Biomedical Signal Processing and Control 59 (2020) 101929

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

journal homepage: www.elsevier.com/locate/bspc

AI-enabled remote and objective quantification of stress at scale

Abdulrhman H. Al-Jebrni[a,b], Brendan Chwyl[c], Xiao Yu Wang[c], Alexander Wong[c,d], Bechara J. Saab[b,*]

[a] Software School, Shanghai Jiao Tong University, 800 Dongchaun Rd., 200241, Minhang District, Shanghai, China
[b] Mobio Interactive Inc., BioMedical Zone, St. Michael's Hospital, M5B 1T8, Toronto, ON, Canada
[c] Vision and Image Processing (VIP) Research Group, Department of Systems Design Engineering, University of Waterloo, 200 University Ave W, N2L 3G1, Waterloo, ON, Canada
[d] Department of Mechanical and Mechatronics Engineering & David R. Cheriton School of Computer Science, University of Waterloo, 200 University Ave W, N2L 3G1, Waterloo, ON, Canada

ARTICLE INFO

Article history:
Received 23 March 2019
Received in revised form 8 January 2020
Accepted 2 March 2020

Keywords:
Artificial intelligence
Digital therapeutics
Photophlethysmography
Heart rate variability
Stress

ABSTRACT

Background: Accurate measurement of human stress at scale is a major mHealth challenge. Here we explore the potential for deep neural networks (DNNs) to improve remote and objective quantification of stress from voluntary selfie videos captured through mobile device front-facing cameras.
Methods: Two DNNs were trained with heart rate (HR) and heart rate variability (HRV) data obtained through photophlethysmographic imaging (PPGI) of 11,823 mobile device selfie videos captured in tandem with self-assessments of stress, and compared to contemporary algorithms used to estimate stress from HR and HRV data.
Results: A classification DNN and predictive DNN determined self-reported stress with 86 % accuracy and a mean absolute error of 0.001, respectively. Both DNNs performed far better than other recently described approaches when applied to the identical dataset.
Conclusions: Well-trained DNNs can objectively and remotely quantify stress at scale. Future efforts may concentrate on the measurement of additional enigmatic cognitive states.

1. Introduction

Peripheral physiology often reflects cerebral activity. One clear example of this relationship is the autonomic shift in heart rate variability (HRV) paralleling shifts in cognitive stress [1]. For this reason, HRV has become an experimental proxy for stress assessment in clinical research [2] and amongst the general population through mHealth tools [3]. One major limitation in the collection of HRV via mobile apps is a dependence on "wearables" or other hardware that is only accessible and desirable to a small fraction of mobile app users. There is therefore increasing interest to develop technology that is able to obtain accurate measures of HRV and stress without the need for additional hardware. Leveraging the increasingly high-quality mobile device camera is one promising approach, since computer vision and photoplethysmographic imaging (PPGI) can extract HRV by measuring volumetric changes in arterial blood flow within the finger or face [4,5]. The persisting challenge however, is making use of the relatively low quality mobile device PPGI to estimate cognitive stress. This challenge represents a special opportunity in artificial intelligence (AI) research.

In this study, deep learning was performed on heart rate (HR), HRV and self–assessment of stress data obtained through a popular meditation app from a broad spectrum of individuals in nearly 100 countries. These data were obtained via the mobile device camera as a standalone voluntary measurement, or as voluntary paired "before and after" assessments in conjunction with a guided mindfulness meditation provided through the same app. Implicit (neural network) models were then employed to predict self-assessment of stress based on the HR and HRV data. Two classes of artificial feed-forward neural networks were compared, and contrasted against current clinical algorithms typically applied to hospital-recorded electrocardiograms (ECGs), to assess their ability to provide an objective measure of cognitive stress from PPGI data obtained from mobile devices.

Current clinical algorithms used to assess stress from HRV originate from the discovery that pacemaker sites in the heart are influenced by the autonomic nervous system (ANS). In particular, the heart's sinoatrial node is innervated by both the

* Corresponding author at: Mobio Interactive, Centre for Social Innovation, 720 Bathurst St., Toronto, ON, Canada.
E-mail address: bechara@mobiointeractive.com (B.J. Saab).

https://doi.org/10.1016/j.bspc.2020.101929

parasympathetic and sympathetic nervous systems (PNS and SNS, respectively). It is precisely because both the PNS and SNS influence HR, that HRV is correlated to cognitive stress. Specifically, the time-courses of PNS- and SNS-mediated HR regulation are distinct, generating two HRV frequency components [6], with the high-frequency HRV components being primarily mediated by the PNS. HRV therefore provides insight into the balance of PNS and SNS control over HR; and also to some extend the stress response circulating hormones, including corticotropin, acetylcholine, epinephrine and norepinephrine [7,8]. Therefore, keeping other factors constant, activation of the SNS, which prepares the body for danger via the so-called "flight-or-flight" response, should decrease the HRV high-frequency/low-frequency component ratio [6]. Indeed, exceptionally low HRV is associated with impaired ANS homeostasis [1,9], which is a serious manifestation of extreme psychological stress [10]. The interpretation of more subtle changes in HRV remains controversial, even for high-quality ECG recordings acquired over several hours or days. There is great promise nonetheless, that with sufficiently advanced algorithms, HRV data may be successfully leveraged to obtain an accurate reading of cognitive stress.

Reliable methods to assess psychological stress from HRV are of high clinical and societal importance due to the critical role played by stress in overall health, especially if these methods leverage the ubiquity of mobile devices. To this end, we have applied the latest advancements in AI to train deep neural networks (DNNs) using real world data (RWD) obtained from thousands of mobile device users throughout the globe.

DNNs are mathematical models that simulate basic features of neurocortical communication, and are considered distinctive acyclic computational graphs [11] containing three basic parts; i) an input layer, ii) several hidden layers and iii) an output layer. Each hidden layer has several units called neurons (N). Each neuron produces a chain of information comprised of real value activation (a) that are computed via connections to the neurons of subsequent layers through weighted (W) connections and biases (b). Through training DNNs, the weights between neurons (and to some extent biases) are continually tuned until the DNN produces the desired output.

Recently, implicit DNN models have been used to predict stress levels from HRV, and provide better accuracy compared with explicit models employing support vector machine (SVM) or logistic regression [12–15]. However, the only implicit model calculating stress from RWD that we could find, in this case a Lomb's estimator-derived algorithm, used PPGI data from a wearable that dramatically limits scalability [15]. Here we describe implicit models that classify and predict stress based on PPGI data from the mobile device camera, addressing the hypothesis that DNNs can facilitate objective and remote assessment of stress at scale.

2. Methods

2.1. Dataset and data preparation

RWD consisted of PPGI of 11,823 front-facing selfie videos voluntarily obtained with free consent (see Terms of Use) from 3,527 users of the mindfulness meditation app *Wildflowers Mindfulness*, a research predecessor to *Am Mindfulness* developed by Mobio Interactive Inc. From these anonymous data, HR and HRV ("high" and "low" frequency bands; HRV_{HIGH} and HRV_{LOW}) were estimated with previously published methods [16]. Briefly, the "SAPPHIRE" algorithm stochastically samples and tracks the cheek region of the face to construct corresponding time series of erythema transforms, and thereby a photoplethysmogram waveform that is estimated via Bayesian minimisation with the required posterior probability inferred using a Monte Carlo approach. All analysed selfie videos were simultaneously obtained with results from a "mood board" that the user used to input a selection of words that corresponded to their current mood and a "stress slider" that the user used to provide an assessment of their current level of stress (Fig. 1). PPGI data (HR, HRV_{HIGH}, HRV_{LOW}, normalised to a range of 0–1) were used in the current study as three inputs to train and test each of the neural network models to the results of the stress slider (a real number between 0 and 1). The training dataset comprised 90 % of the entire dataset, or 90 % of the entire user base (and comprising 92 % of the entire dataset), depending on the validation methods used.

2.2. Mobile app

Data were obtained from the research predecessor of *Am Mindfulness* developed by Mobio Interactive Inc., "*Am*" is a commercially available mobile wellness app that supports the development of a personal mindfulness practice through guided meditations and journaling facilities. *Am* is the first mindfulness apps to use PPGI in the objective measures of stress and the first mindfulness app to outperform an active control in a randomised controlled trial examining wellbeing [17,18].

2.3. Deep learning

We used plain fully connected DNNs, with every neuron in one layer connected to every neuron in the subsequent layer by an activation function, such as the sigmoidal activation function used in the predictive model's output layer (1):

$$a^{[l](i)} = \sigma\left(\sum_{j}^{m} W_j^{[l](i)} a_j^{[l-1]} + b_i^{[l]}\right) \tag{1}$$

Where $a^{[l](i)}$ is the output of the activation function at the layer l of the neuron i, $W_j^{[l](i)}$ is the weight of the path from the neuron on the previous layer to the neuron on the current layer, $a_j^{[l-1]}$ is the activation function output of the neuron j in the previous layer and $b_i^{[l]}$ is the bias of neuron i in the previous layer $l-1$.

2.4. DNN framework and analysis

TensorFlow 1.6.0 was selected as the framework to implement, train and deploy the deep learning models because of its high-level of flexibility [19]. All the tests and experiments were performed in Python 3.6 on an Intel CPU Core i5-4300 running Windows 10 Pro 64-bit. The architecture of the DNNs was based on empirical analysis of performance on different designs.

2.5. Classification model

The classification model featured a deep feedforward neural network that consisted of a three-neuron input layer, four 20-neuron hidden layers and a four-neuron output layer (Fig. 2A, Table 1). The input layer consisted of average HR, HRV_{LOW} and HRV_{HIGH}. The output layer classified the stress value into four categories. A hyperbolic tangent transfer function was employed as the activation function (2) throughout the hidden layers:

$$a = Tansig\,(n) = \frac{2}{1+e^{-2n}} - 1 \tag{2}$$

However, a non-saturating non-linearity activation function, the ReLU (Rectified Linear Unit) function (3), was used as the

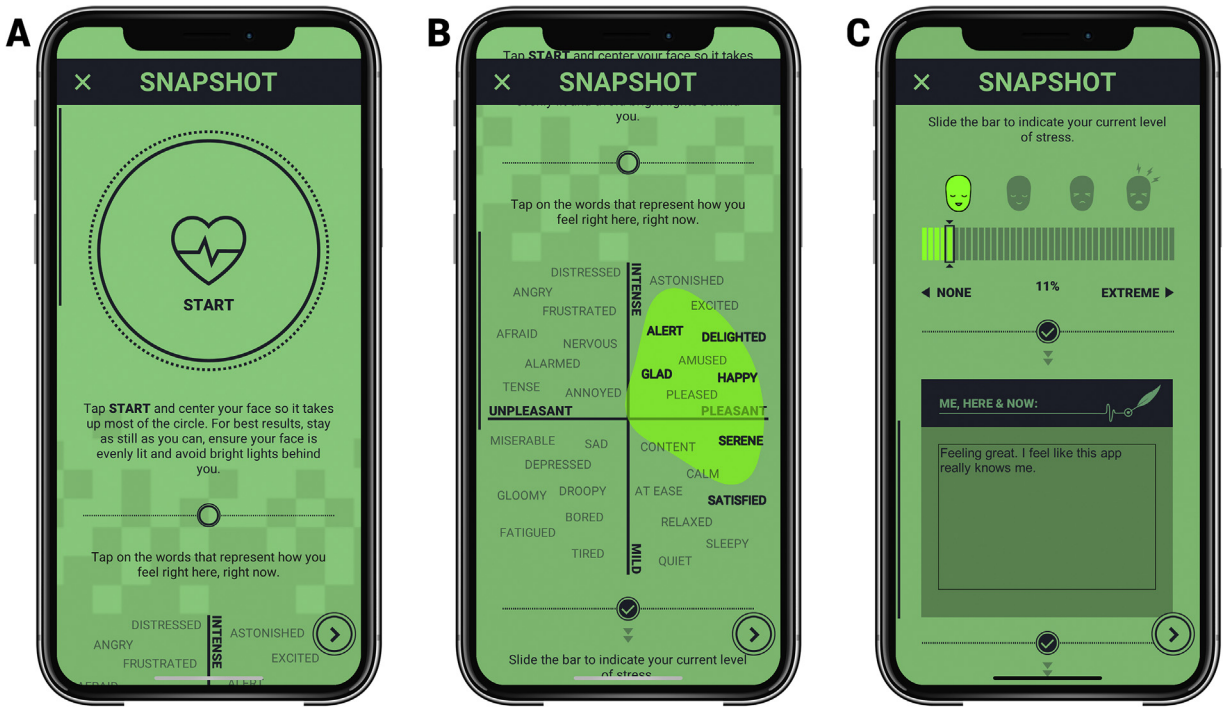

Fig. 1. PPGI and self-assessments.
(A) In-app tool employed by app users to complete selfie videos. (B) In-app tool employed by app users to convey their current mood. (C) In-app tool employed by app users to (top) indicate their level of stress via a stress-slider and (bottom) input notes reflecting on the meditation experience, their state of mind, or anything else they may wish to input and read over later. Note, in the version of the app used for our current analyses, the faces above the stress-slider were not present. The potential impact of adding these faces (each representing one quadrant used in the classification DNN) is under current investigation.

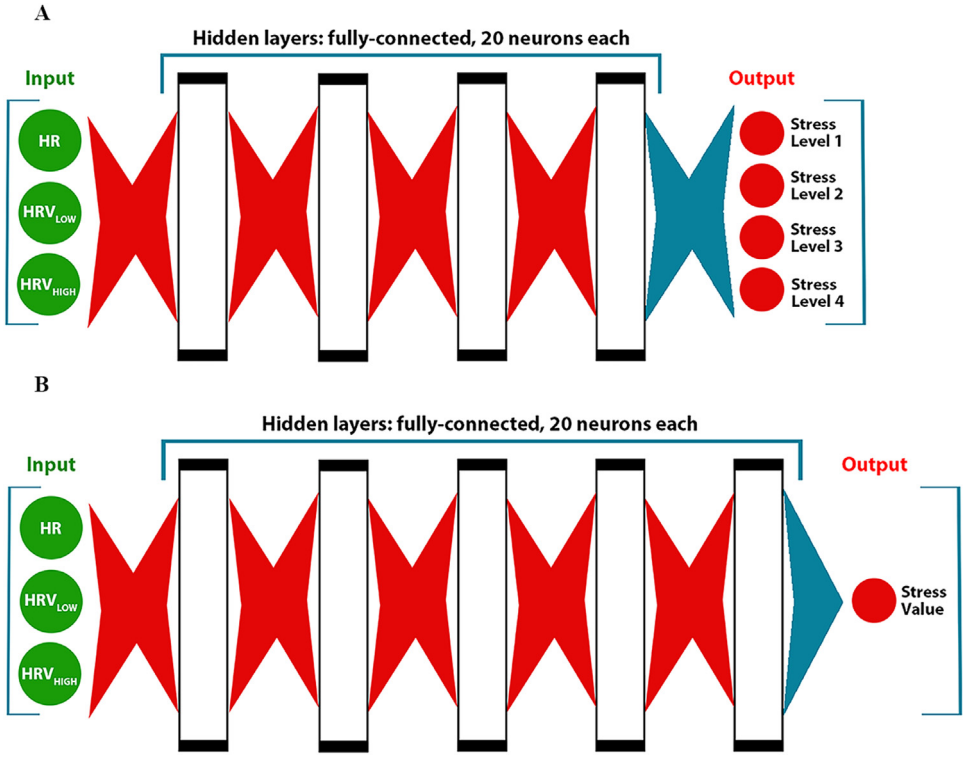

Fig. 2. Deep neural networks. (A) Diagram of classification DNN structure. (B) Diagram of predictive DNN structure.

activation function for the neurons of the output layer due to non-negativity restrictions.

$$a = relu\,(n) = \max(n,\ 0) \tag{3}$$

Where n is the input to the neurons of the output layer. The Tansig function outputted a value in the interval [−1, +1], while the ReLU outputted a nonnegative value. Cost was calculated with the SoftMax layer and cross entropy. The DNN

Table 1
DNN parameters.

Parameters	Predictive Model	Classification Model
Input features	3	3
Hidden layers	5	4
Neurons in every hidden layer	20	20
Neurons in the output layer	1	4
Hidden layers activation function	ReLU	Tansig
Output layer activation function	Sigmoid	ReLU
Learning rate	0.001	0.0001
Batch size	5	32
Number of epochs	100	100

was trained and its parameters updated with a gradient descent optimizer.

2.6. Cross-validation

We employed the holdout method of cross-validation as a model estimation technique to assess the skill of the prediction model because of its ease of implementation, computational cost-effectiveness, elimination of biasedness, conceptual simplicity and appropriateness for datasets in the range of this study [20]. In addition, we assessed the models using 10-fold cross-validation to investigate how well the models generalize.

2.7. Predictive model

The predictive model featured a deep feedforward neural network that consisted of a three-neuron input layer, five 20-neuron hidden layers and a single-neuron output layer (Fig. 2B, Table 1). Like the classification model, the predictive model's input layer consisted of HR, HRV_{HIGH} and HRV_{LOW}. The output layer contained one neuron, corresponding to the predicted stress value (a real number between 0 and 1). The ReLU function (3) was selected for the hidden layers of the predictive model primarily because of its faster training speed and smaller mean squared error (MSE) relative to Tansig. The outputs of the ReLU function through the hidden layers can grow larger than 1, while the range of the expected stress value is bounded between 0 and 1. To avoid this problem, we used a sigmoid activation function (1) in the neuron of the predictive model output layer to produce a bounded output between 0 and 1.

The predictive model DNN parameters were updated with the Adam optimizer [21], since it readily handles large volume datasets, is computationally efficient and provided a smaller MSE relative to other options, such as Gradient Descent.

2.8. Supervised learning model

A support vector machine (SVM) classifier was used to assess accuracy, using the identical inputs and outputs as the classification model DNN.

2.9. Statistical analysis

MSE, mean absolute error (MAE), accuracy, cost and the confusion matrix were applied as evaluation metrics. MSE and MAE were computed according to their standard formulae. Accuracy was calculated by dividing the number of correct predictions by all predictions. Cost was calculated using SoftMax cross entropy between self-reported stress and DNN output. The classification and predictive DNNs were not directly compared since the underlying required evaluation metrics for each are different, given that each has its own type of output and purpose. Instead, we compared each of the models against alternative machine learning algorithms that we identified from the literature. All statistical calculations, including the confusion matrix and random number generation used for chance calculations, were performed using the Numpy and TensorFlow libraries for Python 3.6 on an Intel CPU Core i5-4300 running Windows 10 Pro 64-bit.

Table 2
The probability percentage of all the stress levels of the training split.

Stress Levels	%
None	1.3
Normal	26.5
Medium	41.0
High	31.1

Table 3
Stress level range and designation.

Stress Level Value	Stress Class
1	"none"
2	"normal"
3	"medium"
4	"high"

3. Results

3.1. Real world data

The mobile app examined is available in over 140 political jurisdictions and no geographical preferences were made in data sampling. Irrespective, over three quarters of the data were obtained from English-speaking regions (Fig. 3). Similarly, while no preferences were made for self-reported age or identified gender, women aged 18–34 made up a disproportional segment (46.2 %) of the dataset. Meanwhile, over half of the data (56.7 %) were obtained from mobile devices produced by Apple and running iOS, despite two thirds of mobile devices globally running Android as the operating system. These RWD were therefore skewed as expected towards the app's user base. We also discovered that self-assessments of stress in the real world dataset are not normally distributed (Table 2), but this did not interfere with DNN training since the DNNs are model-free.

3.2. Classification model evaluation

The classification model divided stress into four levels ("none", "normal", "medium", "high"; Table 3) and was assessed by examining accuracy and cost during training, and with a confusion matrix during both training and testing. Training accuracy increased early, reached 85 % at the 23rd epoch and continued to hover around 85 % for the remaining 77 epochs until culminating at 86 %, indicating 100 epochs were sufficient to achieve maximal accuracy (Fig. 4A). The calculated cost during training, or the error in the DNN results compared to the error in the original data, demonstrated a smooth decrease throughout training from 1.25 to 0.59 within 20 epochs to a final cost of 0.50 (Fig. 4B). The training and testing confusion matrixes, which plot the fraction of correct classifications made by the DNN as a function of original and predicted stress levels, both demonstrated a concentration of correct classifications along the diagonal and were most dense for the medium stress level (Fig. 5). We noted a low density of correct predictions for the "none" stress level, which makes sense given that DNNs were provided very limited training with "none" stress level data (Table 2).

In addition, we separately examined the accuracy and cost of the classification model using the holdout cross-validation method. In this evaluation, we trained DNNs on data obtained from 90 % of the app users, and tested the DNNs on the remaining 10 % of the users. This division by user provided 92 % of the total data for training.

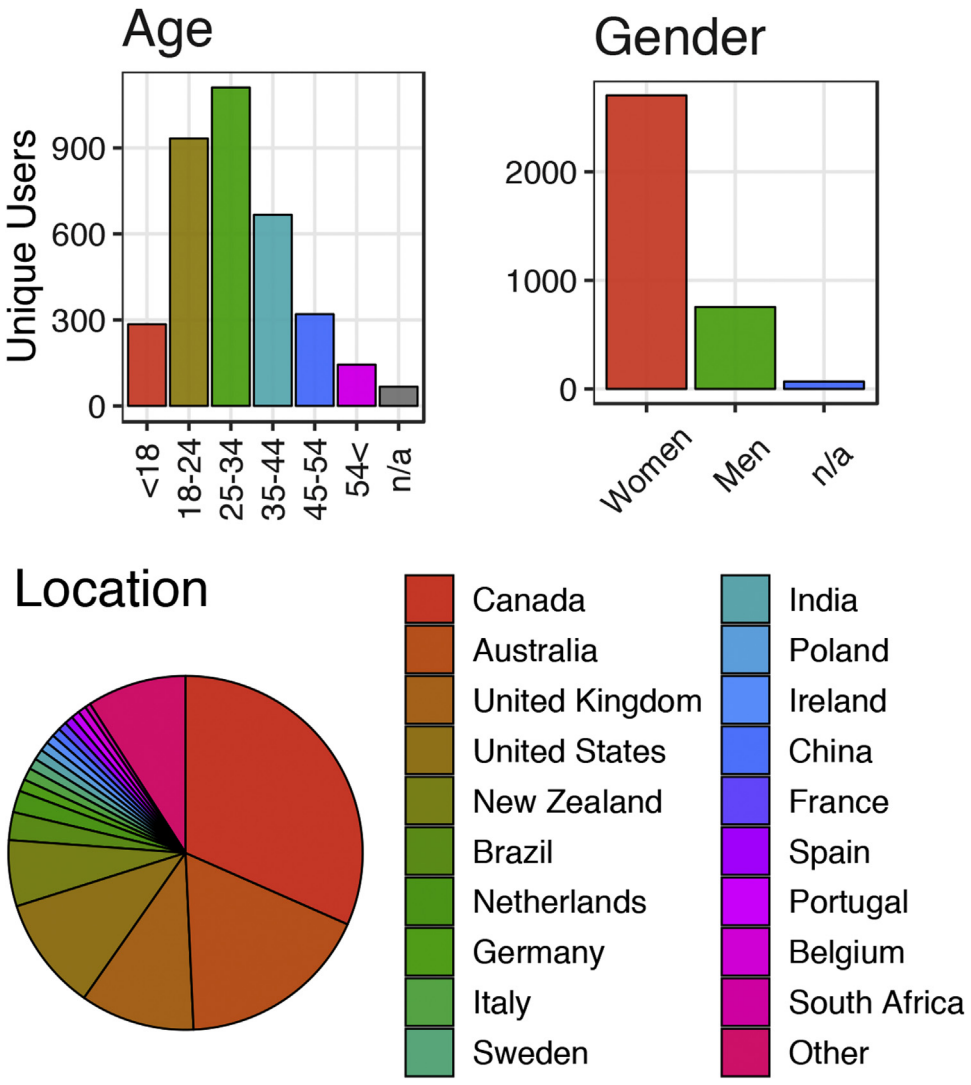

Fig. 3. Demographics of real world dataset.

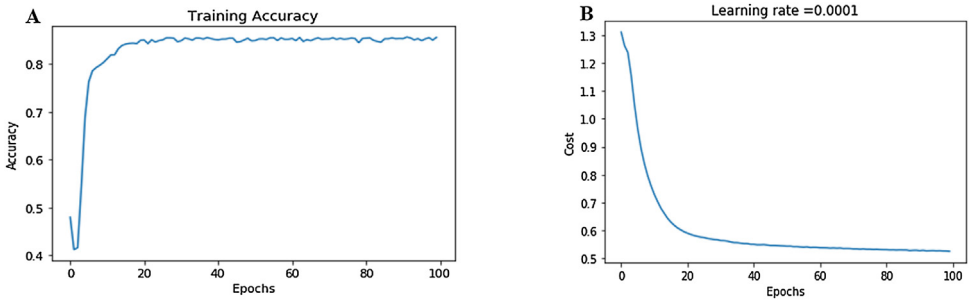

Fig. 4. Classification model evaluation. (A) Classification model accuracy over the 100 training epochs. (B) Classification model cost over the 100 training epochs.

Training the DNNs this way resulted in a slight decrease in accuracy and a slight increase in the cost rates (Table 4) in comparison to a strict user-agnostic 90:10 training:testing data split.

Finally, 10-fold cross-validation demonstrated strong alignment between accuracy and cost (Figs. 6, 7). For example, the three least accurate folds are also the folds with the highest cost rate (Fig. 6), and every drop in accuracy corresponded to a spike in cost (Fig. 7). Meanwhile, the remaining seven folds all achieved above 80 % accuracy and were almost all below a cost of 0.6. In addition, DNNs trained with 10-fold cross-validation yield lower training accuracy and higher training cost when compared to training via the holdout method (Fig. 7, Table 4).

3.3. *Predictive model evaluation*

The predictive model outputted stress as a value within the range of 0–1 and was assessed by examining the MAE and MSE. MAE dropped sharply from the initial epoch, and continued to decrease throughout training, albeit at a much slower rate following the first few epochs (Fig. 8A). MSE demonstrated a very

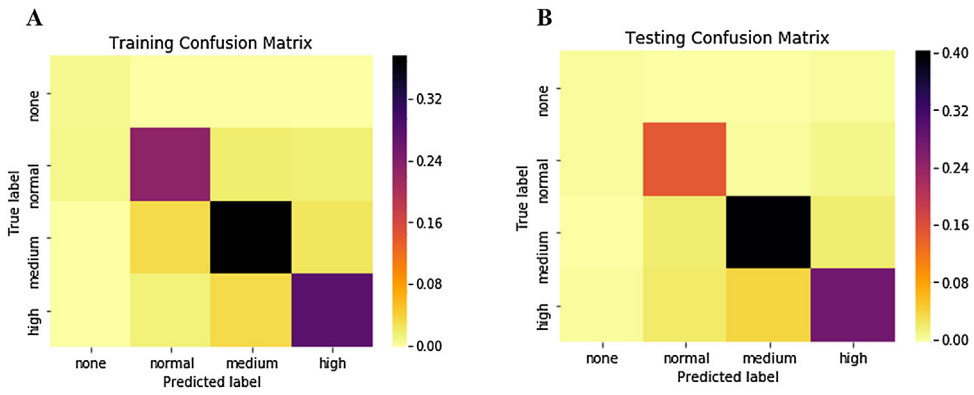

Fig. 5. Confusion matrices from the classification model.
(A) Classification model confusion matrix during training. (B) Classification model confusion matrix during testing.

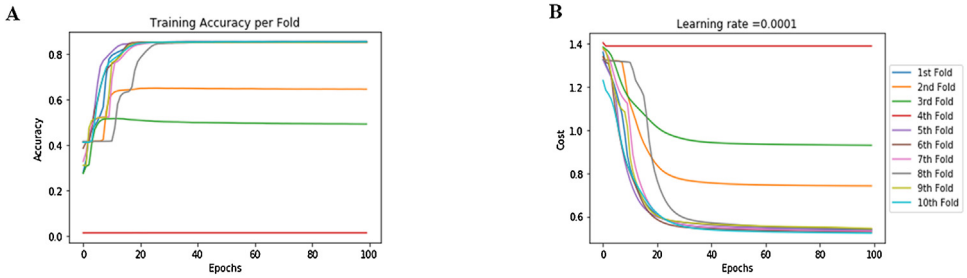

Fig. 6. Classification 10-fold cross-validation evaluation (training).
(A) Classification model accuracy over the 100 training epochs per fold. (B) Classification model cost over the 100 training epochs per fold.

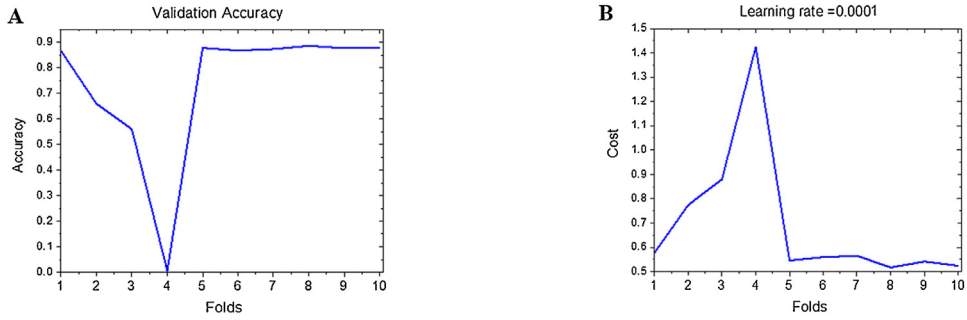

Fig. 7. Classification model 10-fold cross-validation evaluation (validation).
(A) Classification model accuracy over the 10 validation folds. (B) Classification model cost over the 10 validation folds.

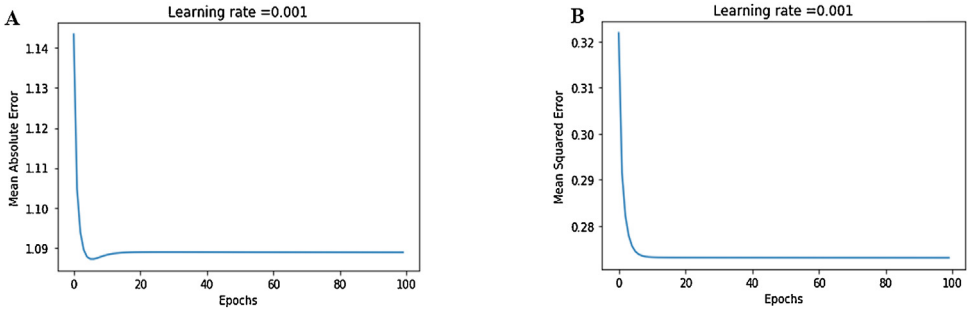

Fig. 8. Predictive model evaluation. (A) Predictive model MAE over the 100 training epochs. (B) Predictive model MSE over the 100 training epochs.

similar profile (Fig. 8B). Final values for MAE and MSE were 0.1844 and 0.0165, respectively. The predictive model underwent two additional evaluations, using holdout cross-validation and 10-fold cross-validation methods (as was performed on the classification model). The holdout method led to an increase in the MSE (Table 5), while the 10-fold cross-validation had no impact on MAE or MSE. Surprisingly however, all folds demonstrated an extremely sharp drop in MAE following the first epoch. Lastly, 10-fold cross-validation of the prediction model demonstrated strong alignment between MAE and MSE (Figs. 9, 10).

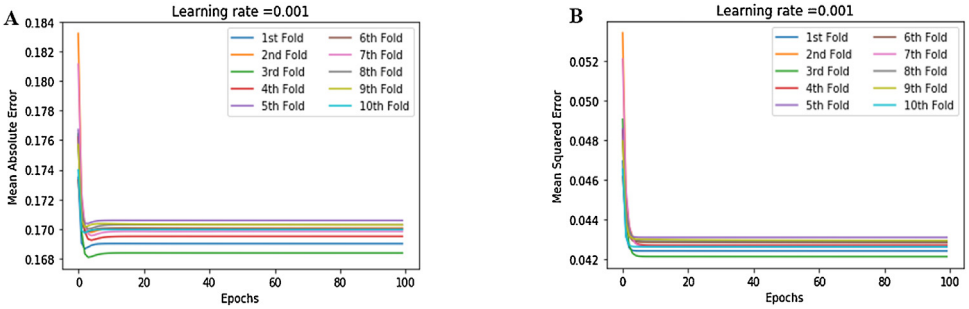

Fig. 9. Predictive model10-fold cross-validation evaluation (training).
(A) Predictive model MAE over the 100 training epochs per fold. (B) Predictive model MSE over the 100 training epochs per fold.

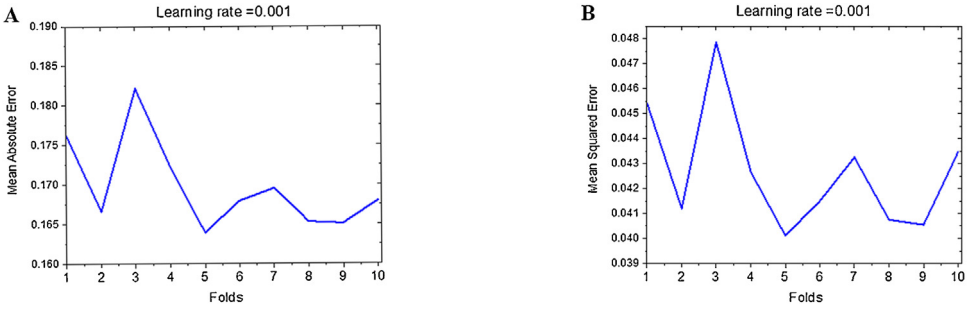

Fig. 10. Predictive model 10-fold cross-validation evaluation (validation).
(A) Predictive model MAE over the 10 validation folds. (B) Predictive model MSE over the 10 validation folds.

Table 4
A comparison between our classification model against other classification machine learning algorithms.

Method	Accuracy
Classification DNN (holdout cross-validation)	86 %
Classification DNN Average of 10-Fold Cross Validation	74 %
Classification DNN (Subject-based splits) with (holdout cross-validation)	83 %
Chance	32 %
Logistic Regression	77 %
SVM	43 %
HRV_{HIGH}/HRV_{LOW}	32 %

Table 5
A comparison between our predictive model against other predictive machine learning algorithms.

Method	MSE
Prediction DNN (holdout cross-validation)	0.01
Prediction DNN Average of 10-Fold Cross Validation	0.04
Prediction DNN (Subject-based splits) with (holdout cross-validation)	0.04
Logistic Regression	0.31
HRV_{HIGH}/HRV_{LOW}	1.45

3.4. Comparison to alternative methods

To gain insight into the utility and potential advantages of DNNs in the context of HRV obtained from the mobile device camera, we compared the accuracy of the classification model to chance and to other algorithms applied previously (Table 4). Random generation of stress levels achieved an accuracy of 32 %, which was slightly higher than expected given that there are four potential stress levels. The discrepancy likely resulted from the uneven distribution of the dataset. With respect to supervised learning, SVM and logistic regression achieved accuracies of 43 % and 77 %, respectively. Meanwhile, the "standard" HRV_{LOW}/HRV_{HIGH} was not better than chance (32 %). By far the best performing method on this particular dataset was the classification model, which achieved 86 % when using holdout cross-validation.

4. Discussion

Objective and remote quantification of stress at scale is a major focus of mHealth, not in the least because of the potential for mHealth tools to establish real world evidence that may reduce clinical trial cost and duration, while extending the lifespan of intervention monitoring beyond clinical trials. Here we present the analysis of two implicit DNNs trained with HR and HRV data obtained from the PPGI output of selfie videos collected with tandem self-reports of stress from individuals around the world through a popular meditation app. Despite the short sampling time (30 s) and rawness of the dataset, the neural networks were able to classify and predict stress with a high degree of accuracy and low amount of error, and did so within a computationally efficient number of epochs.

To understand just how important the use of machine learning is in the context of quantification of stress at scale, we also applied two alternative methods used previously for estimating cognitive stress from HR and HRV. Perhaps unsurprisingly, a HRV_{LOW}/HRV_{HIGH} algorithm that has been applied to assess stress from HRV calculated via 24 h ECG recordings [22], failed to predict stress beyond that which would be expected by chance. While a SVM was superior to chance, it proved only about half as predictive as the DNNs developed and examined here. It is likely that both the alternative algorithms were compromised by the unique nature (and arguably poor quality) of the 30 s PPGI, and these algorithms may have performed better when provided longer time course ECG recordings. This contemplation highlights how incredible it is that the DNNs described here are capable of such high accuracies from relatively crude RWD.

Potential additions to the DNNs directly examined here include forming low and high arousal categories of stress via a small explicit neuro-fuzzy classifier neural network [13], or the use of

Pioncare' plots of inter-point distance (beat-to-beat) measures of HR [23]. However, neither of these approaches is suitable for the dataset obtained from mobile camera PPGI once processed through the SAPHIRE algorithm [16] employed in the meditation app examined. Similarly, an implicit model calculating stress from a Lomb's estimator-derived algorithm is unlikely to be sufficiently robust to deal with the RWD from mobile device cameras.

In future, either after substantially larger datasets are available, or more accurate measures of HR and HRV emerge from mobile devices, one might hope to obtain levels of accuracy in predicting stress that greatly exceed 86 %. However, given the enigmatic nature of "cognitive stress", it may not be practical or even desirable to obtain measures much more accurate than those obtained here. Instead, future development may be better focused on predicting other elements of wellbeing, such as fatigue, distress, or emotional states.

Together our analyses demonstrate that the objective and remote quantification of stress at scale is achievable with relatively simple DNNs trained with RWD. Future work will explore the potential to predict mood from additional measures of the face.

CRediT authorship contribution statement

Abdulrhman H. Al-Jebrni: Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Brendan Chwyl:** Data curation, Formal analysis, Methodology, Software, Writing - review & editing. **Xiao Yu Wang:** Data curation, Formal analysis, Methodology, Software, Writing - review & editing. **Alexander Wong:** Funding acquisition, Investigation, Methodology, Resources, Supervision, Writing - review & editing. **Bechara J. Saab:** Data curation, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Visualization, Writing - original draft, Writing - review & editing.

Acknowledgements

This study was supported by an Ontario Centre for Excellence Grant awarded to Mobio Interactive Inc. and A.W., as well as funds from Mobio Interactive in support of BJS and AA.

Declaration of Competing Interest

BJS is the Chief Scientist and CEO of Mobio Interactive Inc, and he owned approximately 40% of the company at the time of manuscript acceptance. BJS did not directly contribute to, or have influence over, data collection or analysis. AA is an employee of Mobio Interactive Inc, and has been issued stock options amounting to less than 0.1% of the company at the time of manuscript acceptance. XYW is a former employee of Mobio Interactive Inc., and holds no stock options at the time of manuscript acceptance. No other authors have connections to Mobio Interactive Inc. beyond being research collaborators and no other authors in this study received financial compensation or any other form of compensation for the research undertaken herein.

References

[1] G. Bertson, J. Cacioppo, P. Binkley, B. Uchino, K. Quigley, A. Fieldstone, Autonomic cardiac control. III. Psychological stress and cardiac response in autonomic space as revealed by pharmacological blockades, Psychophysiology 31 (6) (1994) 599–608.
[2] G. Billman, H. Huikuri, J. Sacha, K. Trimmel, An introduction to heart rate variability: methodological considerations and clinical applications, Front. Physiol. 6 (2015).
[3] L. Marzano, A. Bardill, B. Fields, et al., The application of mHealth to mental health: opportunities and challenges, Lancet Psychiatry 2 (10) (2015) 942–948.
[4] R. Amelard, C. Scharfenberger, F. Kazemzadeh, et al., Feasibility of long-distance heart rate monitoring using transmittance photoplethysmographic imaging (PPGI), Sci. Rep. 5 (2015) 14637.
[5] M. Malik, J. Bigger, A. Camm, et al., Heart rate variability: standards of measurement, physiological interpretation, and clinical use, Eur. Heart J. 17 (3) (1996) 354–381.
[6] S. Rotenberg, J. McGrath, Inter-relation between autonomic and HPA axis activity in children and adolescents, Biol. Psychol. 117 (2016) 16–25.
[7] A. Noma, W. Trautwein, Relaxation of the ACh-induced potassium current in the rabbit sinoatrial node cell, Pflugers Arch. Eur. J. Physiol. 377 (3) (1978) 193–200.
[8] H. Brown, D. Difrancesco, S. Noble, How does adrenaline accelerate the heart? Nature 280 (5719) (1979) 235–236.
[9] J. Sztajzel, Heart rate variability: a noninvasive electrocardiographic method to measure the autonomic nervous system, Swiss Med. Wkly. 134 (35-36) (2004) 514–522.
[10] S. Porges, Cardiac vagal tone: a physiological index of stress, Neurosci. Biobehav. Rev. 19 (2) (1995) 225–233.
[11] J. Schmidhuber, Deep learning in neural networks: an overview, Neural Netw. 61 (2015) 85–117.
[12] L. Vanitha, G. Suresh, Hierarchical SVM to detect mental stress in human beings using Heart Rate Variability, in: 2014 2nd International Conference on Devices, Circuits and Systems (ICDCS), Combiatore, India: IEEE, 2014.
[13] A. Mand, J. Jia Wen, M. Sayeed, S. Swee, Robust stress classifier using adaptive neuro-fuzzy classifier-linguistic hedges, in: 2017 International Conference on Robotics, Automation and Sciences (ICORAS), Melaka, Malaysia: IEEE, 2017, pp. 1–5.
[14] S. Mayya, V. Jilla, V. Tiwari, M. Nayak, R. Narayanan, Continuous monitoring of stress on smartphone using heart rate variability, in: 2015 IEEE 15th International Conference on Bioinformatics and Bioengineering (BIBE), Belgrade, Serbia: IEEE, 2015, pp. 1–5.
[15] P. Mohan, V. Nagarajan, S. Das, Stress measurement from wearable photoplethysmographic sensor using heart rate variability data, in: 2016 International Conference on Communication and Signal Processing (ICCSP), Melmaruvathur, India: IEEE, 2016, pp. 1141–1144.
[16] B. Chwyl, A. Chung, R. Amerland, J. Deglint, D. Clausi, A. Wong, SAPPHIRE: stochastically acquired photoplethysmogram for heart rate inference in realistic environments, IEEE Int Conf Image Process ICIP 2016 (2016) 1230–1234.
[17] K. Walsh, B. Saab, N. Farb, Effects of a mindfulness-meditation app on subjective well-being: an active randomized controlled trial and experience sampling study, JMIR Ment. Health 6 (1) (2019), e10844.
[18] U. Subnis, N. Farb, K. Piedalue, M. Speca, S. Lupichuk, P. Tang, P. Faris, M. Thoburn, B. Saab, L. Carlson, The SEAMLESS Study: protocol for a clinical trial evaluating a SmartphonE App-based MindfuLnEss intervention for cancer SurvivorS. https://doi.org/10.2196/15178.
[19] M. Abadi, P. Barham, J. Chen, et al., TensorFlow: a system for large-scale machine learning, in: The 12th USENIX Conference on Operating Systems Design and Implementation, Savannah, GA, USA: USENIX Association, Berkeley, 2016, pp. 265–283.
[20] J. Schneider, A. Moore, A Locally Weighted Learning Tutorial Using Vizier 1.0, Carnegie Mellon University, Pittsburgh, PA, 2000, pp. 32–33.
[21] D. Kingma, Ba L. Adam, A method for stochastic optimization, in: International Conference on Learning Representations, San Diego: Ithaca, NY, 2015, pp. 1–13, arXiv.org.
[22] H. Kim, E. Cheon, D. Bai, Y. Lee, B. Koo, Stress and heart rate variability: a meta-analysis and review of the literature, Psychiatry Investig. 15 (3) (2018) 235–245.
[23] N. Bu, Stress evaluation index based on Poincaré plot for wearable health devices, in: IEEE 19th International Conference on e-Health Networking, Applications and Services (Healthcom), Dalian, China: IEEE, 2017, pp. 1–6.