Take a breath: A product aimed to help calm down people suffering from social anxiety in crowded environments

R. van Overbeek	T.J. Veelenturf	G.A.M.T. Walenberg
r.v.overbeek@student.tue.nl	t.j.veelenturf@student.tue.nl	g.a.m.t.walenberg@student.tue.nl
1460641	1426362	1455044

Group 10

ABSTRACT

Many people suffer from social anxiety in public spaces. It is well established that stress has a negative impact on daily life. This report aims to relate the presence of social anxiety to the amount of background noise and the presence and severity of an anxious leg bounce. Subjects were asked to wear an accelerometer while a microphone records environmental noise. By recording the background noise using the microphone and detecting leg movement, a dataset was created where four different classes are identified. Using a learning algorithm, the state of anxiety was classified. Several different algorithms were tried and an LSVC was determined to be optimal. The resulting in-sample accuracy was high enough to predict the presence of anxiety to a high precision. However, the out of sample accuracy was too low. These results suggest that to accurately predict the presence of anxiety, the model should be adapted to each individual user.

INTRODUCTION

On average around 12.5 percent of Americans experience a phobia at some point in their life and in 2009, 7 percent or 15 million American adults were affected by social anxiety disorder every year [1, 9, 12]. These are high numbers even before Covid-19, which made the numbers increase [11]. Such a social anxiety disorder negatively impacts daily life, such as avoiding crowded places or interaction with new people [7, 8] thus have a negative impact on the quality of life [10]. It can also cause bodily symptoms such as the 'anxious leg bounce' [2] . This leg bouncing is very often related to anxiety among other things [4].

A solution is to identify when a person is anxious by measuring the 'anxious leg bounce' and then engage the person in breath exercises as a way of anxiety management [3]. Breathing exercises can help control the breathing and decrease the intensity of the anxious feelings [6, 13, 14]. A machine learning algorithms could offer this solution by learning to classify and recognize this bodily symptom. For this challenge of the course Intelligent Interactive Products, provided by the Eindhoven University of Technology, an intelligent interactive product is made.

A short video of the product and concept can be viewed through the following link: https://youtu.be/mVH029GIou0

TAKE A BREATH PRODUCT

Intelligence

This interactive product needs to identify the context of the situation (how crowded it is) and assess, based on the intensity of the leg bounce, if the person is very anxious, mildly anxious or not anxious. The way of leg bouncing and how the background noise sounds differ per person and situation. Therefore, it is not programmable and requires machine learning. Take a Breath recognized not only leg bouncing intensity but also background crowd noise and fuses the sensor streams to predict one of the classes. Here a multi classification is used with a frequency domain.

Here we defined a learning problem Y=g(x). Where y is the predicted outcome [no anxiety, mildly anxious, highly anxious] (see table 1) and g( ) is the classifier. As the algorithm uses frequency domain, X consists of the maximum frequency intensity values of the axes of the accelerometer (x, y, z) calculated with a FFT (Fast-Fourier Transform) over a stated timeframe. In addition, X also consists of lots of microphone data where x=P*i, where 1 < i < 64. I means the number of bins, which is in this case limited to 64.

Leg	Sound	Output	Prediction
		Mildly anxious	A
		No anxiety	B
		Mildly anxious	C
		Highly anxious	D

Legend
Not present
Mild
Intensive

Table 1: Classification Prediction Table

Tangibility and interactivity

Take a breath measures the anxiety through an accelerometer placed on the leg next to the knee (see figure 1). In addition, a microphone is placed to measure the context (the situation). The microphone is made in the form of a simple, minimalistic box with LEDs (see figure 2). The data will go through the learning algorithm using edited Arduino and Processing code, facilitated by the course [5]. If the user is not anxious the lights will not do anything. If it measures that the user is a little anxious. The lights will mimic a simple minimal breathing exercise through dimming on and off. If the learning algorithm notices that the user is very anxious (lots of noise and lots of bouncing with the leg) it will increase light and try to get the users attention to do the breathing exercise (see figures 3). Then the user can breathe along with the box and control their breathing (see figure 4).

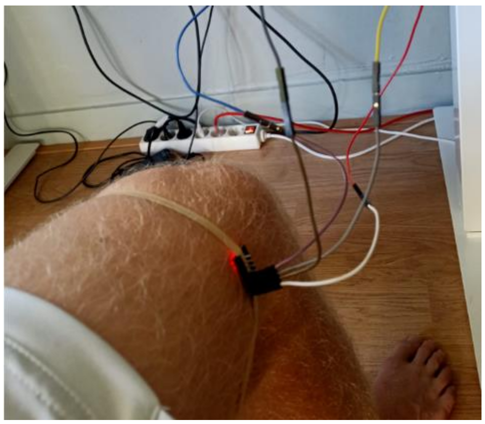

Figure 2: Accelerometer on the right knee

Figure 2: Take a Breath Box

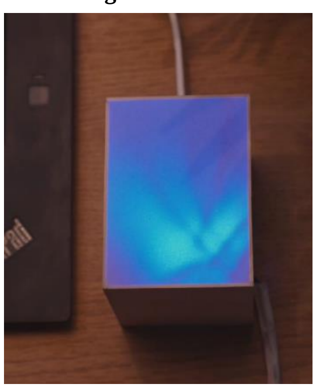

Figure 2: Led breathing exercise

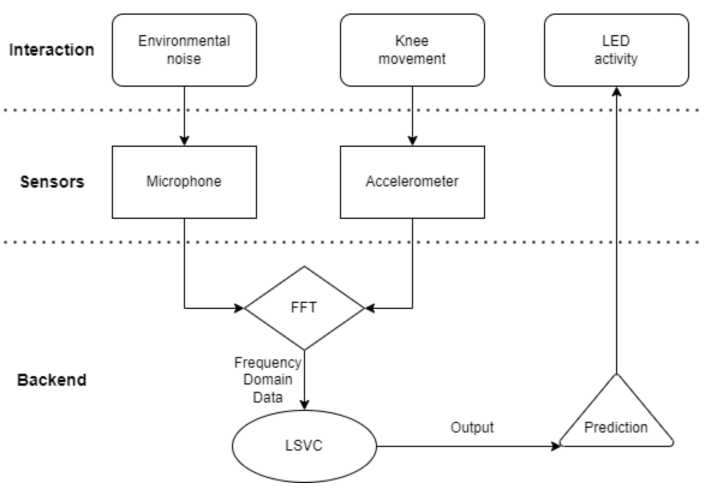

Figure 4: Interaction Flow Chart

REALIZATION

Hardware

The hardware of the product consists of 4 main components. A microcontroller, an accelerometer, a microphone and RGB LEDs. The accelerometer/gyroscope can detect motions and its position. In the case of this product, the MPU6050 was used to detect the intensity of vibrations of the leg. The microphone was used to detect sounds and the RGB LEDs were used to create light in a breathing pattern. In figure 5 a diagram of the circuit is shown.

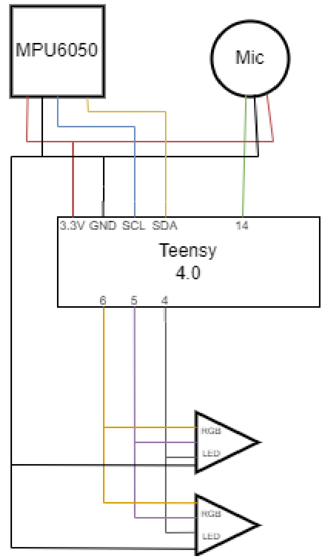

Figure 5: Electrical Circuit

The components were put onto a breadboard (see figure 6), which in turn has a casing around it to protect the wires and main components. The Teensy 4.0 microcontroller was connected to the laptop to create a serial communication with Arduino and Processing. The accelerometers Vin, GND, SCL and SDA were connected to the 3.3V, GND, SCL and SDA pin on the Teensy respectively. The microphones VCC, GND and OUT pin were connected to the 3.3V, GND and pin 14 respectively. The two RGB LEDs were connected to pins 4, 5 and 6 and the GND.

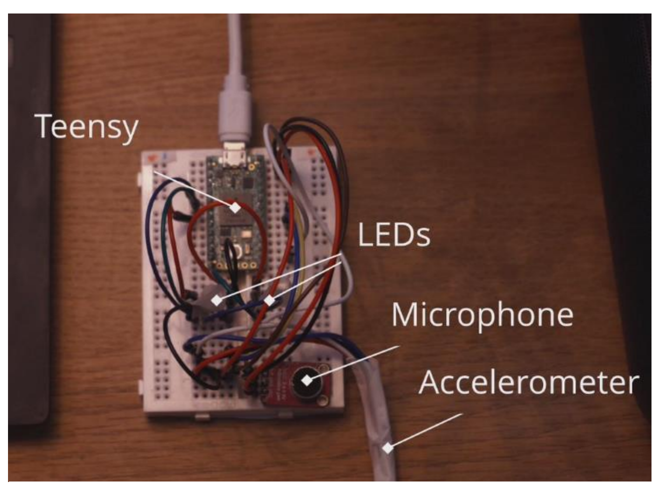

Figure 6: Electronical wiring

The communication protocol of the hardware with the laptop is as follows: The data of the accelerometer and microphone is being captured and is sent (via the serial port) to Arduino and then to Processing. The Learning Algorithm, which will be explained later, goes to work and makes a certain prediction. The prediction is sent back via Processing to Arduino and via the serial port the LEDs will act accordingly.

Data acquisition & Feature extraction

Both the microphone and the accelerometer were sampled at a rate of 500Hz. The output volume of the microphone module is calibrated using the gain adjustment on the back of the microphone to prevent clipping. The analog-to-digital conversion of the microphone signal is done at the Teensy using its built in ADC. The resulting discrete signal has a bit depth of 10. The accelerometer values are communicated digitally as a 16-bit integers, resulting in a value between -32,768 and +32,768 for each of the three axes. All the data is then remapped to a range between 0 and 500 before it is communicated serially with the processing program. There, each second, the received microphone data is separated into 64 bins. Both the audio and accelerometer data are then converted to frequency domain using a fast Fourier transform. The resulting data contains information on the pitch and volume of the surrounding audio and both the speed and amplitude of the leg bounce. This data can be more easily interpreted by the model since the important attributes are more readily visible.

Learning Algorithm

For choosing the best Learning Algorithm, training data was used to optimize and test three models. A kernel support vector classifier with RBF, a linear support vector classifier and a K-nearest neighbor classifier. Table 2 shows for each model the best accuracy and corresponding parameters. These parameters were calculated by filling them in for every different model. Choosing the parameter values was based on the lowest possible C-score obtained during a C-search operation, in combination with the highest accuracy for LSCV and KSCV. For KNN, the 'K' was chosen based on the highest accuracy.

Using a paired T-test showed that between the datasets it cannot be concluded that a significant difference exists. It was therefore decided to use the linear support vector classifier with C=32 and k=5 as it is less prone to overfitting if it is compared to a KNN. It is also computationally less expensive than a Kernel SVC.

	In sample accuracy train	In sample accuracy test	out of sample accuracy
LSVC (C=32)	89.25	86.5	54.5
KNN(K=3)	90.75	85.25	56
Kernel SVC (C=8, y = 1)	90	86.75	54.5

Table 2: Accuracy of different models

```
Results
======

Correctly Classified Instances         346               86.5    %
Incorrectly Classified Instances        54               13.5    %
Kappa statistic                          0.82
Mean absolute error                      0.2617
Root mean squared error                  0.3298
Relative absolute error                 69.7778 %
Root relative squared error             76.1699 %
Total Number of Instances              400

=== Confusion Matrix ===

   a   b   c   d   <-- classified as
  66   0  29   5 |   a = A
   0 100   0   0 |   b = B
  17   1  82   0 |   c = C
   2   0   0  98 |   d = D

=== Detailed Accuracy By Class ===

                 TP Rate  FP Rate  Precision  Recall   F-Measure  MCC      ROC Area  PRC Area  Class
                 0,660    0,063    0,776      0,660    0,714      0,632    0,858     0,629     A
                 1,000    0,003    0,990      1,000    0,995      0,993    0,998     0,990     B
                 0,820    0,097    0,739      0,820    0,777      0,700    0,895     0,668     C
                 0,980    0,017    0,951      0,980    0,966      0,954    0,989     0,945     D
Weighted Avg.    0,865    0,045    0,864      0,865    0,863      0,820    0,935     0,808
```

Figure 7: Results Evaluation Take a Breath with C=32.0 and k=5

EVALUATION

For evaluating the classifier it was decided to use two methods. With the first one being that two team members made train and test data and did an in sample cross validation. This means that the evaluation of the test data is done with the combined training data of the same two persons. The second method was to use training data of two team members and making a test set from the third team member. This was the out of sample method. It can be noted that the in-sample test data, with same members that trained the model, had a much better accuracy as can also be seen in table 2. The out of sample accuracy with using the test data of another team member was not very high. The program was tested using a 5-fold cross validation with a C of 32. This gave an overall accuracy of 86.5% as can be seen in figure 7. This also shows the confusion Matrix together with other evaluation metrics. Here it is clear that the classifications of C and A are not very accurate. This might be explained by the fact that both A and C are mild vibrations and can therefore easily be misclassified. The other evaluation metrics like the TP rate is above the 0.25 (0.865) so it is much better than random. Furthermore, our FP rate is 0.045. Although this is a quite high value, this is because it confuses A and C as already described earlier. As A and C have the same output as shown in table 1, it should not pose a problem for the interaction of the product. In addition, when looking at the ROC-area, its weighted average is 0.935, which is close to 1. This indicates that this is a good classifier.

However, the system would not obtain the skill that well when the product would be given to a user who would immediately start using it, as can be seen in table 2.

CURRENT LIMITATIONS

One major limitation of the study and the product is that the user has to train the data before using it to get the highest accuracy, if the product is to be used without any pretrained data. This could be overcome by adding a lot more training data samples from multiple people, which the team did not have. Furthermore, the long wires of the accelerometer are connected to the teensy so the user would not be able to stand up or do other wild movements with the leg. This way the user can only sit straight up or a bit slumped and would only be able to put the leg in front of them like in a normal chair. Another limitation is that the product will only act on leg movement. For social anxiety, there are a lot more symptoms, while the product only focuses on one. Also, the only interaction with the user is through light in the little box. For a more pleasant, familiar or robust experience, the interaction can be made more extensive and/or rich.

FUTURE WORK

In the future, the user would be able to wear a device that can detect the level of social anxiousness. The first step for this is that the user should be able to move free in the environment. The product can be enhanced with a Bluetooth or Wi-Fi module so there are no wires attached to the body. This allows for more movements and other ways of recognizing social anxiety. In combination with other sensors that measure other symptoms of social anxiety, the product could be a great addition to the already existing forms of dealing with the problem.

Moreover, the product would be more robust when it is trained by more people. This means that it needs more training data, which is more data points, but also a wider variety in people.

Another thing for in the future is that the data should be made more robust. This can be done by applying some form of normalization to the data to make it more accurate for other users. This way, users won’t have to train data themselves, but just use a ready-made product.

CONCLUSION

The purpose of Take a Breath is to detect social anxiety in people who are working in a crowded context by recognizing the surrounding sound and leg bouncing. Using a LSCV turned out to be most optimal for classifying the different states of states of anxiety. It performed similarly to a KNN or Kernel SVC while retaining a relatively low likeliness for overfitting and low computational cost. Using this Learning algorithm, we were able to achieve a high in sample accuracy at 86.5%, but the out of sample accuracy was only 54.5% which is too low for practical application. This means that when a new user presents themselves to the product, the product needs to be retrained to ensure usable levels of accuracy.

REFERENCES

[1] ADAA. 2009. Social Anxiety Disorder. Anxiety and Depression Association of America, ADAA. Anxiety & Depression Association of America. Retrieved from: https://adaa.org/understanding-anxiety/social-anxiety-disorder

[2] Ellen Scott. 2022. The anxious leg bounce: Why it happens and how to deal with it. Metro. Retrieved from https://metro.co.uk/2019/04/20/anxious-leg-bounce-happens-deal-9277508/

[3] Gillian Butler, Anne Cullington, Mary Munby, Peter Amies, and Michael Gelder. 1984. Exposure and anxiety management in the treatment of social phobia. *Journal of Consulting and Clinical Psychology, 52*(4), 642–650. DOI: https://doi.org/10.1037/0022-006x.52.4.642

[4] Hashini Senaratne, Kisrten Ellis, Sharon Oviatt, and Glenn Melvin. 2020. Detecting and differentiating leg bouncing behaviour from everyday movements using tri-axial accelerometer data. In Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers (UbiComp-ISWC '20). Association for Computing Machinery, New York, NY, USA, 127–130. DOI: https://doi.org/10.1145/3410530.3414388

[5] Liang, R. H. (2021). GitHub - howieliang/IIP2021: Example Codes: 2021 DBB220 TU/e Interactive Intelligent Products. [Computer software] GitHub. Retrieved from https://github.com/howieliang/IIP2021

[6] Martin A. Katzman, Patricia L. Gerbarg,, Christina Iorio, Catherine Cameron, Monica Vermani, Richard P. Brown, Michele Davis, and Dina Tsirgielis. 2012. A multicomponent yoga-based, breath intervention program as an adjunctive treatment in patients suffering from generalized anxiety disorder with or without comorbidities. *International Journal of Yoga*, *5*(1), 57–65. DOI: https://doi.org/10.4103/0973-6131.91716

[7] Mental Health America. (n.d.). Social Anxiety Disorder. Retrieved from https://www.mhanational.org/conditions/social-anxiety-disorder#:%7E:text=Fifteen%20million%2C%20or%20seven%20percent,Social%20Anxiety%20Disorder%20%5B1%20%5D.

[8] National Collaborating Centre for Mental Health (UK). 2013. *Social Anxiety Disorder: Recognition, Assessment and Treatment*. Leicester (UK): British Psychological Society (UK); 2013. (NICE Clinical Guidelines, No. 159.) 2, SOCIAL ANXIETY DISORDER. Retrieved from https://www.ncbi.nlm.nih.gov/books/NBK327674/

[9] National Institute of Mental Health. 2017. Specific Phobia. National Institute of Mental Health (NIMH). Retrieved from https://www.nimh.nih.gov/health/statistics/specific-phobia

[10] Ronald C. Kessler. 2003. The impairments caused by social phobia in the general population: implications for intervention. *Acta Psychiatrica Scandinavica*, 108, 19–27. DOI: https://doi.org/10.1034/j.1600-0447.108.s417.2.x

[11] Statista. 2021. Mental health changes due to the coronavirus in the Netherlands 2020. Retrieved from https://www.statista.com/statistics/1115803/dutch-mental-health-coronavirus/

[12] Thomas Christiansen. 2022. Social Anxiety Disorder. The Recovery Village Drug and Alcohol Rehab. Retrieved from https://www.therecoveryvillage.com/mental-health/social-anxiety-disorder/

[13] Traci Pedersen. 2022. All About Enochlophobia (Fear of Crowds). Psych Central. Retrieved from https://psychcentral.com/anxiety/enochlophobia#treatment

[14] Xiao Ma, Zi-Qi Yue, Zhu-Qing Gong, Hong Zhang, Nai-Yue Duan, Yu-Tong Shi, Gao-Xia Wei, and You-Fa Li. 2017. The Effect of Diaphragmatic Breathing on Attention, Negative Affect and Stress in Healthy Adults. *Frontiers in Psychology, 8.* DOI: https://doi.org/10.3389/fpsyg.2017.00874