the students were captured opportunistically every time they responded to in-situ ecological momentary assessment (EMA) [28] guestions. These EMAs were scheduled multiple times per day over the 10-week study period, allowing the system to record the user in response to specific EMAs, including self-reports on stress, affect, and workload. As part of the study participants completed a number of pre and post surveys including PHQ9 depression scale [20, 21, 29], UCLA loneliness scale [27], positive and negative affect schedule (PANAS), perceived stress scale (PSS) [8], big five personality [17], and flourishing scale [11]. These surveys collectively measure mental health in four dimensions: depression, stress, flourishing, and loneliness. EMAs were triggered at different points of a day. For all EMA sequences, the user was first prompted by the Photographic Affect Meter (PAM) [25] to capture their mood as shown in Figure 1 (a), followed by a short survey (*e.g.*, a single-item stress survey). While we changed the scheduled short survey associated with each EMA

throughout the day (e.g., sleep, stress, workload, social, stress), PAM was always common to every EMA: the common presentation of PAM then a survey was used throughout the 10-week study.

When were photos taken of the participants? When the user consented to the study, they were informed that the system would take a single photograph from the front-facing camera at some point between the presentation of PAM on the phone and completing the short survey that followed PAM. As shown in Figure 1(a), PAM asks "Touch how you feel right now". Users can swipe through multiple photos in PAM and select the photo that best represent their current mood. PAM has been validated against positive and negative affect schedule (PANAS) [25]. Once the user completes their interaction with PAM, a short survey (e.g., stress) is shown on the screen. After the user answers the survey, the application exits. We used Google PACO to schedule and manage firing PAM and EMAs on smartphones. What the user did not know in advance was that we took the single photo at the same moment they interacted with the PAM application, answering how they felt right now. We were interested in how someone might feel internally and how they looked externally. Only one student out of the cohort disabled the photo option in the study. Students were told that they could delete any photos anytime during the study. Only two students did this. As a result we have a large and novel dataset of pictures of students taken across the 10 weeks of a busy term. In total we have collected over 5811 photos from 37 participants. The study is approved by IRB. Because of the IRB and to protect the privacy of the students, we cannot show photos in this paper. What we do show is some sample photos of some of the research staff (i.e., a professor and two



research students) just to illustrate the quality of these opportunistic images. The research staff are not part of the study. They, however, were also using the application for the same 10 week period. As shown in Figure 1 (b), the opportunistic photos are challenging to work with because of the different lightening conditions, guality of the image, and orientation of the camera in reference to the user.

Computer Vision

We first examine the feasibility of using computer vision methods to automatically infer a person's mental health from the opportunistic photos. Specifically, we explore three representative computer vision approaches:

facial expression recognition, facial landmarks recognition, and low-level image feature extraction. We discuss the limitations of using these computer vision algorithms for this dataset.

We start with applying a facial expression recognition method based on eigenfaces proposed in [3, 19]. This method works as follows. It first preprocesses the images by normalizing the image size, equalizing the histogram, and converting the image to gay scale. It then calculates the eigenfaces from the preprocessed image and extracts face descriptors from the eigenfaces. Finally, it infers the facial expression using a neural network. After applying this method to our

dataset, we find that it fails to recognize almost all of the facial expressions. We further test the facial expression recognition service provided by [1]. We handpick 10 photos as a test set. These photos are taken in very good lighting conditions (i.e., no blurred faces) where users shows obvious emotional expressions. However, the recognition service [1] fails to recognize the facial expressions even for these good face images. The poor facial expression recognition performance we observed can be attributed to two reasons. First, taken in naturalistic settings, most opportunistic photos have poor image quality. We used the Nexus 4 phone for the study with an 1.3 MP image sensor. The photos were taken with the resolution of 240 by 320, and in regular mode (i.e., no HDR was available). Many photos suffer from over/under exposure and thus the user's face is either too bright or too dark, while others suffer from image blur caused by motion. Furthermore, a large portion (41.9%) of the photos do not contain the participant's full face but potions of the face. Face occlusion is common. Also the photos are taken from different angels so that the face is usually distorted either by the perspective of the face or by the lens distortion. Second, most existing facial expression recognition algorithms are built upon datasets collected from actors posing expressions [22, 23]. The actor's facial expressions are usually exaggerated. The facial expressions in our dataset, however, are more subtle and taken under naturalistic conditions.

We next turn to detecting the facial landmarks, which can be used for face alignment, facial expression analysis, and other potential applications [13]. We hypothesize that the changes in long term facial expression could be a predictor of the participant's

mental health. The literature [9] has shown that depressed patients are less expressive than non-depressed patients. Thus, we extract the facial landmarks from the photos, align the facial landmarks across the photos of each participant, and calculate the variance of the landmarks' positions as a mental health predictor. We apply the existing facial landmark recognition method to find main facial landmarks (i.e., eyebrows, eyes, nose, mouth, and chin) from the participants' face photos. Specifically, we use the face recognition algorithm implemented by [12] to find the face in a photo, and then apply the facial landmarks recognition proposed in [6], which uses robust cascaded pose regression to find all the landmarks and is able to detect landmarks in the presence of occlusions. We scale, rotate, and center all the landmarks according to the positions of eyes and nose, such that we can compare the facial landmarks in different photos. The relative positions of eyes and nose are stable and not likely to be affected by expressions. Therefore, the rotated and scaled landmarks preserve the user's expressions. For each participant, we calculate the average variance of each landmark's position as a mental health predictor. We note, that a number of students in the study were moderately depressed or severely depressed from the pre-post PHQ9 depression scale. However, we fail to find any correlation between the facial landmark predictors and the mental health ground truth. This is because the algorithm can detect landmarks in only 12.5% of the photos. We believe that the low detection rate is due to the poor image guality described previously.

Finally, we extract low-level image features from each photo: brightness, contrast, and saturation. Our

hypothesis is that for participants who suffer from mental health disorders, the environment (e.g., background lighting conditions) where they use their phones might be different from participants with different depression scale scores. For example, depressed patients may have irregular sleeping schedules [4], so they may stay in bed longer, making it more likely to capture photos in dark environments or when they are lying down. We run correlation analysis to examine if any image features correlate with any of the pre and post mental health survey ground truth. For each participant, we calculate the mean and standard deviation of the low-level image features as two mental health predictors. However, we do not find any correlations either.

Human Labeling

In order to understand why the vision approaches do not work and to understand if the pictures are indeed useful to infer the participants' mental health, we use human coders to label the pictures. Our goal is to figure out what aspects of the photos may reveal useful information about mental health. In what follows, we present our image coding system, coding task design, and our findings.

Picture Coding

We build our picture coding system based on PyBossa [2], which is an open source crowdsourcing platform.

Table 1: Picture label questions and label results distribution

Question	Options	
Is the image blurred?	Yes (75.4%), Medium (14.1%), No (10.5%)	
Is the complete picture too dark?	Yes (15.2%), Medium (19.7%), No (65.1%)	
How much of the student's face is in the picture?	Full face (58.1%), Half face (17.2%), Quarter face (7.0%), None(17.7%)	
In general, the student's facial expression is?	Positive (10.3%), Neutral (47.4%), Negative (8.5%), Unknown (33.8%)	
What is the student's expres- sion?	Neutral (46.9%), Happiness (8.8%), Sadness (3.8%), Surprise (1.4%), Fear (0.2%), Anger (1.6%), Disgust (1.0%), Contempt (1.0%), Unknown (35.3%)	
What is the stu- dent doing?	Laying down (4.3%), Sitting (38.4%), Walking/Standing (18.9%), In a vehi- cle (0.7%), Unknown (37.7%)	
Does the student look like they have just awoken from sleep?	Yes (9.4%), No (54.6%), Unknown (36.0%)	
Does the student look stressed?	Yes (9.5%), No (52.2%), Unknown / Not applicable (38.3%)	
Does the student look depressed?	Yes (5.2%), No (55.3%), Unknown / Not applicable (39.5%)	
Does the student look tired?	Yes (18.3%), No (43.5%), Unknown / Not applicable (38.2%)	

We design the picture coding questions to cover four dimensions of the pictures: image quality, facial

expression, context, and mental health ground truth. The questions are shown in Table 1. The image quality questions help us understand the challenges of collecting opportunistic photos and may provides clues to advancing the capturing technology. The facial expression questions are the key questions about inferring mental health from the opportunistic pictures. The context questions ask about what the student is doing when the picture is taken. The mental health ground truth questions ask the coder to label if the participant in the picture appear to be stressed, depressed, or tired.

We recruited 20 human coders to label the 5811 pictures. Amongst the coders, 15 of them have a computer science background, two coders have psychology backgrounds, one coder has an engineering background, and two coders have non-technical backgrounds. In order to evaluate the consensus among coders, we assigned 300 pictures for each coder with 10 pictures overlapping with other coders; that is, we collected 6000 labels for the pictures and 189 pictures have been labeled by two coders. We provide labeling instructions to each human coder. The instruction explains the meaning of each question, and gives examples of how to label pictures.

In what follows, we discuss the distribution of the labels for all the pictures as shown in Table 1. The distributions give information about the overall picture quality, students' facial expressions, what the students are doing, and how our human coders perceive the students' psychological states. In term of image quality, 75.4% of them are labeled as not blur at all, 65.1% are labeled as in good lighting condition, and 58.1% are labeled as having full faces in the picture. The results

show around 30% of the pictures suffer from poor image quality problems. Capturing the participants' full face is also problematic. In term of facial expressions, 47.4% of the pictures are labeled as neutral in the general expression guestion and 46.9% are labeled as neutral in the more detailed facial expression question. Over 30% of the pictures are labeled as unknown in both questions due to either poor image quality or the coder cannot identify the participant's expression. In term of the student's activity, 38.4% of the pictures are labeled as sitting and 18.9% are labeled as walking/standing. In term of sleep related labels, 4.3% of the pictures are labeled as laying down and 9.4% are labeled as looking like just awoke. Similar to facial expression labels, over 30% of the pictures are labeled as unknown for the two questions. In term of the mental health ground truth labels, 38% of the pictures are labeled as unknown for the three question. About 50% of the pictures are labeled as negative in stress, depressed, and tired questions. In summary, we see that the pictures suffer from poor quality in term of blurriness and too dark, which leads to many unknown labels in facial expression, activity, and mental health ground truth questions.

We calculate the consensus of picture labels by comparing the labels given to the same picture by different coders. Specifically, we compare each question's labels given by two coders, and mark 1 if the labels are the same or 0 otherwise. Then, we calculate the percentage of 1s for all the questions as the consensus. For image quality questions, the consensus is 68.75% for blurriness, 67.19% for lighting, and 76.34% for face area. For facial expression questions, the consensus is 76.12% for the general question and 67.97% for the detailed expression question. Since most of the labels for facial expression are neutral and unknown, we removed pictures whose expressions are labeled as neutral or unknown by both coders, and calculate the consensus for the rest of the labels. The consensus drops to 36% for the general facial expression question, and 24.3% for the detailed question. For activity context questions, the consensus for activity is 59.6% and the the consensus for woke up from sleep is 74%. The consensus for the mental health ground truth questions are all over 72%. The consensus shows that the overall labeling quality is acceptable. However, our coders cannot agree on facial expressions other than neutral and unknown.

Correlation Analysis

Table 2: Correlations between picture labels and the mental health ground truth

label	ground truth	r	р
general expression - neutral	PHQ-9 pre	-0.39	0.032
laying down	PHQ-9 pre	0.41	0.023
laying down	PHQ-9 post	0.51	0.006
woken from sleep	PHQ-9 pre	0.44	0.016
not woken from sleep	PHQ-9 pre	-0.37	0.044
expression - contempt	PSS pre	-0.37	0.041
expression - sadness	PSS pre	-0.42	0.018
expression - surprise	PSS post	0.42	0.023
woken from sleep	PSS post	0.46	0.013
general expression - neutral	flourising post	-0.40	0.037
expression - neutral	flourising post	-0.50	0.008
not depressed	flourising post	-0.44	0.021
not stressed	flourising post	-0.41	0.035
in a vehicle	extroversion	0.40	0.034
expression - anger	extroversion	0.39	0.041
stressed	extroversion	0.41	0.032
not tired	neuroticism	0.40	0.040

In what follows, we present the correlations between the mental health ground truth and and the labels given by human coders. We calculate the percentage of each picture label (i.e., each option in each question). It tells us the percentage of pictures that have been labeled as happy, sad, stressed, or depressed. We then calculate the Pierson correlation between the percentage labels and the mental health ground truth. The significant correlations are shown in Table 2. The results show that depression correlates with the students' facial expression and activity (i.e., laying down, sleep); stress correlates with facial expressions and activity; flourishing correlates with facial expressions and coder perceived depression and stress; personality traits correlate with activity, facial expression, and coder perceived stressed and tired. Some of the correlations are intuitive. We, however, cannot explain other correlations. Irregular sleep pattern is one depression symptom. We find students who are laying on a bed or just woke up are more likely to be more depressed. One possible explanation is that more depressed student use their phone more before, during and after sleep. Students who appear to be more depressed and stressed are more likely to be less flourishing, which is in line with the definition that people who are more flourishing have more psychological strength.

Future Directions and Concluding Remarks

In this section, we discuss future directions to automatically assessing mental health using opportunistic photos from smartphones taken in natural environments.

Improving the image quality will greatly help computer vision algorithms to extract mental health cues from photos. More and more Android phones have

optimized their cameras to better capture faces. The new Google Camera app would automatically focus on human faces and adjust exposure to match lighting conditions on the face. To avoid blurred images, we would use the accelerometer to detect motion and determine if we should discard the picture according to the exposure parameters. That is, higher shutter speeds are more tolerant to motion whereas slower shutter speeds are sensitive to motion. The phone would try to capture multiple images or video and keep images that were taken with the least phone motion. The accelerometer reading data can also be used for removing blurriness [18]. We would like to see the user's full face in the image. Many Android phones have implemented face recognition in the camera hardware, and the face information can be accessed by the app. We would use this information to capture the image at the right time such that the user's full face is in the image.

Existing facial expression detection technologies can take advantages of face muscle motions from video to detect expressions more accurately. High frame-rate videos also provide opportunities for micro-expression recognition. However, due to various lighting conditions, video could be in low frame rate in low light environments. In this case, it would be better to capture a single image of the user.

Previous work [32] has shown that pictures taken when the user unlocks the phone are less likely to be able to capture the user's face. In our pilot study, the pictures were taken when the user is answering EMAs. We believe that we have a better chance to capture the user's face images when they are closely interacting with their phones. For example, people's faces would be more expressive when they are reading messages or interacting with social networks. Android provides API to get the running app. We would capture an image or a short video when the user is interacting with a specific set of apps.

Facial expression recognition is an active field of research. However, most of the existing work focuses on detecting facial expressions from a staged lab environment, which perform poorly on our dataset. Our human label results show that the participants' facial expressions in most of the pictures are neutral or subtle. However, our coders have reported that even neutral faces have some subtle expression traits. Therefore, the facial expression recognition algorithm used in opportunistic photo should have a higher sensitively and should be robust enough to work for various facial types, ethnicities, lighting conditions, blurriness, occlusions, and backgrounds.

Detecting facial expressions is not only about inferring the exact facial expressions (i.e., FACS). Detecting facial action units [30] provides more fine grained information about the user's mood. For example, [16] studied depression patients' facial action units. They found nonsuicidal depressive patients showed different expression of emotions compared with suicidal patients. If we can accurately recognize facial action units from images, we could investigate how to recognize emotions such as frustration, tired, stressed, and sleepy.

Our results have shown that users' context could be predictors of mental health. A user's context include activity, body posture, type of location. Vision provides more information to recognize the user context compared with sensor based methods. For example, our results show that depression is correlated with the body posture of laying down. We could apply body posture analysis [10, 24] to detect if the user is laying down.

In this paper, we presented a pilot study to assess mental health using opportunistic photos taken from smartphones. We discussed the limitations of the images themselves, existing computer vision techniques, and our efforts at human labeling and discovering links between the opportunistic photos and mental health outcomes. Finally, we presented some future directions in this area.

References

- [1] nviso developers platform. https://developer.nviso.net.
- [2] Pybossa. http://pybossa.com/.
- [3] AGARWAL, M., JAIN, N., KUMAR, M. M., AND AGRAWAL, H. Face recognition using eigen faces and artificial neural network. *International Journal of Computer Theory and Engineering 2*, 4 (2010), 1793–8201.
- [4] AMERICAN PSYCHIATRIC ASSOCIATION. The Diagnostic and Statistical Manual of Mental Disorders: DSM 5. American Psychiatric Publishing, 2013.
- [5] BARDRAM, J. E., FROST, M., SZÁNTÓ, K., AND MARCU, G. The monarca self-assessment system: a persuasive personal monitoring system for bipolar patients. In *Proceedings of the* 2nd ACM SIGHIT International Health Informatics Symposium (2012), ACM, pp. 21–30.
- [6] BURGOS-ARTIZZU, X. P., PERONA, P., AND DOLLÁR, P. Robust face landmark estimation under occlusion. In *Computer Vision* (*ICCV*), 2013 IEEE International Conference on (2013), IEEE, pp. 1513–1520.
- [7] BURNS, M. N., BEGALE, M., DUFFECY, J., GERGLE, D., KARR, C. J., GIANGRANDE, E., AND MOHR, D. C. Harnessing context sensing to develop a mobile intervention for depression. *Journal of medical Internet research 13*, 3 (2011).

- [8] COHEN, S., KAMARCK, T., AND MERMELSTEIN, R. A global measure of perceived stress. *Journal of health and social behavior* (1983), 385–396.
- [9] COHN, J. F., KRUEZ, T. S., MATTHEWS, I., YANG, Y., NGUYEN, M. H., PADILLA, M. T., ZHOU, F., AND DE LA TORRE, F. Detecting depression from facial actions and vocal prosody. In Affective Computing and Intelligent Interaction and Workshops, 2009. ACII 2009. 3rd International Conference on (2009), IEEE, pp. 1–7.
- [10] COULSON, M. Attributing emotion to static body postures: Recognition accuracy, confusions, and viewpoint dependence. *Journal of nonverbal behavior 28*, 2 (2004), 117–139.
- [11] DIENER, E., WIRTZ, D., TOV, W., KIM-PRIETO, C., CHOI, D.-W., OISHI, S., AND BISWAS-DIENER, R. New well-being measures: Short scales to assess flourishing and positive and negative feelings. *Social Indicators Research 97*, 2 (2010), 143–156.
- [12] DOLLÁR, P. Piotr's Computer Vision Matlab Toolbox (PMT). http://vision.ucsd.edu/~pdollar/toolbox/ doc/index.html.
- [13] EFRATY, B. A., PAPADAKIS, M., PROFITT, A., SHAH, S., AND KAKADIARIS, I. A. Facial component-landmark detection. In Automatic Face & Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on (2011), IEEE, pp. 278–285.
- [14] FROST, M., DORYAB, A., FAURHOLT-JEPSEN, M., KESSING, L. V., AND BARDRAM, J. E. Supporting disease insight through data analysis: refinements of the monarca self-assessment system. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (2013), ACM, pp. 133–142.
- [15] HALFIN, A. Depression: the benefits of early and appropriate treatment. *The American journal of managed care 13*, 4 Suppl (2007), S92–7.
- [16] HELLER, M., AND HAYNAL, V. Depression and suicide faces. What the face revealsâĂŹ(Oxford, 1997) (1997), 398–407.
- [17] JOHN, O. P., AND SRIVASTAVA, S. The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook* of personality: Theory and research 2 (1999), 102–138.
- [18] JOSHI, N., KANG, S. B., ZITNICK, C. L., AND SZELISKI, R. Image deblurring using inertial measurement sensors. In ACM Transactions on Graphics (TOG) (2010), vol. 29, ACM, p. 30.

- [19] KAUR, M., VASHISHT, R., AND NEERU, N. Recognition of facial expressions with principal component analysis and singular value decomposition. *International Journal of Computer Applications 9*, 12 (2010), 36–40.
- [20] KROENKE, K., AND SPITZER, R. L. The phq-9: a new depression diagnostic and severity measure. *Psychiatric Annals 32*, 9 (2002), 509–515.
- [21] KROENKE, K., SPITZER, R. L., AND WILLIAMS, J. B. The phq-9. Journal of general internal medicine 16, 9 (2001), 606–613.
- [22] LYONS, M., AKAMATSU, S., KAMACHI, M., AND GYOBA, J. Coding facial expressions with gabor wavelets. In Automatic Face and Gesture Recognition, 1998. Proceedings. Third IEEE International Conference on (Apr 1998), pp. 200–205.
- [23] LYONS, M. J., BUDYNEK, J., AND AKAMATSU, S. Automatic classification of single facial images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 21*, 12 (1999), 1357–1362.
- [24] MOESLUND, T. B., AND GRANUM, E. A survey of computer vision-based human motion capture. *Computer vision and image understanding 81*, 3 (2001), 231–268.
- [25] POLLAK, J. P., ADAMS, P., AND GAY, G. PAM: a photographic affect meter for frequent, in situ measurement of affect.
- [26] RABBI, M., ALI, S., CHOUDHURY, T., AND BERKE, E. Passive and in-situ assessment of mental and physical well-being using mobile sensors.
- [27] RUSSELL, D. W. UCLA loneliness scale (version 3): Reliability, validity, and factor structure. *Journal of personality assessment* 66, 1 (1996), 20–40.
- [28] SHIFFMAN, S., STONE, A. A., AND HUFFORD, M. R. Ecological momentary assessment. *Annu. Rev. Clin. Psychol.* 4 (2008), 1–32.
- [29] SPITZER, R. L., KROENKE, K., WILLIAMS, J. B., GROUP, P. H. Q. P. C. S., ET AL. Validation and utility of a self-report version of prime-md: the phq primary care study. *Jama 282*, 18 (1999), 1737–1744.
- [30] TIAN, Y.-L., KANADE, T., AND COHN, J. F. Recognizing action units for facial expression analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 23*, 2 (2001), 97–115.

- [31] WANG, R., CHEN, F., CHEN, Z., LI, T., HARARI, G., TIGNOR, S., ZHOU, X., BEN-ZEEV, D., AND CAMPBELL, A. T. Studentlife: Assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (New York, NY, USA, 2014), UbiComp '14, ACM, pp. 3–14.
- [32] YOO, C., SHIN, J., HWANG, I., AND SONG, J. Facelog: Capturing user's everyday face using mobile devices. In Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication (New York, NY, USA, 2013), UbiComp '13 Adjunct, ACM, pp. 163–166.