

Towards a Mobile and Wearable System for Predicting Panic Attacks

Jonathan Rubin Hoda Eldardiry Rui Abreu Shane Ahern

Honglu Du Ashish Pattekar Daniel G. Bobrow

{Jonathan.Rubin, Hoda.Eldardiry, rui, Shane.Ahern, Honglu.Du, Ashish.Pattekar, bobrow}@parc.com
Palo Alto Research Center, Inc. California, United States

ABSTRACT

In this paper, we present first steps towards a mobile and wearable system intended to help people who experience regular and spontaneous panic attacks due to panic disorder. The goal of the system is to predict oncoming panic attacks and to deliver in-the-moment interventions on a smartphone device. Interventions are intended to reduce symptom severity by enabling a user to respond to approaching panic episodes. An initial feasibility study is described where a small real-world data set was collected. Personalized prediction models were trained which take, as input, physiological data and output a binary classification of either pre-panic or non-panic. We demonstrate proof-of-concept of episode prediction on this small dataset.

Author Keywords

Physiological monitoring; affective computing; mHealth

ACM Classification Keywords

I.2.1 Artificial Intelligence: Applications and Expert Systems; J.4 Social and Behavioral Sciences: Psychology

INTRODUCTION

This work investigates one way that mobile and wearable technology can be used to improve people's mental health and wellbeing. In particular, we focus on individuals suffering from panic disorder who experience regular panic attacks. In the United States alone, six million adults each year experience recurrent, spontaneous panic attacks due to panic disorder [12]. While the exact cause of panic disorder is not universally agreed upon, a number of studies have identified a relationship between abnormal respiration and panic attack onset [5, 10, 11]. Panic attacks are also a component of other anxiety disorders, including Post Traumatic Stress Disorder and Generalized Anxiety Disorder, which further increases the number of sufferers. During a panic attack, a sufferer may experience shortness of breath, increased heart rate, dizziness, chest pain and sweating, as well as feelings of unreality and fear of dying. Previous research [9, 15] has

shown that changes in physiology begin to take place *up to one hour* before the onset of a panic attack and sufferers are not aware that these changes are taking place until they experience the “*spontaneous*” symptoms of a panic attack. Many of these changes are physiological in nature, and hence detectable through wearable physiological measurements, such as heart rate, respiration rate, perspiration and skin temperature. In this paper, we describe first steps towards a system that combines a personal wearable device with a smartphone application for reacting to approaching panic attacks.

The components of the system are as follows:

1. A **wearable device**, which continuously monitors a user's physiological signals. In the present work we have used a chest-worn device. However, any wearable device capable of recording appropriate physiological information can be used, e.g. wrist-worn devices (provided reliability and precision were comparable to chest-worn devices).
2. A **mobile application**, which communicates via Bluetooth with the wearable device. Algorithms on a server, or on the mobile application itself, detect whether the physiological data indicates the user will experience an upcoming panic attack. In the event of a detection, mobile based interventions are delivered via a notification to the user's phone.
3. An **intervention**, which can be selected by the application user. One effective intervention is *respiratory biofeedback* [10]. In respiratory biofeedback, breathing and relaxation exercise instructions are delivered to a user who can visually monitor the effect breathing exercises have on their vital signs using their mobile device.

We detail the progress that has been made in developing the system described above. Components 1 and 2 are presented as part of the current work, while component 3 is left for future work. The main contribution of this paper is the presentation of results from an initial detection feasibility study, conducted on a small real-world dataset that was collected. Individuals who suffer from panic disorder were asked to use a wearable device that continuously monitored their vital signs in day-to-day life for up to three weeks. During that time subjects were asked to report, using a mobile device, whenever they experienced a panic episode. We describe the physiological dataset that was generated and how it has been used within a supervised anomaly detection algorithm to predict approaching panic episodes. We present experimental results on this small dataset and show that on average our predictions

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 the author(s) 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 '15, September 7–11, 2015, Osaka, Japan.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3574-4/15/09...\$15.00

<http://dx.doi.org/10.1145/2750858.2805834>

have a precision of 93.8% and a recall value of 83.8%, giving an overall F_1 score of 88.5%. While the dataset produced from this initial feasibility study is small, the results demonstrate the potential of predicting upcoming panic episodes and how this can be used within a mobile and wearable system.

RELATED WORK

There are many challenges involved with conducting real-world studies that attempt to continuously monitor the physiological signals of individuals. Motion artifacts and sensor dropout are issues that are required to be dealt with, as well as ensuring accuracy of reported information and subject compliance. Given these challenges, relatively few *in-the-wild* studies have been conducted with individuals who suffer from panic disorder. Instead, previous studies have attempted to record physiological responses to panic attacks in a laboratory setting. Both spontaneous [6, 18] and artificially induced [4, 7] panic attacks have been studied in-the-lab. However, in-the-lab studies are limited in the amount of time they are able to monitor the physiological signals of a subject – typically short amounts of time preceding and directly after episodes occur.

Despite the challenges of in-the-wild studies mentioned above, Meuret *et al.* [9] present one such example where individuals with panic disorder were monitored continuously in a real-world setting for 24 hours at a time. Physiological measurements were recorded for heart rate, PCO₂ (the amount of carbon dioxide in the blood), tidal volume (the volume of air moved into or out of the lungs), breathing rate and heart rate variability, among others. Subjects self-reported when a panic attack took place. In total, physiological data was captured for thirteen naturally occurring panic attacks. Using this data, [9] showed that a series of physiological changes occurred up to one hour before an individual experienced any symptoms related to a panic episode. During this period subjects were not aware of any symptoms until the moment they experienced an 'out-of-the-blue' panic attack. As part of the same study, Rosenfield *et al.* [15] used change-point analysis [13] to show that all of the vital sign measurements mentioned above, exhibited one or more change points in the hour preceding a reported panic episode. When compared with regular intervals, where a panic episode was not reported, none or relatively few change points were exhibited.

Our work extends the work of [9, 15]. First, the work presented in [9] required subjects to wear multiple types of intrusive physiological monitoring equipment, which would not be suitable for everyday life, whereas the physiological monitoring conducted in our work required a single wireless wearable device that is much less intrusive. Second, rather than monitor subjects over a 24 hour time period, subjects in our study wore the device for up to three weeks. Lastly (and most importantly), rather than just show that changes between the two states exist, the goal of our work is to develop a predictive algorithm that can be utilized within a mobile and wearable system to allow users to respond to approaching panic episodes via interventions delivered on their mobile device. As far as we are aware, the work we describe is the first attempt to create such a system.

DATA COLLECTION

To obtain data to train prediction models, subjects who suffer from panic disorder were sought to participate in a data collection study. Subjects were recruited from local Meetup groups where the group's purpose had to do with coping with panic and anxiety. In addition, a Google Adwords campaign was run and classified advertisements were placed on a website (Craigslis) seeking individuals who suffer from panic disorder. In total, ten Zephyr BioPatch™ wearable devices [3] were distributed to study participants. Each subject self identified as suffering from panic disorder and answered yes to the pre-screen question 'Do you suffer from frequent panic attacks'. Each subject was supplied a wearable device, along with instructions to download the corresponding mobile application that was used to report when panic episodes occurred. Subjects were asked to wear the device as often as possible for 3 weeks and to use the corresponding mobile application/widget to report when they experienced a panic attack. A widget that resided on the subject's mobile phone home screen was used to allow easy access to episode reporting. Participants agreed to tap the widget's start button as soon as they became aware of experiencing the symptoms of a panic attack. Once the start button was tapped this began a counter that recorded how long the attack lasted. After symptoms of the attack had subsided, participants tapped the widget's stop button to stop recording the episode. At the completion of an episode subjects were asked to rate the severity of symptoms they had just experienced. In total, subjects were asked to rate the severity of 15 symptoms using a severity rating of 1 (None) to 5 (Extreme). Fig. 1 depicts a screenshot from the mobile application. The symptoms listed are from the DSM-IV (Diagnostic and Statistical Manual of Mental Disorders) [2] and the Panic Disorder Severity Scale standard instrument [16].

The study was approved by Palo Alto Research Center's Institutional Review Board and all subjects signed informed consent forms.

Dataset Summary

Out of the ten individuals who signed up for the study, five were female, four male and one subject identified as trans-male. The minimum age was 19 and the maximum was 53. From the initial ten participants, three subjects did not log any physiological data. The remaining seven subjects logged data and four of the seven reported at least one panic episode while wearing the device. In total 19 panic episodes were recorded with corresponding physiological data and symptom severity information. The minimum number of reports made by an individual, who logged data, was 1 and the maximum was 12 panic reports made by one individual. On average each panic episode lasted approximately 3 minutes and 40 seconds.

Data Measurements

The Zephyr BioPatch™ device allows continuous monitoring of a user's physiological information and has a sampling frequency of 250Hz for ECG, 25Hz for breathing rate and 100Hz for (3-axis) acceleration. The following subset of data measurements were chosen for use within the analysis: heart rate (HR), breathing rate (BR), heart rate variability (HRV),



Figure 1. Screenshot of mobile reporting application

core temperature (Temp) and activity. Each of the above measurements were summarized at 1Hz. Activity was represented using vector magnitude units, calculated over the three axial acceleration magnitudes, i.e. $vmu = \sqrt{x^2 + y^2 + z^2}$. This results in approximate values for activity types [17], such as: no activity ($vmu = 0.0$), low activity ($vmu \leq 0.2$), moderate activity ($0.2 < vmu < 0.8$) and high activity ($vmu \geq 0.8$).

PANIC EPISODE PREDICTION

Data Processing

For each of the 19 reported panic episodes, 60 minutes of physiological data leading up to the episode start time was extracted. From this data, summary statistics were generated that captured the number and location of significant change points [13] in the time-series, as well as the mean and standard deviation for each of the recorded physiological measurements (HR, BR, HRV, Temp). Feature vectors were then composed consisting of the captured summary statistic variables. Equivalent feature vectors were also created for non-panic intervals by extracting the same amount of physiological data and deriving summary statistics from sessions when no panic episodes were reported. In total 19 vectors were constructed that summarized the hour before a panic attack occurred (pre-panic) and 280 vectors were constructed that summarized normal (non-panic) hours.

Physical Activity Analysis

Given that we are attempting to infer psychological state information from physiological data, it is important that physical activity be accounted for to ensure that any changes observed in the signal are not merely due to motion artifacts. In order to account for physical activity we constructed individual regression models for each subject from the data collected when they had not experienced any panic episodes. Activity vector magnitude units (vmu) were used as regressors to predict physiological response values (HR, BR, HRV, Temp). These regression models predicted the value that would be expected, given that a subject was exerting a certain amount

of physical activity and was not experiencing panic. Given predictions from the personalized regression models, physiological data measurements were adjusted to represent the difference between the expected value of the physiological signal and the actual physiological value observed. Summary statistics that made up the feature vectors (described previously) were computed based on this difference.

Supervised Anomaly Detection

Finally, the labeled feature vectors were run through a supervised learning algorithm and classified as either pre-panic or non-panic. Personalized models were constructed for each individual. There were many more instances of negative class labels (i.e. non-panic) compared to positive class labels (i.e. pre-panic). This meant that the collected dataset was skewed towards negative instances. As such, we chose to perform anomaly detection, as opposed to standard classification, as anomaly detection is suited to handle skewed classes [14]. Separate datasets were constructed for each subject and a Gaussian probability density distribution was estimated, given a collection of *normal* (non-panic) training examples. After a Gaussian had been fit, any instance that produced sufficiently low probability values (i.e. less than a fitted threshold parameter, ϵ) was considered to be an outlier – outside the *normal* distribution.

RESULTS

Given that there was a limited amount of data available, we chose an analysis based on an adjusted leave- k -out cross-validation method [1]. This allowed all positive (pre-panic) instances a chance to be evaluated by including them in a held-out test set, while still training models on separate training/validation data. Training, validation and held-out test sets were constructed as follows:

Training set: 80% of the non-panic instances were used to construct a *normal* model.

Validation set: In cases where panic episodes were reported, 50% of the pre-panic instances were included in the validation set. The training and validation set were combined to perform threshold (ϵ) selection.

Held out test set: The remaining 20% of non-panic instances and 50% of pre-panic instances were held out and used to evaluate the learned model.

Where it was possible, two separate runs were performed in which the 50% of pre-panic instances were swapped between the validation and test set. Individual personalized models were constructed for each subject.

The total sum of classifications made by the two runs are reported in Table 1 for each subject that recorded physiological data ($n=7$). All panic episodes that consisted of at least 15 minutes of physiological data preceding the report were included. This requirement led to two of the panic episodes being removed from the analysis. Subjects 1 – 4 reported panic episodes while wearing the device, whereas Subjects 5 – 7 collected physiological data, but did not report any usable panic episodes. True positives, true negatives, false positives and false negatives recorded from the held-out test-set

	True Positives	True Negatives	False Positives	False Negatives		Precision	Recall	F_1 Score
Subject 1	1	13	0	0		1.0	1.0	1.0
Subject 2	3	8	0	1		1.0	0.75	0.86
Subject 3	2	2	0	0		1.0	1.0	1.0
Subject 4	6	16	2	4		0.75	0.6	0.67
Subject 5	0	10	2	0	Micro-Average	0.857	0.706	0.774
Subject 6	0	28	0	0	Macro-Average	0.938	0.838	0.885
Subject 7	0	22	0	0		–	–	–
Total	12	99	4	5		–	–	–

Table 1. Results of panic episode prediction on held-out test-set for all subjects

	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	Subject 7	Total
Non-Panic Instances	65	20	5	42	28	69	51	280
Pre-Panic Instances	1	4	2	10	0	0	0	17
Total	66	24	7	52	28	69	51	297

Table 2. Total pre-panic and non-panic instances per subject used to create training, validation and test sets.

are listed for each subject. Also displayed, are the precision, recall and F_1 score for subjects who reported panic episodes (i.e. Subjects 1 – 4).

For Subject 1 and Subject 3, there were no false positives or false negatives. However, the number of instances in the test-set was limited, so care must be taken in interpreting these results. Subject 2 recorded no false positives, but misclassified 1 pre-panic instance out of 4 positive instances, giving a precision of 100% and recall of 75%. Subject 4 recorded the largest number of panic episodes. Out of 10 reported episodes, 6 were classified correctly and 4 were misclassified as non-panic, giving a precision and recall of 75% and 60%, respectively. For Subjects 5 – 7 all, but 2, instances in the test-set were correctly classified as non-panic episodes.

Table 1 also displays the micro and macro averages [8] for Subjects 1 – 4¹. The micro-average is calculated by first summing the true positives, true negatives, false positives and false negatives and then computing the precision, recall and F_1 score on the sum of values. The macro-average is the ordinary mean of the individual precision, recall and F_1 values. According to [8] micro-averaging allows large classes to dominate over small classes, whereas macro-averaging gives equal weight to each class. Therefore, [8] advocates macro-averaging be used to get a sense of effectiveness on small classes. Table 2 lists the total number of pre-panic and non-panic instances that were used to train each model.

DISCUSSION

By examining the hour of physiological data preceding a reported panic episode it could be seen that heart rate, respiration rate and core temperature were elevated and heart rate variability was lower than expected, compared with non-panic periods. In general, this led to prediction models that produced a low number of false positives and false negatives. For one subject (Subject 4), larger numbers of false negatives were recorded due to not being able to entirely separate pre-panic and non-panic episodes using the given inputs of HR, BR, HRV and Temp. Collecting further data from

¹Including these results for Subjects 5 – 7 did not make sense as there is no possibility of true positives.

additional subjects will allow us to determine if this is a common trend, or specific to certain individuals. The trade-off between reducing false-positives (notification overload) versus false negatives (missed interventions) is also an interesting question for future research.

In previous work [9, 15], it was shown that physiological measurements exhibited multiple significant change points in the hour preceding a reported panic attack. While change points were witnessed in the feature vectors used within our analysis, the main difference observed had to do with the magnitude of the recorded measurements between pre-panic and non-panic intervals. In their analysis, [9] aggregated the measurements from all individuals in their study, whereas in our work we constructed separate models for each individual. This may explain why the existence of distinguishing change points was not as pronounced between pre-panic and non-panic time periods in our results.

Limitations

The main limitation of the presented work is the scarce amount of data available per subject to train personalized prediction models. In particular, for some subjects only one or two panic episodes were available for evaluation. While the amount of data was limited, the results from this initial feasibility study show the potential of performing panic attack prediction from physiological data. However, further data will need to be collected to have more confidence in the results.

CONCLUSION AND FUTURE WORK

In this work, we introduced the idea of a mobile and wearable system for predicting and responding to upcoming panic attacks and presented results from an initial detection feasibility study. Although the dataset collected from this initial study was small, the results from our proof-of-concept analysis highlight the potential of such a system. For future work, we intend to collect further data in order to fully evaluate the efficacy of model predictions. Further effort is also required to allow our prediction algorithm to run continuously in real-time. Finally, we wish to evaluate the efficacy of the complete system by running intervention studies on further populations.

REFERENCES

1. Sylvain Arlot and Alain Celisse. 2010. A survey of cross-validation procedures for model selection. *Statistics Surveys* 4 (2010), 40–79.
2. American Psychiatric Association. 1994. *Diagnostic and statistical manual of mental disorders : DSM-IV* (4th ed. ed.). American Psychiatric Association Washington, DC. xxvii, 886 p. ; pages.
3. Zephyr Technology Corporation. 2015. (2015). <http://www.zephyranywhere.com>.
4. Raymond R Goetz, Donald F Klein, Robert Gully, Jeffrey Kahn, Michael R Liebowitz, Abby J Fyer, and Jack M Gorman. 1993. Panic attacks during placebo procedures in the laboratory: physiology and symptomatology. *Archives of general psychiatry* 50, 4 (1993), 280–285.
5. GA Hibbert. 1984. Hyperventilation as a cause of panic attacks. *British Medical Journal* 288, 6413 (1984), 263–264.
6. Malcolm Lader and Andrew Mathews. 1970. Physiological changes during spontaneous panic attacks. *Journal of Psychosomatic Research* 14, 4 (1970), 377–382.
7. Michael R Liebowitz, Abby J Fyer, Jack M Gorman, Donald Dillon, Ilana L Appleby, Gail Levy, Samuel Anderson, Morton Levitt, Michael Palij, Sharon O Davies, and others. 1984. Lactate provocation of panic attacks: I. Clinical and behavioral findings. *Archives of General Psychiatry* 41, 8 (1984), 764–770.
8. Christopher D Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Vol. 1. Cambridge university press Cambridge.
9. Alicia E. Meuret, David Rosenfield, Frank H. Wilhelm, Enlu Zhou, Ansgar Conrad, Thomas Ritz, and Walton T. Roth. 2011. Do unexpected panic attacks occur spontaneously? *Biological psychiatry* 70, 10 (2011), 985–991.
10. Alicia E Meuret, Frank H Wilhelm, and Walton T Roth. 2001. Respiratory biofeedback-assisted therapy in panic disorder. *Behavior Modification* 25, 4 (2001), 584–605.
11. Antonio E Nardi, Rafael C Freire, and Walter A Zin. 2009. Panic disorder and control of breathing. *Respiratory physiology & neurobiology* 167, 1 (2009), 133–143.
12. National Institute of Mental Health. 2015. Panic Disorder. (2015). <http://www.nimh.nih.gov/health/topics/panic-disorder/index.shtml>.
13. E. S. Page. 1957. On problems in which a change in a parameter occurs at an unknown point. *Biometrika* (1957), 248–252.
14. William K. Robertson, Federico Maggi, Christopher Kruegel, and Giovanni Vigna. 2010. Effective Anomaly Detection with Scarce Training Data. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2010, San Diego, California, USA, 28th February - 3rd March 2010*. <http://www.isoc.org/isoc/conferences/ndss/10/pdf/07.pdf>
15. David Rosenfield, Enlu Zhou, Frank H. Wilhelm, Ansgar Conrad, Walton T. Roth, and Alicia E. Meuret. 2010. Change point analysis for longitudinal physiological data: detection of cardio-respiratory changes preceding panic attacks. *Biological psychology* 84, 1 (2010), 112–120.
16. M Katherine Shear, Timothy A Brown, David H Barlow, Roy Money, Diane E Sholomskas, Scott W Woods, Jack M Gorman, and Laszlo A Papp. 1997. Multicenter collaborative panic disorder severity scale. *American Journal of Psychiatry* 154, 11 (1997), 1571–1575.
17. Zephyr Technology. 2013. Zephyr BioHarness 3 Data Sheet. (2013).
18. Dominic JC Wilkinson, Jane M Thompson, Gavin W Lambert, Garry L Jennings, Rosemary G Schwarz, Don Jefferys, Andrea G Turner, and Murray D Esler. 1998. Sympathetic activity in patients with panic disorder at rest, under laboratory mental stress, and during panic attacks. *Archives of General Psychiatry* 55, 6 (1998), 511–520.