Human Interaction and Emerging Technologies (IHIET 2024), Vol. 157, 2024, 230–239
https://doi.org/10.54941/ahfe1005482

An Exploration of Machine Learning and Reinforcement Learning for Emotional Well-Being

Kenneth Y. T. Lim[1], Ryan Thaddeus Shi Eng Chio[2], and Zhi Wen Lim[3]

[1]National Institute of Education, Singapore
[2]National Junior College, Singapore
[3]Independent Scholar, Singapore

ABSTRACT

With the high levels of stress in Singapore, mental and emotional well-being is an important health and social issue today. Research has shown the positive effects of pet ownership on mental and emotional well-being, however challenges of owning a pet in Singapore such as pet licensing restrictions, high costs, fear of losing a pet, a busy lifestyle and even allergies may deter pet lovers from owning a pet. Thus, we propose a technology-driven solution to emulate the useful effects of pets while mitigating the challenges of pet ownership. This project focuses on designing emotion recognition and reinforcement learning models as a stepping stone to individualise responses to a person's emotions. Our approach utilises the output from the emotion recognition model as an input in the proposed reinforcement learning algorithm. Hence, the paper first compares pre-trained and custom trained facial recognition models, and postulates the use of physiological signals via hardware sensors to further enhance the emotion recognition model. This is inspired from the ability of pets to perceive and respond to different emotions based on facial expressions and physiological signals like heart rate. The paper then outlines the development of novel K-Bandit algorithms in reinforcement learning tested on simulated reward functions, with the aim of optimising parameters for individualised responses to a person's emotions. Since reinforcement learning is typically used in simulation scenarios, this paper works towards developing a model that will eventually learn a person's preferences in real time by monitoring their emotional changes. To conclude, this project has showcased the feasibility of facial expressions and physiological signals for emotion recognition, and established the effectiveness of our proposed parameter optimisation functions in the K armed bandit reinforcement learning model to customise responses based on an individual's emotions. We hope this paper can act as a basis for future works in creating a human-friendly prototype to emulate man's best friend.

Keywords: Human systems integration, Machine learning, Reinforcement learning, Empathetic robotics, Human robot interaction

INTRODUCTION

Studies have shown that Singaporeans face high levels of stress (Cigna Singapore, 2022), so measures to improve mental and emotional well-being are becoming increasingly important. Singapore's recent launch of a national mental health strategy highlights this issue's significance (Teo, 2023). A study conducted by Boehringer Ingelheim Animal Health Singapore found that pets had a positive impact on the mental well-being of the majority of participants (Goh, 2021). Moreover, research has found that the presence of pets contributed to a lower resting heart rate as well as a faster recovery to resting heart rate when feeling stressed (Allen et al., 2002). However, Singaporeans may be deterred from pet ownership due to pet licensing restrictions for Housing & Development Board (HDB) flats (Housing & Development Board, n.d.), the high cost of owning a pet, estimated to be around $102,000 over a pet's lifespan (Wong, 2022), and the rising cost of veterinary fees and services by at least 10 - 20% due to inflation (Ee, 2023). Moreover, allergy considerations could discourage pet ownership, with allergic rhinitis typically triggered by pet dander affecting up to 13.1% of the Singapore population (Pharmaceutical Society of Singapore, 2023).

AIMS & OBJECTIVES

We propose a solution to incorporate artificial intelligence and robotics to emulate the useful effects of pets while mitigating the challenges of pet ownership, focusing on designing emotion recognition and reinforcement learning models as a stepping stone to individualise responses to a person's emotions. This paper aims to (1) contribute to existing literature regarding emotion recognition and (2) explore the use of reinforcement learning in individualising responses based on emotions. The paper will first delve into the emotion recognition details before investigating the use of the K Armed Bandit reinforcement learning model in individualising actions based on emotions detected. Given time and resource constraints, our research is focused on the theoretical aspects of the larger idea. Since the hardware can be developed from pre-existing frameworks, we scoped our research to the computational approach of emotion recognition and reinforcement learning taken for simulating pet ownership, which should be taken as the primary focus in this research endeavour.

METHODOLOGY

In the development of an emotional recognition model, no active participation was required from those asked to contribute to the dataset. Hence, there were no practices constituting "direct human participation" or involving their associated safety concerns in this stage of the investigation. Friends and family who consented, submitted images of their interpretations of the seven pre-identified basic emotions (anger, fear, happiness, sadness, surprise, neutral and disgust). The use of custom training data was hypothesised to allow for a wider range of possible expressions to be classified as the same emotion for more accurate emotion recognition in practical implementations.

The reinforcement learning aspect of research aimed to create an algorithm that allowed our robot prototypes to effectively and efficiently adapt to feedback provided on the effectiveness of its selected emotional response in real-time constraints. This focus was selected due to a gap of knowledge identified in current RL algorithms, whose training times were not emulative of actual learning rates. (The design of many RL algorithms and hence what was considered "efficient" in current literature was done for systems where training and interaction processes were done separately, as opposed to concurrently in our proposed design application.) We first create a testing ground for our algorithm by parameterising human-pet companion interactions. A basic three stage framework for an interaction was proposed (E.g. Approach > Interaction > Retreat) with each stage being composed of various mutable parameters (E.g. Approach considerations may be the speed of the approach, or the distance from the human a companion stops at) Background research regarding how interaction parameters generally evolved with respect to varying conditions allowed constraints to be made for the Reinforcement Learning Reward equation (the function to be maximised during the training process) allowing for more optimisations in our algorithms to be proposed. Each algorithm would then be tested with the metrics of accuracy and efficiency. Environmental conditions (noise and error in the input information) was also simulated to assess the stability of our proposed algorithm in applied conditions.

MACHINE LEARNING IN EMOTION RECOGNITION

A study conducted by the University of Veterinary Medicine Vienna indicates that dogs can distinguish and respond to 6 basic facial expressions: anger, fear, happiness, sadness, surprise, and disgust (Siniscalchi, d'Ingeo and Quaranta, 2018). Hence, we chose to explore facial expressions as a form of emotion recognition. Psychologist Dr Paul Ekman (1992) identified the same 6 emotions as universal emotions. Another study corroborated Ekman's claim of universal facial expressions, finding that people from different cultures share approximately 70% of facial expressions (Cowen et al., 2021). However, Jack et al. (2012) shed light on the differences in facial expressions across cultures, such as Chinese participants conveying emotion through the eyes or Western Caucasian participants expressing emotions with the eyebrows and mouth.

The FER2013 facial expression dataset comprises mainly of Caucasian and African Americans (Lukac et al., 2023). A custom dataset for Singapore was created through mass data collection to test the accuracy of a pre-trained model DeepFace against a custom trained Teachable Machine model (Teachable Machine, n.d.). The Teachable Machine model was neural network based, trained using the Tensorflow-Keras library in python. A 50/50 split was used for training and testing data. The custom trained model had 75 images for the training dataset of each emotion to ensure fair distribution amongst the different classes of emotion.

PHYSIOLOGICAL SIGNALS IN EMOTION RECOGNITION

Pets such as dogs are able to detect the rise or fall of heart rate (Deangelis, 2020), and can be trained to monitor blood oxygen saturation (SPO_2)

levels (Leomiti and Ausman, 2023). Hence, we chose to explore the use of physiological signals in addition to facial expression to increase the accuracy of emotion recognition. A study on the reliability of physiological responses induced by basic emotions concluded that physiological features, such as skin conductance level (SCL), heart rate (HR) and blood volume pulse (BVP), would be important for emotion recognition in the human computer interface area (Jang et al., 2019). Other studies have found that blood oxygen levels (SPO2) were also a possible physiological feature to classify emotions (Alkawaz et al., 2015). Stress, typically associated with fear and anger (Gu et al., 2019), as well as sadness which may lead to crying (Raypole, 2020), can lower blood oxygen level (Cox, 2022). Emotions can lead to differences in blood oxygen levels (Alkawaz et al., 2015). When interfacing the MAX30102 heart rate sensor on the Arduino, output data of SPO2 levels and heart rate was obtained. This shows the feasibility of including physiological signals to enhance emotion recognition.

From the accuracy results of both pre-trained and custom trained emotion recognition models, it is shown that the custom trained Teachable Machine model has a higher average accuracy than that of the pre-trained DeepFace model (as shown in Table 1).

Table 1. Accuracy of emotion recognition models.

Emotion Recognition	Deep Face	Teachable Machine
Happy	45.3%	78.7%
Sad	28.0%	73.3%
Surprise	42.7%	57.3%
Angry	28.0%	53.3%
Fear	57.3%	74.7%
Disgust	8.0%	12.0%
Neutral	77.3%	48.0%
Average	40.9%	56.8%

It should be noted that disgust is the emotion with the lowest accuracy score. An article exploring the use of a convolutional neural network trained on the same FER2013 dataset also had poor accuracy in classifying disgust, citing the smaller number of disgust facial expressions in the training dataset as the limitation (Tang, 2019). However, the custom trained Teachable Machine model had the same number of images for all classes of emotion, and still had a significantly low accuracy score for disgust. Thus, further research could be undertaken in this area to investigate the difficulty in classifying this emotion.

REINFORCEMENT LEARNING ALGORITHMS TO SIMULATE ORGANIC LEARNING BEHAVIOUR

Reinforcement learning algorithms are fundamentally maximisation problems: they evolve policy algorithms to select actions that optimise reward functions derived from intended outcome behaviours. We aim to make this process more efficient for practical simulation of organic learning behaviour. With the parameter framework in our proposed robot design, our proposed use of reinforcement learning intends to tune the action parameters based on a Reward function that reflects positive changes in sensor data when a specific action configuration is executed.

For the proposed inputs into the standard reinforcement learning model, the State is determined through the facial detection model described above. The Reward is determined by the reward function R(x) which encompasses the "positive changes" detected in sensors. Determining the actual reward functions and sensor response correlations are not discussed in this paper and are proposed as future work to be completed when progressing towards eventual testing and deployment.

We discuss assumptions and constraints that can be made to the reward function R(x) where optimisations to the action selection algorithm may be proposed. It is asserted that there is only one optimal parameter value at any given state of companion owner interaction. While the optimal parameter itself may shift with a greater number of encounters, multiple configurations of a robot's parameters cannot be equally preferred. It is assumed that for reward functions where parameters may appear to have two or more similarly preferred values, that only one of such is eventually optimised. In a reward function that has only one stationary maxima, and when provided with the range where this point may lie, the optimum point can be estimated through analysing the function value at regular intervals k within the specified range. The subsequent removal of values that return lower values of the function, and addition of values between those who return higher values of the function, would shift the range in a manner such that the exact maxima point would eventually be found. This principle of successive halving is often applied in the hyper-parameter optimization of neural network configurations. This algorithm can reliably determine the optimum point in single maxima functions within the time constraints suitable for authentic pet-human interactions. Research on human-robot interaction has emphasised the need for individualised technology development (Søraa et al., 2022), with a study investigating living with robots noting that participants became increasingly comfortable with the robot approaching them closely over time (Mehta and Losey, 2023). The assumption made is that there is a single optimising value for each mutable parameter in the robots action array, resulting in a reward function with a single maximum. This allows for parameter optimization through a sequential decision-making algorithm.

THE K ARMED BANDIT MODEL

The Stochastic K-Armed Bandit is a Single State Markov Decision Process that frames *sequential decision-making under uncertainty problems*. It abstracts the action selection as: within a K number of slot machines that return Reward R drawn randomly from a fixed probability distribution, and maximises the cumulative rewards for N slot machines and an array of actions $\{a_{i=1}, a_{i=2}, \dots a_{i=N}\}$, where the action selected at a time step t is denoted as A_t. This results in reward R_t, the mean reward of the action number i, v_i is denoted as:[28]

$$v_i = v(a_i) = E[R_k|A_k = a_i] \tag{1}$$

v_i is initially unknown and is approximated with greater time steps t as:

$$\widehat{v}_{i,\,t} = \widehat{v}_t(a_i) = \frac{\sum_{t-1}^{j=1} R_{i,j}}{n_{i,k}} \tag{2}$$

where $R_{i,j}$ is the reward from action a_i at step t and $n_{i,k}$ is the selection count of the action a_i prior to 2. Rearranging (2), a recurrence relation is obtained, to update the estimated value $\widehat{v}_{\mathrm{t}}$ with each step of the algorithm:

$$\widehat{v}_k = \frac{1}{t + 1}(R_{t-1} - \widehat{v}_{k-1}) \tag{3}$$

The selection of an action uses the epsilon greedy approach: for specified number ϵ and randomly generated number p, if $p \succeq \epsilon$, an action is selected by taking the largest reward $A_t = argmax_{a_i}\{v_{i,t}, v_{i+1,t}\ldots\}$, whereas otherwise randomly selected from the action array. In this way, a balance is struck between the maximisation of historically rewarded actions, and the exploration of new potential strategies. Such an approach can be applied to the algorithm described in 3.1 to approximate Reward function maximas, with the range divided into k intervals being taken as the action array and with the Reward value R taking up the value of the reward function $R(x)$ instead of a static fixed value. The subsequent section describes this implementation. For the Reward Function $R(x)$ with no local maximas and given an initial search range $[a, a_n]$, this range is divided into a k array of "arms" $\mathbf{A} = \{a_1, a_2, \ldots a_n\}$, then fed into a k number of slot machines, which returns their corresponding input from the reward function e.g. $\{R(a_1), R(a_2), \ldots R(a_n)\}$. The K bandit model in 3.2 selects from the array of slot machines and "plays" until the proportion of the most historically selected actions from all choices exceeds a specified confidence value C. Thereafter, the average rewards for each "arm" would be sorted, discarding "arms" from a lower performing half and pruning the array down to a size $\lfloor \frac{k}{2} \rfloor$. The array is then resorted, with new arms being inserted between neighbours as their mean, obtaining a new set of arms $\mathbf{A}$ of original length k. The K bandit model then plays and prunes recursively, until the range of $\mathbf{A}$ falls below a specified convergence value, μ. As such, an accurate range of values for where the Reward maxima lies is obtained. The *successive halving* algorithm was tested: Provided with an initial range from **1 to 60**, the hidden maxima value of 31.5 was reached by the algorithm after **6** interactions (translated as 6 different companion, pet encounters), demonstrating the capability of *successive halving* in tuning parameters within a relatively few number of iterations.

ALGORITHM 2 GRADIENT ASCENT

While *Successive Halving* is an efficient and effective algorithm for determining maximas of reward functions that remain unchanged throughout the duration of interaction, it is unable to adjust for changes should the optimal value of the reward function shift during the course of the interaction. This presents an important drawback to be considered, as the optimum parameter value seldom remains unchanged with respect to the time step of the model. For example, it can be reasonably asserted that human companion interactions become more trusting with an increased number of encounters. In such a case, the optimum value for parameters such as the distance of approach would vary with the extent of trust. As such we draw from various gradient descent techniques commonly applied in machine learning to allow for our model to dynamically shift along with the Reward equation when a change is detected, thus taking into account that human preferences may change over time. For the Reward Function $R(x)$ including an additional parameter t, $R(x,\ t)$ whereby for all parameters t in $R(x,\ t)$ the corresponding $R(x)$ contains no local maximas, and provided an initial array of parameters P, gradient ascent can be used to recursively shift the parameter selection range towards the new function maxima: $R_{x+1,t} = R_{x,t} + \alpha\frac{\partial R}{\partial x}$ (where α is the "learning rate"). The value of $\frac{\partial R}{\partial x}$ is approximated by taking the gradient between the Reward function at the maximum and minimum values in the parameter array **P**: $\frac{\partial R}{\partial x} \approx \frac{R(argmax(P),\ t)\ -\ R(argmin(P),\ t)}{argmax(P)\ -\ argmin(P)}$. The model would default to responding either as the current maximum or minimum in the parameter array, depending on the remaining value required to approximate the gradient. α is taken as the range(**P**). The bounds of **P** are then shifted by the value $\alpha\frac{\partial R}{\partial x}$, with an additional step added in the step direction (positive or negative and taken as the opposite sign to the initial gradient: $\frac{\partial R}{\partial x}_1$). As such values of **P** progresses towards the maxima, resizing in the direction of search to favour values closer to convergence at later recursions. When $\frac{\partial R}{\partial x}$ falls below a "find condition value" λ ($\frac{\partial R}{\partial x} \to 0$ approaching the maxima), *successive halving* with input P is then used to determine the final new maxima range.

Models utilising sensor data as input are subject to significant amounts of noise and distortion, owing to minute uncontrolled variations in the environment of choice, and inherent uncertainties within our chosen instruments, as such it is imperative that the proposed algorithms maintain a high level of accuracy outside of the assumption of a smoothed reward function. We select the reward equation for successive halving arbitrarily as a distribution function with a maxima of 0.40 at parameter 51.5. Noise was simulated by randomly scattering inputs about the range of the true function value, whereby distortion could be varied through changing the range variable r. The convergence and error scatter plots for 5 trials, r = [0.01, 0.02,... 0.05] are below. The range of R(x) in our test selection was [0.00, 0.40] with a peak distortion r = 0.05 translating to a 12.5% error overall. Despite this, successive Halving maintained a high accuracy range between 99.2 - 99.8% when varying the distortion values, with all configurations converging after 4–7 interactions. This demonstrates the resilience in the algorithm to a reasonable level of distortion within input data. Gradient Ascent was similarly tested where a quadratic reward function R(x) was switched after the 10th stable interaction from a maxima of 31.5 to 11.5, with a similar 99.0% - 99.8% accuracy range, converging 10–14 interactions after the change in reward function, demonstrating also a resilience to a reasonable level of distortion within input data.

Variations in the reward function's shape were initially hypothesised to have significant effects on the number of interactions required for models to converge towards maximas: this is significant given the amount of variability that reward functions may take between individuals in our selected use case, and may present a significant design flaw to be imperatively addressed. We similarly arbitrarily select a distribution function with a maxima of 0.40 at parameter 51.5, but include additional parameters to allow variations in the equation shape, in the dimensions of translation and scaling. Notably scaling parallel to x and y axis had no effect on the number of interactions required to converge to stability, when translation between the range of the function was performed, an interesting pattern emerged, indicating that the number of interactions required for convergence increases as the maxima approaches the bounds of the specified function range but also contains periodic regions at certain translation values where an increase in the number of interaction is noted.

EVALUATION AND CONCLUSION

The single maxima constraints may result in reward functions being unsupported by the gradient ascent algorithm, as there is currently a possibility of approaching a local maxima. Additionally, both algorithms are susceptible to butterfly effect conditions, particularly the convergence value in Algorithm 1 and "find condition" in Algorithm 2 as values greater than a specified value would often result in highly inaccurate final parameters. While a solution was found in setting both variables to be extremely small, there was a significant tradeoff in the number of interactions required. Tests however demonstrate that successive halving is capable of converging parameter values towards an unknown optimal value within a reasonable number of iterations. Additionally, further optimisations to the gradient ascent algorithm are proposed to minimise its susceptibility and improve its performance. Further research can be conducted to address limitations in this project. First, further experimentation to determine the extent of effectiveness of incorporating physiological signals in emotion recognition. Second, more would need to be done for the development of reward equations with sensor data as input in real world conditions. Finally, further research of a physical prototype for mass testing on human participants.

To conclude, this project has achieved its goals of showcasing the feasibility of facial expressions and physiological signals for emotion recognition. This project has also established the effectiveness of our proposed parameter optimisation functions in the K armed bandit reinforcement learning model to customise responses based on an individual's emotions. We hope that this project has provided a new perspective as to how artificial intelligence and robotics can be used to emulate the positive effects of pets, and acts as a basis for future works in creating a human-friendly prototype to emulate man's best friend.

REFERENCES

Alkawaz, M. H., Mohamad, D., Basori, A. H., Saba, T. and Rehman, A. (2015). 'The Correlation between Blood Oxygenation Effects and Human Emotion towards Facial Skin Colour of Virtual Human', *3DR Express*, **6**(2), article no: 13. doi: 10.1007/s13319-015-0044-9.

Allen, K., Blascovich, J. and Mendes, WB. (2002). 'Cardiovascular reactivity and the presence of pets, friends, and spouses: The truth about cats and dogs', *Psychosomatic Medicine*, **64**(5), pp. 727–39. doi: 10.1097/01.psy.0000024236.11538.41.

Cigna Singapore (2022). *Stressed in Singapore*. Available at: https://www.cigna.com.sg/static/docs/pdf/Cigna_Singapore_360_Well-being_Report_2022.pdf (Accessed: 5 January 2024).

Cowen, A. S., Keltner, D., Schroff, F., Jou, B., Adam, H. and Prasad, G. (2021). 'Sixteen facial expressions occur in similar contexts worldwide', *Nature*, **589**(7841), pp. 251–257. doi: 10.1038/s41586-020-3037-7.

Cox, J. (2022). 'Can Stress Cause Low Oxygen Levels?', *PsychCentral*, 12 August. Available at: https://psychcentral.com/stress/can-stress-cause-low-oxygen-levels (Accessed 5 January 2024).

Deangelis, S. (2020) 'What Can Dogs Detect?', *Wag!*, 6 April. Available at: https://wagwalking.com/sense/what-can-dogs-detect (Accessed: 5 January 2024).

Ee, A. (2023). 'More Abandoning Pets as Costs to Take Care of Them Rise', *The Straits Times*, 11 May. Available at: https://www.straitstimes.com/singapore/more-abandoning-pets-as-costs-to-take-care-of-them-rise (Accessed: 05 January 2024).

Ekman, P. (1992). 'An argument for basic emotions', *Cognition and Emotion*, **6**(3–4), pp. 169–200. doi: 10.1080/02699939208411068.

Goh, T. (2021). '89% of owners surveyed feel pets help improve mental health during Covid-19 pandemic', *The Straits Times*, 24 December. Available at: https://www.straitstimes.com/singapore/health/89-of-owners-surveyed-feel-pets-help-improve-mental-health-during-covid-19-pandemic (Accessed: 05 January 2024).

Gu, S., Wang, F., Patel, N. P., Bourgeois, J. A. and Huang, J. H. (2019). 'A Model for Basic Emotions Using Observations of Behavior in Drosophila', Frontiers in Psychology, 10(781). doi: 10.3389/fpsyg.2019.00781.

Housing & Development Board. (n.d.). *Keeping Pets*. Available at: https://www.hdb.gov.sg/residential/living-in-an-hdb-flat/keeping-pets (Accessed: 5 January 2024).

Jack, R. E., Caldara, R., & Schyns, P. G. (2012). 'Internal representations reveal cultural diversity in expectations of facial expressions of emotion', *Journal of experimental psychology. General*, **141**(1), pp. 19–25. doi: 10.1037/a0023463.

Jang, EH., Byun, S., Park, MS. and Sohn, JH. (2019) 'Reliability of Physiological Responses Induced by Basic Emotions: A Pilot Study', *Journal of Physiological Anthropology*, **38**(1), article no: 15. doi:10.1186/s40101-019-0209-y.

Leomiti MD, Ausman C. (2023). 'EMS Canine Evaluation of Oxygen Saturation (SpO2)', in StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing. Available at: https://www.ncbi.nlm.nih.gov/books/NBK597366/ (Accessed 5 January 2024).

Lukac, M., Zhambulova, G., Abdiyeva, K. and Lewis, M. (2023). 'Study on Emotion Recognition Bias in Different Regional Groups', *Scientific Reports*, 13 (1), article no: 8414. doi:10.1038/s41598-023-34932-z.

Mehta, S. A. and Losey, D. P. (2023). 'Unified Learning from Demonstrations, Corrections, and Preferences during Physical Human-Robot Interaction', To be published in *ACM Transactisons on Human-Robot Interaction* [Preprint]. doi: 10.1145/3623384 (Accessed 5 January 2024).

Pharmaceutical Society of Singapore (2023) Allergic Rhinitis. Available at: https://www.healthhub.sg/a-z/diseases-and-conditions/allergic-rhinitis (Accessed: 5 January 2024).

Raypole, C. (2020). '6 Surprisingly Practical Reasons We Cry', *Healthline*, 28 August. Available at: https://www.healthline.com/health/why-do-we-cry (Accessed: 5 January 2024).

Siniscalchi, M., d'Ingeo, S. and Quaranta, A. (2018) 'Orienting Asymmetries and Physiological Reactivity in Dogs' Response to Human Emotional Faces', *Learning & Behavior*, **46**(4), pp. 574–85. doi: 10.3758/s13420-018-0325-2.

Søraa, R. A., Tøndel, G., Kharas, M. W. and Serrano, J. A. (2022) 'What Do Older Adults Want from Social Robots? A Qualitative Research Approach to Human-Robot Interaction (HRI) Studies', *International Journal of Social Robotics*, **15**(3), pp. 411–424. doi: 10.1007/s12369-022-00914-w.

Tang, T. (2019). 'How Are You Feeling? Answered with A Real-Time Emotion Detector', *Medium*, 25 December. Available at: https://medium.com/analytics-vidhya/building-a-real-time-emotion-detector-towards-machine-with-e-q-c20b17f89220 (Accessed: 5 January 2024).

Teachable Machine (n.d.) *Teachable Machine*. Available at: https://teachablemachine.withgoogle.com/ (Accessed: 5 January 2024).

Teo, J. (2023) 'S'pore Launches National Mental Health and Well-Being Strategy', *The Straits Times*, 6 October. Accessed: https://www.straitstimes.com/singapore/health/s-pore-launches-national-mental-health-and-well-being-strategy (Accessed: 5 January 2024).

Wong, S. (2022) 'Want a Dog? Be Prepared to Set aside over $100k', *The Straits Times*, 27 November. Available at: https://www.straitstimes.com/business/invest/i-ve-always-wanted-a-dog-but-it-could-cost-over-100k (Accesed: 5 January 2024).