A Comparison of Humans and Machine Learning Classifiers Categorizing Emotion from Faces with Different Coverings

Harisu Abdullahi Shehu^{a,*}, Will N. Browne^b, Hedwig Eisenbarth^c

^a*School of Engineering and Computer Science, Victoria University of Wellington,
New Zealand*

^b*School of Electrical Engineering and Robotics, Queensland University of Technology,
Australia*

^c*School of Psychology, Victoria University of Wellington, New Zealand*

Abstract

Partial face coverings such as sunglasses and face masks unintentionally obscure facial expressions, causing a loss of accuracy when humans and computer systems attempt to categorise emotion. With the rise of soft computing techniques interacting with humans, it is important to know not just their accuracy, but also the confusion errors being made—do humans make less random/damaging errors than soft computing? We analyzed the impact of sunglasses and different face masks on the ability to categorize emotional facial expressions in humans and computer systems. Computer systems, represented by VGG19, ResNet50, and InceptionV3 deep learning algorithms, and humans assessed images of people with varying emotional facial expressions and with four different types of coverings, i.e. unmasked, with a mask covering the lower face, a partial mask with transparent mouth window, and with sunglasses. The first contribution of this work is that computer systems were found to be better classifiers (98.48%) than humans (82.72%) for faces without covering (>15% difference). This difference is due to the significantly lower accuracy in categorizing anger, disgust, and fear expressions by humans ($p's < .001$). However, the most novel aspect of the work is identifying how soft computing systems make different mistakes

*Corresponding author

Email address: harisushehu@ecs.vuw.ac.nz (Harisu Abdullahi Shehu)

to humans on the same data. Humans mainly confuse unclear expressions as neutral emotion, which minimizes affective effects. Conversely, soft techniques often confuse unclear expressions as other emotion categories, which could lead to opposing decisions being made, e.g. a robot categorizing a fearful user as happy. Importantly, the variation in the misclassification can be adjusted by variations in the balance of categories in the training set.

Keywords:

Emotion classification, Emotion categorization, Face masks for COVID-19, Facial expression, Sunglasses

1. Introduction

Facial expressions convey information regarding human emotion in non-verbal communication. Thus, an accurate categorization of a person’s emotional facial expression contributes to understanding their needs, intentions, and potential actions. This is a valuable insight for soft computing applications such as when robots are introduced into shared workspaces. Thus, computational methods are needed to analyze facial expressions, but how do they differ from human classifiers, especially in different circumstances?

Face coverings such as sunglasses and face masks are used to cover the face for different reasons. For instance, people use sunglasses to either protect the eyes from sunlight or to improve appearance, and people use face masks often to prevent the spread of infectious diseases such as the coronavirus (COVID-19) or for hygienic reasons in hospitals or food preparation. However, these coverings not only cover parts of a face but might also affect social interaction as they obscure parts of facial expressions, which might contain socially relevant information even if the expression does not represent the emotional state of a person. While sunglasses cover an estimated 10-15% of the face, approximately 60-65% of the face is covered with face masks — exact numbers are hard to measure due to variance across different people, dependent on the size of the face [1]. More importantly, these coverings obscure different parts of the face,

which convey information about emotional expressions to a varying extent [2].

Different studies have shown that humans are far from perfect in accessing the emotional *states* of people by just looking at their faces [3]. Further, humans are also far from perfect when categorizing emotional *labels* from facial expressions [4]. Occlusion of the face leads to an even higher decrease in the performance of humans in categorizing emotion from the faces of people. For instance, it has been observed that partial occlusion of the eyes [5] impaired the achieved performance by humans especially for categorizing sad expressions [6]–[8]. While negative emotions such as sadness and fear are the most difficult to categorize by humans [9], [10] compared with positive emotions like happiness [11]–[14], obscuring the mouth leads to less accurate categorization even for happiness [6].

More specifically, Carbon [1] investigated the impact of face masks on humans categorizing six different categories of emotion (anger, disgust, fear, happy, neutral, and sad) from images of people. They found a significant decrease in performance of their participants across all expressions except for fear and neutral expressions. In addition, emotional expressions such as happy, sad, and anger were repeatedly confused with neutral expressions whereas other emotions like disgust were confused with anger expressions. In accordance with Carbon, several studies confirmed the detrimental effect of face masks on emotion categorization by humans [5], [15], [16]. However, none of these studies directly compared the impact of face masks on humans with machine learning classifiers (computer systems).

In contrast to humans, computer systems have done well in categorizing emotion from in-distribution faces images (images from the same dataset) of people, achieving an accuracy of over 95% [17]. Nevertheless, the performance of these systems is greatly affected when it comes to categorizing emotion from natural images where the sun is free to shine and people are free to move, changing dramatically the appearance of faces in these images [18]. For instance, a change in the colour information of the pixels, e.g. foreground colour of the whole or a small region of the face, which are nonlinearly confounded with

emotional expressions, can lead to a decrease in the performance of these models by up to 25% [19]–[21]. Furthermore, face coverings such as sunglasses and face masks affect the performance of computer systems by up to 74% [22]. However,
 55 it is unknown to which extent the same type of sunglasses and face masks would affect the performance of humans in direct comparison to computer systems in emotion categorization of the same stimuli.

We aim to analyze the limitations of humans and computer systems in categorizing emotion of covered and uncovered faces, as well as the misclassification
 60 cost of the computer systems.

The contributions of this work are through the following tasks:

- Compare the impact of sunglasses and different types of face masks (i.e. fully covered and recently introduced face masks with a transparent window [23]) on humans’ and computer systems’ ability to categorize emotional facial expressions. This direct comparison using the same stimuli
 65 across different emotion categories and different standardized face coverings addresses open questions about the different importance of facial features for both types of observers.
- Analyze the misclassification cost for the computer systems. This has
 70 practical importance as when positive expressions are misclassified as negative emotional expressions, artificial intelligence systems might refrain from interacting with a person. Conversely, the systems might continuously try to interact with a person if an expression such as anger or sadness were misclassified as happiness, which could lead to potential problems in
 75 the interaction of humans with these systems.
- Introduce an approach to minimize the misclassification costs for the computer systems while maintaining an equivalent (or minimally reduced) accuracy.

The main reason for analyzing the misclassification cost for the computer
 80 systems is because classifier induction often assumes that the best distribution

of classes within the data is an equal distribution, i.e. where there is no class imbalance [24]–[27]. However, research has shown that this assumption is often not true and might lead to a model that is very costly when it misclassifies an instance [28]–[30]. This might also be true for emotion classification [22], e.g.
 85 when the fully covering mask was used to cover the face, emotion categories like happiness were classified badly by the computer systems due to the majority of the happy expression images being misclassified as anger expression. Thus, a secondary objective is to investigate how data distribution affects misclassification for computer systems.

90 This is important because it is not the error rate that matters. In most cases, what matters most is “where” the misclassification occurs, i.e. the negative effect that may be caused by computer systems misclassifying human facial expressions. For these reasons, the cost of misclassification needs to be assessed for computer systems trained with data that is oversampled based on an ex-
 95 pression that is considered to be of less threat. This is anticipated to change the behaviour of the computer systems such that if they misclassify unclear expressions to expressions, it will cause the least harm. Therefore, as additional analysis, we will introduce an approach to minimize the misclassification costs for the computer systems.

100 Considering that the face coverings balance ecological validity with standardization, by obscuring facial expressions as it happens with varieties of face masks and sunglasses, while still being standardized across the individuals and images in the dataset. Understanding these limitations and regulating the cost of misclassifying faces with these coverings can inform the production of artificial
 105 emotionally intelligent systems [31] when it comes to real-world situations.

The rest of the paper is organized as: Section 2 describes the machine learning classifiers used as examples of artificial emotion classification systems and provides details on how the classifiers were set up in each experiment. The section also describes the human sample, the survey set up, the tasks needed to
 110 be performed by the participants during the survey, and the analysis procedure used. Section 3 presents the obtained results, along with the statistical analysis

carried out. Section 4 provides a further discussion on the obtained results, and finally, Section 5 concludes the paper.

2. Method

115 This study involved collecting and analyzing data from humans and computer systems to compare their accuracy and confidence in categorizing emotion from covered and uncovered faces.

2.1. Human sample

Ethical approval for this study has been obtained from the Human Ethics
120 Committee (HEC) of Victoria University of Wellington and the study was pre-registered on OSF (<https://osf.io/mgx9p>). Participants gave informed consent for participation in the study and could enter a raffle/draw to win a \$50 voucher for compensation.

All participants were recruited online via SurveyCircle and social media sites
125 such as Facebook, Twitter, Instagram, etc. through snowball sampling. A power calculation based on the effect size obtained from categorizing the “*afraid*” class from Eisenbarth et al’s [32] study was performed to estimate the required sample size, which was found to be 52. However, the study aimed to recruit at least 100 participants to prepare for potential missing data.

130 Of 154 people who started participating in the survey, 82 participants completed the full study. Of those 82 participants, 43 were male, 37 were female, and 2 either did not specify their gender or identified themselves to be from other groups. The age of the participants ranged from 19 to 70, with a mean of 27.88 (SD = 7.53). 28 of the participants self-identified as white, 11 as Black or
135 African American, 24 as Asian, 4 as Arab, 3 as multicultural, and the remaining 12 either did not answer the question or identified themselves to be from other groups.

Data were collected via SoSciSurvey (<https://www.soscisurvey.de/>) between the 21st of December 2020 to the 21st of February 2021 during the
140 COVID-19 pandemic.

2.2. Computer systems sample

The computer systems sample was represented by the artificial emotion classification systems. Three of the most commonly used deep learning (i.e. the Visual Geometry Group (VGG) [33], the Residual Neural Network (ResNet) [34], and the Inception (Inception) [35]) models were used as examples of a machine learning algorithms (computer systems).

The Visual Geometry Group (VGG) model is one of the most popular convolutional neural network architectures. The network made significant improvements over AlexNet [36] by replacing the large kernel-sized filters with several 3x3 kernel-sized filters. There are two types of VGG models; the VGG16 and the VGG19 model. VGG19 [33] was chosen to be used as it has more weight layers (i.e. 19) compared to its pair (VGG16) with 16 weight layers. Another reason for selecting VGG19 is because it is a well-studied standard model that has achieved high performance [37], [38].

The residual neural network (ResNet) is a powerful neural network that utilizes skip connections, or jumps over certain layers, to avoid the problem of vanishing gradient and mitigates the degradation problem, which causes poor learning for deep networks. There are different types of ResNet such as ResNet18, ResNet34, ResNet50 etc. ResNet50 [34] was chosen as it improves the efficiency of the network with more layers while minimizing the percentage classification error.

InceptionV3 is a widely used convolutional neural network architecture that is made up of symmetric and asymmetric building blocks. The model applied batch normalization to activation functions throughout the model to propagate label information across the network, which stabilizes the learning process. There are different types of the Inception network such as Inception version 1 (InceptionV1), Inception version 2 (InceptionV2), Inception version 3 (InceptionV3), etc. InceptionV3 [35] was chosen because it provides several improvements, e.g. use of auxiliary classifier and factorization of convolutions, to its previous versions. As a result, the network is faster and lighter compared to versions 1 and 2.

The models were set up using the Keras API [39], [40] to use the same number of layers as originally introduced in each corresponding paper [33]–[35]. The models were all set up to run for 200 epochs, starting with an initial
 175 learning rate of 0.001. The learning rate was reduced by 10% after 80, 100, 120, 160 epochs, and by 5% after 180 epochs to avoid overfitting¹ the data. Note that the computer systems were trained with images from the last-half frames of the CK+ dataset based on the technique developed by Shehu et al. [41].

The accuracy of the deep models was obtained by dividing the number of cor-
 180 rectly predicted test images by the total number of presented test images, which in this case is 35 for each type of covering. In deep models, the prediction of each image returns the probability of associating the image to a particular emotion category where the summation of all the probabilities is equal to 1 ($\sum_{i=1}^n P_i = 1$). The predicted class is the class with the highest probability ($\max_{x \in [1, \dots, n]} P(x)$)
 185 where n is the number of emotion categories). These probabilities are usually considered as the confidence of the model.

The three computer systems (VGG19, ResNet50, and InceptionV3) algorithms were included to test generalizability. However, detailed results of only the VGG19 have been provided as it has achieved higher performance com-
 190 pared to ResNet50 and the InceptionV3 models on partially covered images of the CK+ dataset.

Due to the stochastic nature of processes and the non-deterministic nature of the deep models, the model was run 30 times and the results presented are an average obtained from the 30 runs. Thus, the study compared the survey
 195 results obtained from 82 human participants with the 30 computer systems.

¹Overfitting occurs when a model learns details, including noise in the training data to such an extent that it has a negative impact on the performance of the model on new data. As a result, the model performs too well on the training data, but poorly on the test data.

2.3. Stimulus Material

Seven categories of emotion, i.e. the six basic emotions (*anger*, *disgust*, *fear*, *happy*, *sad*, and *surprise*) defined by Ekman [42], as well as the neutral expression were included. The experiment consisted of four blocks presenting facial expressions unmasked, masked, partially masked, and with sunglasses. The unmasked block consisted of images with nothing added to the face. Five images were used for each of the seven categories, i.e. a total of 35 images are used. The images used mainly comprised images of different people such that participants could not classify the facial expression from a previous image seen. Similarly, the masked block consisted of 35 emotional images of people with face masks added to the images. The partially masked block consisted of 35 emotional images of people with a partial mask, i.e. a face mask with a transparent window added to the images and finally, the sunglasses block consisted of 35 images of people with sunglasses added to the images. All the images used in the experiment are emotional images from the CK+ dataset as images from the dataset have recently been used in numerous and wide ranging studies [43]–[47]. However, only one section (unmasked) consisted of original images from the CK+ dataset. The remaining three blocks (mask, partial mask, and sunglasses) consisted of simulated images with different kinds of face masks, as well as sunglasses. Face masks and sunglasses were added to the images using a strategy that was implemented in our previous research [22]. Fig. 1a shows sample original and simulated emotional images of the six basic emotions, plus the neutral expression used here.

This research used a total of 3018 images from the last-half frame images of the CK+ dataset to train the computer systems based on the technique developed by Shehu et al. [17], [41]. The test images were randomly selected from the last-half frame images and not used in system training. Thus, the test images are a combination of images from mid to peak or highest emotional intensity frames.

(a)

What emotion do you think this person is showing?

Anger	Fear	Neutral	Surprise
Disgust	Happy	Sad	

How confident are you that the image belongs to the category you have selected?

Next

(b)

Figure 1: Sample images of faces: a) unmasked (first row), masked (second row), partial mask (third row), and sunglasses (last row) from the CK+ dataset. Left to right: anger, disgust, fear, happy, neutral, sad, and surprise expressions. [41] b) example trial for human participant.

235 2.4. Descriptive questions

Human participants were asked to answer a set of questions such as their age, gender, and ethnicity. Furthermore, we asked the participants three additional questions to understand the extent of their prior experience with people wearing masks.

230 The three questions asked to the participants were; “1) For how long did you interact with people wearing masks on a daily basis during the past 3 months?”, “2) On average, for how much time per day (in percentage) did you interact with people wearing masks during the past 3 months?”, and “3) Of the people you saw on a day-to-day basis during the past 3 months, approximately [what]
235 percentage [was] wearing masks on average?”. All questions were asked prior to taking the survey and participants answered using a slider ranging from 0 (“not at all”) to 100 (“very often”). A mean exposure score (Q_{mean}) was computed, averaging the three questions ($M = 36.15$, $SD = 27.62$).

2.5. Emotion categorization task

240 Fig. 1b shows a sample trial. For each trial, participants were required to choose the emotion category that the image most closely related to, as well as their confidence in choosing the category of the emotion. There was a total of 140 trials across the 7 categories. All emotion categorization questions were required to be answered and participants were not taken to the next page unless
245 the emotion category and the confidence rating were made (see Supplementary Materials).

The data collected, as well as the stimulus set used, can be accessed via OSF (<https://osf.io/mgx9p/>).

2.6. Analysis

250 Group and stimulus effects on accuracy and confidence ratings were tested using a $2 \times 4 \times 7$ repeated-measures ANOVA with the factor Emotion category with 7 levels (anger, disgust, fear, happy, neutral, sad, and surprise), the factor group (computer systems vs humans) and the factor covering (unmasked,

masked, partial mask, and sunglasses). Post hoc tests included 2×4 repeated
 255 measures ANOVAs and t-tests.

For all multiple tests such as post hoc tests after the initial overall ANOVAs,
 the alpha level was adjusted using *Bonferroni corrections* [48].

3. Results

This section presents the accuracy ratings obtained from humans compared
 260 with computer systems (i.e. VGG19, ResNet50, and InceptionV3) in categoriz-
 ing emotion from covered and uncovered faces.

3.1. Emotion categorization

Initially, a 3-way repeated-measures ANOVA ($2 \times 4 \times 7$) was performed with
 the between-subject factor Group (i.e. humans vs computer systems), the
 265 within-subject factors Covering (i.e. unmasked, mask, partial mask, and sun-
 glasses), and Emotion (i.e. anger, disgust, fear, happy, neutral, sad, and sur-
 prise) to analyze the achieved accuracy. We found significant main effects for
 Group, $F(1, 3080) = 80.25, p < .001, \eta^2 = 0.01$, Covering, $F(3, 3080) = 40.44,$
 $p < .001, \eta^2 = 0.09$ and Emotion, $F(6, 3080) = 81.87, p < .001, \eta^2 = 0.02$. All
 270 interactions between the factors were significant (see Table 1a), as well as the
 three-way interaction of Group $\times$ Covering $\times$ Emotion $F(18, 3080) = 28.11, p$
 $< .001, \eta^2 = 0.09$.

Thereafter, the alpha level was adjusted using *Bonferroni correction* ($\alpha =$
 0.0083) before post hoc tests. A significant interaction between the emotion
 275 categories (i.e. anger, disgust, fear, happy, neutral, sad, and surprise) was found
 for humans, $F(6, 567) = 33.18, p < .001, \eta^2 = 0.26$ and computer systems, $F(6,$
 $203) = 33.14, p < .001, \eta^2 = 0.50$ for the unmasked emotion (see Fig. 2a).

Fig. 2a represents the violin plots for categorization accuracy of unmasked
 images of people by both humans and computer systems. As can be seen from
 280 Fig. 2a, the computer systems classify almost all the classes correctly, achieving
 an overall accuracy of 98.48% compared to humans who achieved an overall ac-
 curacy of 82.72%. This decrease in the overall performance of humans is based

(a) unmasked

(b) masked

(c) partial mask

(d) sunglasses

Figure 2: Violin plots showing accuracy obtained from categorizing a) unmasked, b) masked, c) partial mask, and d) sunglasses emotions by humans and computer systems (VGG19).

on significantly lower accuracy in categorizing anger, disgust, and fear expressions, with fear contributing the most (Error rate $> 30\%$) to the significantly worst performance (see Table 1b) in humans. However, *Bonferroni corrected* post hoc t-tests showed no significant difference in categorization accuracies achieved by the computer systems across emotion categories except for the neutral expression (see Table 1c), which was categorized with a significantly lower accuracy compared to other emotion categories.

The classification accuracy achieved by the computer systems significantly decreased from an average accuracy of 98.48% when there was no mask to 24.0% when there were masks on the face (see Fig. 2b). Not only did the average accuracy decrease, but also the accuracy achieved from categorizing each class decreased significantly ($> 68\%$), except for anger (less than 16%). Although the overall accuracy achieved by humans for masks on the face also decreases in each category (from when there was no mask), the decrease is not large ($< 26\%$) compared to the decrease seen in the computer systems' accuracy (up to 74.48%). Also, post hoc tests showed that differences between emotion categories were significant for both humans, $F(6, 567) = 64.85$, $p < .001$, $\eta^2 = 0.41$ and computer systems, $F(6, 203) = 125.79$, $p < .001$, $\eta^2 = 0.79$ ($\alpha = 0.0083$) for the masked faces. Post hoc two-sample t-tests for the accuracies in each category obtained from humans and the computer systems showed significant differences ($p < .001$) for most of the emotion categories at $\alpha = 0.0014$ (see Tables 1d and 1e).

The violin plots presented in Fig. 2c show accuracies obtained from categorizing images by humans and computer systems when a partial mask (mask with transparent window) was used to cover the face. In contrast to the use of a fully covering mask, the classification accuracy obtained by computer systems saw a great increase in categorizing happiness and surprise expressions with partial masks (0% vs 66.67% and 20% vs 73.33%). Also, while the accuracy achieved by humans had also increased across all the emotion categories, the largest increase was seen for happy, disgust, sad, and surprise expressions (all $> 19\%$ increase). Post hoc tests revealed significance differences for both humans,

	Sum Square	df	F	<i>p</i>	Eta Squared
Intercept	1.76E+05	1.0	438.05	<.001	0.08
C(Group)	3.23E+04	1.0	80.25	<.001	0.01
C(Covering)	4.88E+04	3.0	40.44	<.001	0.09
C(Emotion)	1.98E+05	6.0	81.87	<.001	0.02
C(Group): C(Covering)	1.78E+04	3.0	14.73	<.001	0.09
C(Group): C(Emotion)	1.94E+05	6.0	80.44	<.001	0.01
C(Covering): C(Emotion)	1.56E+05	18.0	21.55	<.001	0.07
C(Group): C(Covering):C(Emotion)	2.03E+05	18.0	28.11	<.001	0.09
Residual	1.24E+06	3080.0	-	-	-

(a) ANOVA comparing humans and computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.123 (-0.4)	0.341 (0.25)	<.001 (-1.63)	<.001 (-1.00)	<.001 (-0.96)	<.001 (-1.20)
Disgust	0.123 (-0.40)	-	0.025 (0.60)	<.001 (-1.08)	0.049 (-0.52)	0.048 (-0.52)	0.008 (-0.72)
Fear	0.341 (0.25)	0.025 (0.60)	-	<.001 (-1.59)	<.001 (-1.11)	1.00 (-1.08)	1.00 (-1.27)
Happy	<.001 (-1.63)	<.001 (-1.08)	<.001 (-1.59)	-	0.008 (0.71)	0.056 (0.51)	0.113 (0.42)
Neutral	<.001 (-1.00)	0.049 (-0.52)	<.001 (-1.11)	0.008 (0.71)	-	0.832 (-0.06)	0.354 (-0.24)
Sad	<.001 (-0.96)	0.048 (-0.52)	<.001 (-1.08)	0.056 (0.51)	0.832 (-0.06)	-	0.558 (-0.15)
Surprise	<.001 (-1.20)	0.008 (-0.72)	<.001 (-1.27)	0.113 (0.42)	0.354 (-0.24)	0.558 (-0.15)	-

(b) unmasked by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	1.00 (0)	1.00 (0)	1.00 (0)	<.001 (1.51)	1.00 (0)	1.00 (0)
Disgust	1.00 (0)	-	1.00 (0)	1.00 (0)	<.001 (1.51)	1.00 (0)	1.00 (0)
Fear	1.00 (0)	1.00 (0)	-	1.00 (0)	<.001 (1.51)	1.00 (0)	1.00 (0)
Happy	1.00 (0)	1.00 (0)	1.00 (0)	-	<.001 (1.51)	1.00 (0)	1.00 (0)
Neutral	<.001 (1.51)	<.001 (1.51)	<.001 (1.51)	<.001 (1.51)	-	<.001 (1.51)	<.001 (1.51)
Sad	1.00 (0)	1.00 (0)	1.00 (0)	1.00 (0)	<.001 (1.51)	-	1.00 (0)
Surprise	1.00 (0)	1.00 (0)	1.00 (0)	1.00 (0)	<.001 (1.51)	1.00 (0)	-

(c) unmasked by computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.027 (0.59)	0.230 (-0.32)	<.001 (-1.36)	<.001 (-2.00)	0.276 (0.29)	<.001 (-0.91)
Disgust	0.027 (0.59)	-	0.002 (-0.85)	<.001 (-2.07)	<.001 (-2.78)	0.327 (-0.26)	<.001 (-1.67)
Fear	0.230 (-0.32)	0.002 (-0.85)	-	0.002 (-0.85)	<.001 (-1.40)	0.036 (0.56)	0.106 (0.43)
Happy	<.001 (-1.36)	<.001 (-2.07)	0.002 (-0.85)	-	0.014 (-0.66)	<.001 (1.58)	0.020 (0.62)
Neutral	<.001 (-2.00)	<.001 (-2.78)	<.001 (-1.40)	0.014 (-0.66)	-	<.001 (2.18)	<.001 (1.38)
Sad	0.276 (0.29)	0.327 (-0.26)	0.036 (0.56)	<.001 (1.58)	<.001 (2.18)	-	<.001 (-1.17)
Surprise	<.001 (-0.91)	<.001 (-1.67)	0.106 (0.43)	0.020 (0.62)	<.001 (1.38)	<.001 (-1.17)	-

(d) masked by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	<.001 (5.03)	<.001 (2.90)	<.001 (6.05)	<.001 (3.61)	<.001 (4.69)	<.001 (4.19)
Disgust	<.001 (5.03)	-	<.001 (-1.80)	<.001 (1.00)	0.017 (-0.64)	0.072 (-0.48)	<.001 (-1.12)
Fear	<.001 (2.90)	<.001 (-1.80)	-	<.001 (2.63)	0.002 (0.87)	<.001 (1.45)	<.001 (0.94)
Happy	<.001 (6.05)	<.001 (1.00)	<.001 (2.63)	-	<.001 (-1.24)	<.001 (-1.62)	<.001 (-2.36)
Neutral	<.001 (3.61)	0.017 (-0.64)	0.002 (0.87)	<.001 (-1.24)	-	0.230 (0.32)	0.605 (-0.13)
Sad	<.001 (4.69)	0.072 (-0.48)	<.001 (1.45)	<.001 (-1.62)	0.230 (0.32)	-	0.017 (-0.64)
Surprise	<.001 (4.19)	<.001 (-1.12)	<.001 (0.94)	<.001 (-2.36)	0.605 (-0.13)	0.017 (-0.64)	-

(e) masked by computer systems

Table 1: Statistical test results presenting a) ANOVA obtained from comparing humans and computer systems based on their achieved accuracy from categorizing people's emotions with different types of coverings. b) Post hoc *t*-test results showing the *p*-value and Cohen's *d* effect size obtained from comparing unmasked emotion categories by humans and c) computer systems. d) Post hoc *t*-test results showing the *p*-value and Cohen's *d* effect size obtained from comparing masked emotion categories by humans and e) computer systems. Note that the α level in all the post hoc comparisons has been adjusted to 0.0014 using *Bonferroni correction*.

$F(6, 567) = 64.01, p < .001, \eta^2 = 0.40$ and computer systems, $F(6, 203) =$
 315 235.28, $p < .001, \eta^2 = 0.87 (\alpha = 0.0083)$. Post hoc two-sample unpaired t-tests
 showed a significant differences ($< .001$) for most emotion categories (see Tables
 2a and 2b).

The violin plots presented in Fig. 2d show the accuracy obtained from
 categorizing faces with sunglasses by humans and computer systems. As can
 320 be seen, the accuracy achieved by the computer systems decreased most for sad
 and surprise expressions with sunglasses (all $> 27\%$). Notably, the accuracy
 achieved by humans in categorizing happiness remains the same both with or
 without sunglasses. This shows that putting on sunglasses does not reduce
 humans' happiness detection. As expected, there was a large decrease ($> 16\%$)
 325 in the accuracy for fear expressions when the eyes were covered. Post hoc tests
 showed that these differences between emotion categories were significant for
 both humans, $F(6, 567) = 74.15, p < .001, \eta^2 = 0.44$ and computer systems,
 $F(6, 203) = 13.85, p < .001, \eta^2 = 0.29 (\alpha = 0.0083)$ for faces with sunglasses.
 Post hoc two-sample unpaired t-test showed significant differences ($< .001$) for
 330 most of the emotion categories (see Tables 2c and 2d).

We also tested for differences within each emotion category across different
 types of covering (i.e. unmasked, mask, partial mask, and sunglasses). A sig-
 nificant main effect was found for all emotion categories (all $p < .001$), except
 for the neutral expression, $F(3, 324) = 0.037, p = .99, \eta^2 = 0.003$ (see Fig. 3a).
 335 Although, there was a significant, $F(3, 324) = 5.36, p < .001 (p = 0.00129)$,
 $\eta^2 = 0.05$ main effect for humans categorizing the fear expression comparing
 the different types of coverings (see Fig. 3b), the *effect size* was very small
 ($\eta^2 = 0.05$) compared with other emotion categories (all $\eta^2 > 0.2$) like anger
 (see Fig. 3c), etc. that have a larger effect.

340 Thus, humans might be equally worse (see Fig. 3b) in categorizing the fear
 emotion category with or without covering on the face as the *effect size* is not
 as large as it is in other expressions (all $\eta^2 > 0.2$ except neutral expression η^2
 = 0.003).

Based on the three additional questions asked about exposure to face masked

Figure 3: Violin plot showing the accuracy obtained a) from classifying neutral expression by humans across four different coverings, b) from categorizing fear expression by humans across four different coverings, c) from categorizing anger expression by humans across four different coverings, and d) by computer systems from categorizing neutral and anger expression before and after the label changed.

345 people, Q_{mean} was added as a factor to a linear regression and then investi-
 gate if people's experience with masked individuals might help them achieve
 higher accuracy in categorizing mask emotion images. The results (see Table
 2e) showed significant main effects for the factors Covering, $F(3, 6) = 122.19$,
 $p < .001$, $\eta^2 = 0.09$ and Q_{mean} , $F(6, 70) = 3.73$, $p < .001$, $\eta^2 = 0.07$, where
 350 Q_{mean} represents the average of the three questions (higher accuracy related
 to higher Q_{mean} mean more exposure). However, there was no significant in-
 teraction between Covering and Q_{mean} , $F(6, 210) = 0.62$, $p = 1.00$, $\eta^2 = 0.03$.

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.051 (-0.52)	0.015 (-0.65)	<.001 (-2.49)	<.001 (-2.02)	<.001 (-1.13)	<.001 (-2.24)
Disgust	0.051 (-0.52)	-	0.695 (-0.10)	<.001 (-1.66)	<.001 (-1.29)	0.047 (-0.53)	<.001 (-1.49)
Fear	0.015 (-0.65)	0.695 (-0.10)	-	<.001 (-1.66)	<.001 (-1.25)	0.096 (-0.44)	<.001 (-1.46)
Happy	<.001 (-2.49)	<.001 (-1.66)	<.001 (-1.66)	-	0.112 (0.42)	<.001 (1.26)	0.621 (0.13)
Neutral	<.001 (-2.02)	<.001 (-1.29)	<.001 (-1.25)	0.112 (0.42)	-	0.002 (0.83)	0.330 (-0.26)
Sad	<.001 (-1.13)	0.047 (-0.53)	0.096 (-0.44)	<.001 (1.26)	0.002 (0.83)	-	<.001 (-1.06)
Surprise	<.001 (-2.24)	<.001 (-1.49)	<.001 (-1.46)	0.621 (0.13)	0.330 (-0.26)	<.001 (-1.06)	-

(a) partial mask by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	<.001 (7.29)	<.001 (6.11)	<.001 (-2.35)	0.566 (0.15)	<.001 (8.22)	<.001 (-2.16)
Disgust	<.001 (7.29)	-	0.310 (-0.27)	<.001 (-7.49)	<.001 (-3.30)	0.319 (0.26)	<.001 (-5.56)
Fear	<.001 (6.11)	0.310 (-0.27)	-	<.001 (-6.85)	<.001 (-3.07)	0.078 (0.47)	<.001 (-5.28)
Happy	<.001 (-2.35)	<.001 (-7.49)	<.001 (-6.85)	-	<.001 (1.72)	<.001 (7.91)	0.104 (-0.43)
Neutral	0.566 (0.15)	<.001 (-3.30)	<.001 (-3.07)	<.001 (1.72)	-	<.001 (3.43)	<.001 (-1.82)
Sad	<.001 (8.22)	0.319 (0.26)	0.078 (0.47)	<.001 (7.91)	<.001 (3.43)	-	<.001 (-5.72)
Surprise	<.001 (-2.16)	<.001 (-5.56)	<.001 (-5.28)	0.104 (-0.43)	<.001 (-1.82)	<.001 (-5.72)	-

(b) partial mask by computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.274 (0.29)	0.002 (0.86)	<.001 (-1.64)	<.001 (-0.95)	0.024 (0.60)	<.001 (-1.30)
Disgust	0.274 (0.29)	-	0.031 (0.57)	<.001 (-1.87)	<.001 (-1.21)	0.292 (0.28)	<.001 (-1.55)
Fear	0.002 (0.86)	0.031 (0.57)	-	<.001 (-2.34)	<.001 (-1.73)	0.169 (-0.36)	<.001 (-2.05)
Happy	<.001 (-1.64)	<.001 (-1.87)	<.001 (-2.34)	-	0.012 (0.68)	<.001 (2.57)	0.158 (0.37)
Neutral	<.001 (-0.95)	<.001 (-1.21)	<.001 (-1.73)	0.012 (0.68)	-	<.001 (1.68)	0.207 (-0.33)
Sad	0.024 (0.60)	0.292 (0.28)	0.169 (-0.36)	<.001 (2.57)	<.001 (1.68)	-	<.001 (-2.12)
Surprise	<.001 (-1.30)	<.001 (-1.55)	<.001 (-2.05)	0.158 (0.37)	0.207 (-0.33)	<.001 (-2.12)	-

(c) sunglasses by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.570 (0.15)	0.244 (0.31)	0.576 (-0.15)	<.001 (1.27)	<.001 (1.72)	<.001 (1.49)
Disgust	0.570 (0.15)	-	0.566 (0.15)	0.354 (-0.24)	<.001 (1.12)	<.001 (1.22)	<.001 (1.24)
Fear	0.244 (0.31)	0.566 (0.15)	-	0.158 (-0.38)	<.001 (0.99)	<.001 (0.94)	<.001 (1.05)
Happy	0.576 (-0.15)	0.354 (-0.24)	0.158 (-0.38)	-	<.001 (1.27)	<.001 (1.47)	<.001 (1.42)
Neutral	<.001 (1.27)	<.001 (1.12)	<.001 (0.99)	<.001 (1.27)	-	0.084 (-0.46)	0.550 (-0.16)
Sad	<.001 (1.72)	<.001 (1.22)	<.001 (0.94)	<.001 (1.47)	0.084 (-0.46)	-	0.156 (0.38)
Surprise	<.001 (1.49)	<.001 (1.24)	<.001 (1.05)	<.001 (1.42)	0.550 (-0.16)	0.156 (0.38)	-

(d) sunglasses by computer systems

	Sum Square	df	F	p	Eta Squared
C(Emotion)	5.28E+05	6.0	172.42	<.001	0.27
C(Covering)	1.87E+05	3.0	122.19	<.001	0.09
C(Qmean)	1.33+E05	70.0	3.73	<.001	0.07
C(1/ID)	5.13+E04	81.0	1.24	0.076	0.03
C(Covering):C(Qmean)	6.62E+04	210.0	0.62	1.00	0.03
Residual	1.02+E06	1995.0	-	-	-

(e) ANOVA for human sample including Qmean

Table 2: Statistical test results presenting a) Post hoc test (t -test) results showing the p -value and Cohen's d effect size obtained from comparing partial mask emotion categories by humans and b) computer systems. c) Post hoc test (t -test) results showing the p -value (and Cohen's d effect size) obtained from comparing sunglasses emotion categories by humans and d) computer systems. e) Regression results from comparing humans based on their understanding of masked people and their achieved accuracy from categorizing people's emotions with different coverings. Note that the α level in all the post hoc comparisons has been adjusted to 0.0014 using *Bonferroni correction*.

3.2. Misclassification Analysis

As an additional analysis, this research investigated the types of misclassification made by humans and computer systems.

The results (see Section 3.1) showed that humans mainly misclassify expressions with full covering (see Fig. 4a) and therefore, unclear expressions as neutral. However, the misclassifications varied for the computer systems. For instance, the computer systems mainly classified full face mask covered expressions as anger, while partially covered facial expressions were mainly classified as happiness or surprise (see Fig. 4d).

As such, misclassifications can result in costly negative outcomes and complications for humans interacting with computer systems. Therefore, this work aimed to regulate the cost of misclassification in these computer systems. To make them resemble human classification behaviour, we aimed to alter the training methods to also lead to misclassifying unclear expressions as neutral.

A possible explanation of why humans mainly misclassify unclear emotions as neutral is because it is the emotion that they are exposed to predominantly on a day-to-day basis. In other words, humans would classify unclear expressions as neutral because it is the expression they see on people’s faces most of the time. Therefore, by oversampling the neutral expression images used to train the computer systems, we anticipated that the models would behave more similar to humans as they see more neutral expressions compared to other emotion categories.

3.2.1. Method

Since classifying any emotion category as a neutral expression is considered less costly (see Section 3.2) compared to when they are classified as other emotional categories such as anger, happiness, sadness, etc., the misclassification cost of the model was evaluated without taking the samples that were misclassified as neutral into account. In particular, the quality of the classifier is not measured based on total misclassification but rather using equation 1 below:

$$cost = \frac{Err - Err_{neutral}}{No. test samples} \quad (1)$$

where

- Err represents the misclassified samples or the total misclassification
- $Err_{neutral}$ represents samples misclassified as neutral
- $No. test samples$ represents the total number of test sample images

3.2.2. Results

In Table 3, the column *oversampling ratio* represents the oversampling ratio of the category neutral against all other emotion categories. For instance, the value “2:2” means no oversampling whereas the ratio “4:2” means there were twice as much neutral expression images in the training set compared to images from all other emotion categories. Different ratios were chosen based on the number of available images in the CK+ dataset so as to find the best ratio that would give a much lower misclassification cost while maintaining good accuracy.

The results here are presented with the three computer systems (i.e. VGG19, ResNet50, and InceptionV3) to test generalizability. However, a detailed explanation of only the VGG19 model is provided as the model achieved higher performance compared with the other models (see Table 3).

As can be seen from Table 3, the average misclassification cost obtained in each of the oversampled cases is less than the average cost without oversampling. Thus, oversampling reduced the cost of misclassification, with the least misclassification cost for the models trained with oversampled data at a 5:2 ratio.

Fig. 4a, 4b, and 4c present the confusion matrices² obtained from classifying mask, partial mask, and sunglasses images by humans, Fig. 4d, 4e, and

²A confusion matrix is a table to visualize the performance of an algorithm. In a confusion matrix, on the basis whereby each row represents the number of instances in a predicted class, the columns will represent the total number of instances in an actual class (or vice-versa).

Figure 4: Confusion matrices obtained from classifying a) mask, b) partial mask, and c) sunglasses images by humans, from classifying d) mask, e) partial mask, and f) sunglasses images by computer systems without oversampling, from classifying g) mask, h) partial mask, and i) sunglasses images by computer systems with oversampling [5:2]. Note that the 82 human participants all classified a sum of 35 images for each covering, i.e. a total of 410 images for each emotion category and 2870 for each covering. Similarly, 30 computer systems (i.e. VGG19) all classified a sum of 35 images for each covering, i.e. a total of 150 images for each emotion category and 1050 for each covering. *An* = anger, *Di* = disgust, *Fe* = fear, *Ha* = happy, *Ne* = neutral, *Sa* = sadness and *Su* = surprise.

Table 3: Table presenting accuracies and the misclassification costs obtained from classifying covered images with and without oversampling neutral expression images.

Oversample ratio (Neutral:All)	Misclassification cost			Avg cost	Accuracy			Avg accuracy	Model
	Mask	Partial mask	Sunglasses		Mask	Partial mask	Sunglasses		
2:2	0.61	0.47	0.18	0.41	0.24	0.33	0.80	0.46	VGG19
3:2	0.40	0.54	0.23	0.39	0.20	0.14	0.68	0.34	
4:2	0.26	0.34	0.17	0.26	0.17	0.23	0.71	0.37	
5:2	0.25	0.39	0.10	0.25	0.25	0.17	0.66	0.36	
2:2	0.41	0.35	0.72	0.49	0.22	0.20	0.23	0.22	ResNet50
3:2	0.33	0.31	0.65	0.43	0.20	0.19	0.22	0.20	
4:2	0.30	0.24	0.56	0.37	0.18	0.15	0.20	0.18	
5:2	0.23	0.19	0.50	0.31	0.16	0.16	0.21	0.18	
2:2	0.69	0.40	0.62	0.57	0.22	0.23	0.34	0.26	InceptionV3
3:2	0.63	0.24	0.64	0.51	0.22	0.19	0.27	0.23	
4:2	0.59	0.21	0.51	0.44	0.23	0.19	0.32	0.25	
5:2	0.54	0.18	0.41	0.38	0.21	0.18	0.38	0.26	

4f present the confusion matrices obtained from classifying mask, partial mask, and sunglasses images when no oversampling was applied, and Fig. 4g, 4h, and 4i present the confusion matrices obtained from classifying mask, partial mask, and sunglasses images with the best oversampling ratio (i.e. 5:2) applied to the training images (of VGG19). Humans mainly misclassified unclear expressions as neutral when the face was fully mask covered. However, while the misclassification varied for different coverings when no oversampling was applied to computer systems, after oversampling, the majority of the misclassified images were misclassified as neutral. Although the average accuracy achieved by the model without oversampling is larger ($> 10\%$) when compared with the average accuracy achieved by the model with the best oversampling ratio (5:2), the difference in the average misclassification cost achieved by the model with the best oversampling ratio was significantly lower (16% lower) than the average misclassification cost of the model with no oversampling, ($t(58) = 30.98$, $p < .001$).

4. Discussion

The work aimed to compare the performance of humans and computer systems in categorizing emotion from faces of people with sunglasses and face

masks, as well as compare the decrease caused by each covering. The work also introduces an approach that reduces the misclassification cost for computer
 425 systems.

In addition to the replication of previous findings that computer systems (with 98.48% accuracy) perform better than humans (with 82.72% accuracy) in categorizing emotions from uncovered/unmasked images in the same dataset [49]–[52], we found that face masks (both partially covered and fully covering
 430 masks) and sunglasses reduced the accuracy of both humans and computer systems in emotion categorization. Consistent with Carbon [1], we also found that humans mainly classify faces with fully covering masks as neutral.

While the face masks impair emotion categorization in both humans and computer systems by a significant amount, emotion categories are affected differently depending on the type of covering used. For instance, the use of a
 435 fully covering mask mainly obscured disgust, happiness, and sadness for computer systems whereas anger, disgust, and sadness were mainly more difficult to categorize for humans, which is in line with previous experimental work using stimuli that varied in missing features [6], [53]. The partial masks mainly im-
 440 paired anger categorization in humans, whereas computer systems were mainly impaired through those for disgust, fear, and sadness. And while the use of sunglasses mainly reduced fear categorization ability in humans, it was neutral and surprise for computer systems.

Previous work [20], [54] has shown the relevance of the mouth for correctly
 445 categorizing happiness and surprise in both humans and computer systems. Therefore, since the partial masks allow the mouth to be visible through the transparent window, it was anticipated that happiness and surprise would be categorized more accurately when the partial mask is applied. Indeed, we found a large increase in the classification accuracy of happiness and surprise expres-
 450 sion in both computer systems and humans. While there was an increase of 19% (for happy) and 28% (for the surprise) in the human sample, there was an increase of up to 66% for the category happy and up to 55% for surprise expression in computer systems when the partial mask was used.

While we found a significant difference in the accuracy of humans categorizing each emotion category across the different coverings, no significant difference
 455 was found for the neutral expression across all coverings. As humans classify unclear emotions as neutral, they subsequently were highly accurate in categorizing neutral expressions (see Fig. 3a).

In contrast to humans, the classification of unclear emotions significantly
 460 varies in computer systems. For instance, the computer systems mainly classify full mask covered expressions as anger. However, unclear emotions were mainly classified as happiness or surprise when there was a partial mask on the face. An initial observation of the confusion matrix obtained from categorizing masked emotion images (see Fig. 4d) suggested that the computer systems mainly
 465 classify unclear emotions as anger expressions just because the anger expression appears first after *one-hot encoding*³ the labels that were fed to the classifier. Therefore, to test this hypothesis, we changed the label name of the neutral expression to ‘*anneutral*’ such that the neutral expression appears first after the labels are one-hot encoded. The violin plots in Fig. 3d show the accuracy
 470 obtained from neutral and anger expressions before and after the labels were changed. As can be seen from Fig. 3d, the accuracy obtained before and after the label change was nearly the same or relatively similar for both anger and neutral expression. In addition, no significant difference was found in the achieved accuracies before and after the label was changed. Thus, label order
 475 was not important.

Although the achieved accuracy by the computer systems is relatively low when the fully covering masks were used, the above-mentioned observation suggests that the prediction of the computer systems was not fully random. One possible reason why the images of the fully covering masks were mainly classified
 480 as anger could be because the produced (extracted) features⁴ by the computer

³One hot encoding is a process whereby categorical variables are converted into a form that can be provided to machine learning algorithms.

⁴Feature extraction is the process by which an initial set of raw data (in this case image

systems when the fully covering masks were used have high similarities with the anger expression images.

Depending on the type of covering, the classification accuracy in humans changes. There is no general order of increase or decrease in the achieved accuracy. However, based on the results obtained in this study, the classification
 485 accuracy of the computer systems correlates with the occlusion. The larger the occlusion, the higher the decrease and vice-versa. For instance, in the case of VGG19, the highest accuracy was achieved with nothing added to the image. The achieved accuracy decreases (by up to 18%) when sunglasses (covers 10-
 490 15% of the face) were used. And although a further decrease (up to 65%) has been observed when partial masks (covers 30-35% of the face) were used, the largest decrease (up to 74.48%) was experienced when the fully covering (covers 60-65% of the face) face masks were used. It is important to bear in mind that this observation is only valid for the type of coverings used in this research as it
 495 is unknown whether the decrease is based on the percentage of the face that is covered or based on important features that are covered. This is an important question for future research.

We also found that people’s exposure to people wearing masks (see Section 3.1) does not affect their performance in categorizing emotion. Although more
 500 exposure to people wearing mask was related to increased accuracy across all facial expressions, we found no evidence for people’s exposure to masked individuals helping them to be specifically better in categorizing masked emotions (see Table 2e). Taken together, these results suggest that people perform worst in categorizing masked emotion images independent of their exposure to masked
 505 individuals. In other words, being around more masked people does not increase the ability to categorize partially covered emotional facial expressions.

We trained the computer systems with the last-half frames of images from the CK+ dataset. This means that the computer systems have seen everyone in the dataset (even though testing was performed only on unseen images) whereas

pixels) is reduced to subset of features for processing.

human participants had not seen any of those images before. On the other hand, human beings have seen a number of different expressions of people from the time they were born that the computer systems have not seen. Nevertheless, this is on par with recent studies conducted to perform a similar type of comparison as it seems not feasible to train the computer systems with all images the human participants have seen in the past or the other way around.

Prior work [26], [27] conducted to tackle the problem of imbalanced training at the data-level⁵ suggests that it is crucial to balance the distribution of data within the train set in order to have a reduced misclassification cost, as well as improve classification accuracy. However, the findings based on the misclassification analysis suggest that it might be important to feed imbalanced data to classifiers. For instance, in situations like this where using balanced data might lead to differences in misclassifications, feeding the classifier with data that is oversampled based on a certain expression with less cost of incorrect identification (see Section 3.2.1) is more important than using data with an equal distribution from each expression as it results in a high cost of misclassification.

Not only did oversampling reduce the misclassification cost of the covered images when a specified class is targeted, at the same time, it also increased the classification accuracy when classifying the fully covering mask images (in the case of VGG19). The accuracy achieved by the oversampling method (ratio 5:2) was greater than the accuracy achieved when no oversampling was used. Nevertheless, oversampling the data might also come at a cost of decreasing the classification accuracy, e.g. the accuracy obtained from classifying the partial mask and the sunglasses images were reduced after the neutral expression was oversampled for all the three (VGG19, ResNet50, and InceptionV3) algorithms.

This gives rise to the question of which performance measure to prefer? Is the accuracy or the misclassification cost more relevant, taking into account the

⁵The data-level way of handling imbalanced data tries to balance the training set by increasing the number of samples from smaller classes through the addition of new samples and reducing larger classes.

protective factor, mediating, and moderating effects [55], [56] of these measures? While there is no general answer to this question, both the accuracy and the misclassification cost need to be considered. For instance, misclassifying anger or
 540 sadness as happiness could have extreme consequences in future decision making for artificial intelligence systems like robots. At the same time, an organization will more likely introduce robots that can accurately classify people’s emotions into shared workspaces compared to less accurate ones.

It is also worth stating that the finding that oversampling reduced the cost
 545 of misclassification cannot be generalized to all the different types of face coverings. For instance, the misclassification cost with no oversampling for faces with sunglasses was lower (18%) compared to when oversampling (ratio 3:2) was used (23%).

In general, given a set of image features, the computer systems are more
 550 likely to predict the emotion category of an image as the same, while humans seem not to show that tendency. For instance, in Section 3.1, without coverings on the face, the computer systems (from all 30 runs) categorized six out of seven emotion categories (except neutral) correctly. However, there is a high variation in the accuracy of humans across emotion categories. An image that is seen to
 555 represent a particular emotion by one person might be seen to have a different emotion by another person. This is due to the different perspectives that we all have, and based on the way our brain processes information and mediates actions [57].

Several reports have shown that machine learning classifiers are mainly used
 560 nowadays to produce emotion categorization devices [45], [58]–[60]. However, the findings of this research provide insights that these classifiers do not perform well in categorizing emotion when the face is partially covered. Therefore, these types of classifiers should be implemented carefully in contexts where there are coverings on people’s faces, especially with the recent challenge of the coronavirus. A possible way to improve the performance of these models is to train
 565 them with facial images with coverings.

It is also important to note that the overall findings of what is referred to as

computer systems in this research may be somewhat limited to end-to-end deep learning models as different machine learning techniques might lead to achieving
 570 different results depending on the feature extraction technique, as well as the algorithm used to categorize the emotion.

Across all images, humans spent approximately 0.14 mins to predict the category of an image and 19.77 mins for the complete survey, while it took the computer systems 0.00012 mins (less than 1 sec) to classify each image and
 575 approximately 0.017 mins (1 sec) to classify all 140 images. Nevertheless, the computer systems took approximately 80 mins to train on a 24GB Graphical Processing Unit (GPU) machine with Quadro RTX 6000 CUDA version 11.2. Thus, computer systems can predict the expression of the images faster than humans provided they are already trained. Humans train throughout their
 580 lifetime, although they are already ‘experts’ at recognizing emotions at an early age [61], [62].

5. Conclusion

This study compared the performance of humans and computer systems in categorizing emotion from faces with different coverings. The results obtained
 585 from the study showed that while the computer systems (98.48%) outperformed humans (82.78%) with no covering on the face, the performance of the computer systems showed a larger decrease (up to 74.48% decrease) in comparison to humans (with 25.29% decrease) when the face is partially covered. Further, as these computer systems make different mistakes in comparison to humans, this
 590 work introduced a method to reduce the cost of misclassification for computer systems in order to avoid having extreme consequences in future decision making. This highlights how partial face coverings affect the performance of humans and computer systems in classifying emotions. The findings contribute to our understanding of how computer systems make mistakes, as well as how the
 595 introduced approach reduces the misclassification cost for computer systems. Together, these findings suggest that greater efforts are needed to develop arti-

ificial emotion classification systems that can adapt to more complex situations, e.g. taking into account contextual features. Thus, care is needed prior to using results from computer systems for emotion recognition in soft computing experiments.

The study also has some potential limitations. For instance, the overall findings of what is referred to as computer systems in this research may be somewhat limited to end-to-end deep learning models as different machine learning techniques might lead to achieving different results depending on the feature extraction technique, as well as the algorithm used to categorize the emotion. In addition, the study was performed on a single stimulus set, therefore, the question of whether the findings would generalize to other stimulus sets remains open. Future studies should address this issue by including other stimulus sets, as well as other artificial emotion classification methods.

610

Acknowledgements

Authors would like to acknowledge the participants for putting in the effort to complete the survey.

6. Declarations

Funding

This research was not supported by any organization.

Competing interests

The authors declare no competing interests.

620

Ethics Approval

Ethics approval has been obtained from the Human Ethics Committee (HEC) of Victoria University of Wellington (Application no: 28949).

Consent to participate

All participants gave informed consent for participation in the study.

Consent for publication

All authors have approved the manuscript and agree with submission to Applied
 630 Soft Computing journal

Data availability

Data used in this study can be downloaded from OSF (<https://osf.io/mgx9p>)
 and Code Ocean (<https://codeocean.com/capsule/2976335/tree>)

635

Code availability

Code implemented in this study can be downloaded from OSF (<https://osf.io/mgx9p>) and Code Ocean (<https://codeocean.com/capsule/2976335/tree>)

Authors Contribution

Harisu Abdullahi Shehu: Conceptualization, Methodology, Validation, Formal
 Analysis, Investigation, Data Curation, Writing – Original Draft, Visualization,
 Project administration; Will Browne: Supervision, Writing – Review & Editing;
 Hedwig Eisenbarth: Supervision, Writing – Review & Editing

645

References

- [1] C.-C. Carbon, “Wearing face masks strongly confuses counterparts in
 reading emotions,” *Frontiers in Psychology*, vol. 11, p. 2526, 2020.
- [2] D. Roberson, M. Kikutani, P. Döge, L. Whitaker, and A. Majid, “Shades of
 650 emotion: What the addition of sunglasses or masks to faces reveals about
 the development of facial expression processing,” *Cognition*, vol. 125, no. 2,
 pp. 195–206, 2012.
- [3] L. F. Barrett, “How emotions are made: The secret life of the brain.,”
 2017.

- 655 [4] P. Lewinski, T. M. den Uyl, and C. Butler, “Automated facial coding: Validation of basic emotions and faces aus in facereader.,” *Journal of Neuroscience, Psychology, and Economics*, vol. 7, no. 4, p. 227, 2014.
- [5] E. Noyes, J. P. Davis, N. Petrov, K. L. Gray, and K. L. Ritchie, “The effect of face masks and sunglasses on identity and expression recognition with super-recognizers and typical observers,” *Royal Society Open Science*, 660 vol. 8, no. 3, p. 201 169, 2021.
- [6] M. Wegrzyn, M. Vogt, B. Kireclioglu, J. Schneider, and J. Kissler, “Mapping the emotional face. how individual face parts contribute to successful emotion recognition,” *PloS one*, vol. 12, no. 5, e0177239, 2017.
- 665 [7] M. Yuki, W. W. Maddux, and T. Masuda, “Are the windows to the soul the same in the east and west? Cultural differences in using the eyes and mouth as cues to recognize emotions in Japan and the United States,” *Journal of Experimental Social Psychology*, vol. 43, no. 2, pp. 303–311, 2007.
- 670 [8] E. J. Miller, E. G. Krumhuber, and A. Dawel, “Observers perceive the duchenne marker as signaling only intensity for sad expressions, not genuine emotion.,” *Emotion*, 2020.
- [9] L. A. Camras and K. Allison, “Children’s understanding of emotional facial expressions and verbal labels,” *Journal of nonverbal Behavior*, vol. 9, 675 no. 2, pp. 84–94, 1985.
- [10] M. Guarnera, P. Magnano, M. Pellerone, M. I. Cascio, V. Squatrito, and S. L. Buccheri, “Facial expressions and the ability to recognize emotions from the eyes or mouth: A comparison among old adults, young adults, and children,” *The Journal of genetic psychology*, vol. 179, no. 5, pp. 297– 680 310, 2018.
- [11] J. Wacker, M. Heldmann, and G. Stemmler, “Separating emotion and motivational direction in fear and anger: Effects on frontal asymmetry.,” *Emotion*, vol. 3, no. 2, p. 167, 2003.

- [12] M. G. Calvo and D. Lundqvist, “Facial expressions of emotion (kdef):
685 Identification under different display-duration conditions,” *Behavior re-
search methods*, vol. 40, no. 1, pp. 109–115, 2008.
- [13] K. Hugenberg, “Social categorization and the perception of facial affect:
Target race moderates the response latency advantage for happy faces,”
Emotion, vol. 5, no. 3, p. 267, 2005.
- [14] J. Van den Stock, R. Righart, and B. De Gelder, “Body expressions in-
690 fluence recognition of emotions in the face and voice,” *Emotion*, vol. 7,
no. 3, p. 487, 2007.
- [15] M. Marini, A. Ansani, F. Paglieri, F. Caruana, and M. Viola, “The impact
of facemasks on emotion recognition, trust attribution and re-identification,”
695 *Scientific Reports*, vol. 11, no. 1, pp. 1–14, 2021.
- [16] F. Grundmann, K. Epstude, and S. Scheibe, “Face masks reduce emotion-
recognition accuracy and perceived closeness,” *PloS one*, vol. 16, no. 4,
e0249792, 2021.
- [17] H. A. Shehu, W. N. Browne, and H. Eisenbarth, “An out-of-distribution
700 attack resistance approach to emotion categorization,” *IEEE Transactions
on Artificial Intelligence*, vol. 2, no. 6, pp. 564–573, 2021.
- [18] R. W. Picard, E. Vyzas, and J. Healey, “Toward machine emotional in-
telligence: Analysis of affective physiological state,” *IEEE transactions on
pattern analysis and machine intelligence*, vol. 23, no. 10, pp. 1175–1191,
705 2001.
- [19] H. A. Shehu, W. N. Browne, and H. Eisenbarth, “An adversarial attacks
resistance-based approach to emotion recognition from images using facial
landmarks,” in *2020 29th IEEE International Conference on Robot and
Human Interactive Communication (RO-MAN)*, IEEE, 2020, pp. 1307–
710 1314.

- [20] H. A. Shehu, A. Siddique, W. N. Browne, and H. Eisenbarth, "Lateralized approach for robustness against attacks in emotion categorization from images," in *24th International Conference on Applications of Evolutionary Computation (EvoApplications 2021)*, Springer, 2021.
- 715 [21] H. A. Shehu, W. N. Browne, and H. Eisenbarth, "An anti-attack method for emotion categorization from images," *Applied Soft Computing*, vol. 128, p. 109 456, 2022.
- [22] —, "Emotion categorization from faces of people with sunglasses and facemasks," Springer, 2021. DOI: 10.21203/rs.3.rs-375319/v1.
- 720 [23] T. J. Coleman, "Coronavirus: Call for clear face masks to be 'the norm'," *BBC News*. Accessed: 26 May, 2020, 2020. [Online]. Available: <https://www.bbc.com/news/world-52764355>.
- [24] N. Japkowicz *et al.*, "Learning from imbalanced data sets: A comparison of various strategies," in *AAAI workshop on learning from imbalanced data sets*, AAAI Press Menlo Park, CA, vol. 68, 2000, pp. 10–15.
- 725 [25] J. W. Grzymala-Busse, "Rule induction," in *Data mining and knowledge discovery handbook*, Springer, 2005, pp. 277–294.
- [26] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: Synthetic minority over-sampling technique," *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- 730 [27] E. Ramentol, Y. Caballero, R. Bello, and F. Herrera, "Smote-rs b*: A hybrid preprocessing approach based on oversampling and undersampling for high imbalanced data-sets using smote and rough sets theory," *Knowledge and information systems*, vol. 33, no. 2, pp. 245–265, 2012.
- 735 [28] M. Pazzani, C. Merz, P. Murphy, K. Ali, T. Hume, and C. Brunk, "Reducing misclassification costs," in *Machine Learning Proceedings 1994*, Elsevier, 1994, pp. 217–225.

- [29] M. Kukar, I. Kononenko, C. Grošelj, K. Kralj, and J. Fettich, “Analysing and improving the diagnosis of ischaemic heart disease with machine learning,” *Artificial intelligence in medicine*, vol. 16, no. 1, pp. 25–50, 1999.
- [30] M. Ciraco, M. Rogalewski, and G. Weiss, “Improving classifier utility by altering the misclassification cost ratio,” in *Proceedings of the 1st international workshop on Utility-based data mining*, 2005, pp. 46–52.
- [31] J. Berryhill, K. K. Heang, R. Clogher, and K. McBride, “Hello, world: Artificial intelligence and its use in the public sector,” 2019.
- [32] H. Eisenbarth, G. W. Alpers, D. Segrè, A. Calogero, and A. Angrilli, “Categorization and evaluation of emotional faces in psychopathic women,” *Psychiatry research*, vol. 159, no. 1-2, pp. 189–195, 2008.
- [33] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- [35] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [36] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Advances in neural information processing systems*, vol. 25, 2012.
- [37] S. Cheng and G. Zhou, “Facial expression recognition method based on improved vgg convolutional neural network,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 34, no. 07, p. 2056003, 2020.

- [38] K. Mohan, A. Seal, O. Krejcar, and A. Yazidi, “Fer-net: Facial expression recognition using deep neural net,” *Neural Computing and Applications*, pp. 1–12, 2021.
- [39] “Keras api,” 2021. [Online]. Available: <https://keras.io/>.
- [40] H. A. Shehu, A. Rabie, and M. H. Sharif, “Artificial intelligence tools and their capabilities,” *PLOMS AI*, vol. 1, no. 1, 2021.
- [41] H. A. Shehu, W. N. Browne, and H. Eisenbarth, “Emotion categorization from video-frame images using a novel sequential voting technique,” in *International Symposium on Visual Computing*, Springer, 2020, pp. 618–632.
- [42] P. Ekman and W. V. Friesen, “Constants across cultures in the face and emotion,” *Journal of personality and social psychology*, vol. 7, no. 2, p. 124, 1971. DOI: 10.1109/AFGR.2000.840611.
- [43] A. Mollahosseini, D. Chan, and M. H. Mahoor, “Going deeper in facial expression recognition using deep neural networks,” in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016, pp. 1–10. DOI: 10.1109/WACV.2016.7477450.
- [44] M. Li, H. Xu, X. Huang, Z. Song, X. Liu, and X. Li, “Facial expression recognition with identity and emotion joint learning,” *IEEE Transactions on Affective Computing*, pp. 1–1, 2018. DOI: 10.1109/TAFFC.2018.2880201.
- [45] T. Zhang, W. Zheng, Z. Cui, Y. Zong, and Y. Li, “Spatial-temporal recurrent neural network for emotion recognition,” *IEEE Transactions on Cybernetics*, vol. 49, no. 3, pp. 839–847, 2019. DOI: 10.1109/TCYB.2017.2788081.
- [46] H. He and S. Chen, “Identification of facial expression using a multiple impression feedback recognition model,” *Applied Soft Computing*, vol. 113, p. 107930, 2021.

- [47] M. N. Islam, M. Seera, and C. K. Loo, “A robust incremental clustering-based facial feature tracking,” *Applied Soft Computing*, vol. 53, pp. 34–44, 2017.
- [48] R. J. Cabin and R. J. Mitchell, “To bonferroni or not to bonferroni: When and how are the questions,” *Bulletin of the ecological society of America*, vol. 81, no. 3, pp. 246–248, 2000.
- [49] D. Matsumoto and H. S. Hwang, “Evidence for training the ability to read microexpressions of emotion,” *Motivation and emotion*, vol. 35, no. 2, pp. 181–191, 2011.
- [50] E. G. Krumhuber, D. Küster, S. Namba, D. Shah, and M. G. Calvo, “Emotion recognition from posed and spontaneous dynamic expressions: Human observers versus machine analysis,” *Emotion*, 2019.
- [51] L. F. Barrett, R. Adolphs, S. Marsella, A. M. Martinez, and S. D. Pollak, “Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements,” *Psychological science in the public interest*, vol. 20, no. 1, pp. 1–68, 2019.
- [52] E. G. Krumhuber, D. Küster, S. Namba, and L. Skora, “Human and machine validation of 14 databases of dynamic facial expressions,” *Behavior research methods*, pp. 1–16, 2020.
- [53] M. Hausmann, B. R. Innes, Y. K. Birch, and R. W. Kentridge, “Laterality and (in) visibility in emotional face perception: Manipulations in spatial frequency content,” *Emotion*, vol. 21, no. 1, p. 175, 2021.
- [54] L. S. Delicato and R. Mason, “Happiness is in the mouth of the beholder and fear in the eyes,” *Journal of vision*, vol. 15, no. 1378, 2015.
- [55] R. Rogers, “The uncritical acceptance of risk assessment in forensic practice,” *Law and human behavior*, vol. 24, no. 5, pp. 595–605, 2000.
- [56] J. Brown and J. P. Singh, “Forensic risk assessment: A beginner’s guide,” *Archives of Forensic Psychology*, vol. 1, no. 1, pp. 49–59, 2014.

- [57] A. Chakravarti, “Perspectives on human variation through the lens of diversity and race,” *Cold Spring Harbor perspectives in biology*, vol. 7, no. 9, a023358, 2015.
- [58] W. M. Alenazy and A. S. Alqahtani, “Gravitational search algorithm
825 based optimized deep learning model with diverse set of features for facial expression recognition,” *Journal of Ambient Intelligence and Humanized Computing*, vol. 12, no. 2, pp. 1631–1646, 2021.
- [59] D. Y. Choi and B. C. Song, “Semi-supervised learning for continuous emotion recognition based on metric learning,” *IEEE Access*, vol. 8, pp. 113 443–
830 113 455, 2020. DOI: 10.1109/ACCESS.2020.3003125.
- [60] W. Hayale, P. Negi, and M. Mahoor, “Facial expression recognition using deep siamese neural networks with a supervised loss function,” in *2019 14th IEEE International Conference on Automatic Face Gesture Recognition (FG 2019)*, 2019, pp. 1–7. DOI: 10.1109/FG.2019.8756571.
- [61] S. D. Pollak and D. J. Kistler, “Early experience is associated with the
835 development of categorical representations for facial expressions of emotion,” *Proceedings of the National Academy of Sciences*, vol. 99, no. 13, pp. 9072–9076, 2002.
- [62] J. Decety, “The neurodevelopment of empathy in humans,” *Developmental
840 neuroscience*, vol. 32, no. 4, pp. 257–267, 2010.
- [63] H. A. Shehu, W. N. Browne, and H. Eisenbarth, “Attention-based methods for emotion categorization from partially covered faces,” *IEEE Transactions on Emerging Topics in Computational Intelligence (submitted, under review)*, 2022.
- [64] H. Jiang, K. Huang, T. Mu, R. Zhang, T. Ting, and C. Wang, “Robust one-
845 shot facial expression recognition with sunglasses,” *International Journal of Machine Learning and Computing*, vol. 6, no. 2, p. 80, 2016.

7. Supplementary Materials

7.1. Confidence ratings

850 With regard to the confidence, just like for the emotion categories, a 3-way repeated-measures ANOVA ($2 \times 4 \times 7$) was performed with the between-subject factor Group (i.e. humans vs computer systems), the within-subject factors Covering (i.e. unmasked, mask, partial mask, and sunglasses), and Confidence obtained from categorizing each emotion (i.e. for anger, disgust, fear, happy, 855 neutral, sad, and surprise). We found significant main effects for Group, $F(1, 3080) = 139.85, p < .001, \eta^2 = 0.03$, Covering, $F(3, 3080) = 39.49, p < .001, \eta^2 = 0.07$, and Emotion, $F(6, 3080) = 58.96, p < .001, \eta^2 = 0.02$. All interactions between the factors were significant (see Fig. 4a), as well as the three-way interaction of Group $\times$ Covering $\times$ Emotion, $F(18, 3080) = 28.86, p < .001, \eta^2 = 0.01$. 860

Fig. 5a presents the classification confidence obtained in categorizing unmasked emotion images by humans and computer systems. Compared to humans that classify the unmasked images with maximum confidence of only 89% (achieved in the happy expression), the computer systems classify the unmasked 865 emotion images with very high confidence, achieving classification confidence greater than 97% for all the emotion categories except for the neutral expression, which was classified with the confidence of 88.98%. This is understandable as the artificial system classified all images from all expressions correctly (except for neutral expression) with no mask on the face, achieving a classification 870 accuracy of 100%, except for neutral expression with a classification accuracy of 89.33%. Post hoc tests showed that there was a significant difference in the classification confidence for both humans, $F(6, 567) = 36.28, p < .001, \eta^2 = 0.28$ and computer systems, $F(6, 203) = 110.71, p < .001, \eta^2 = 0.77$ ($\alpha = 0.0083$) across different emotion categories when there was no mask on the face. Post 875 hoc t-test for the significance of the classification confidence of each category for the unmasked emotional images by both humans and the computer systems showed a significant difference ($< .001$) for most of the emotion categories. (see

Tables 4b and 4c).

Fig. 5b presents the confidence obtained from categorizing masked emotion
 880 images by humans and computer systems. Here, the computer systems achieved
 higher classification confidence than the humans only in classifying the anger
 expression, achieving classification confidence of 81.98% in contrast to humans
 that classify the anger class images with a classification confidence of 59.06%.
 Surprisingly, while humans classify the happiness expression images with very
 885 high confidence ($> 89\%$) compared to other emotion categories (all $< 84\%$),
 the classification confidence of the computer systems is lowest for happiness
 expression ($< 0.001\%$). This shows that the computer systems have nearly
 zero confidence in the classification of happiness when the mouth is covered
 with a full covering face mask. Also, post hoc tests showed that there was a
 890 significant difference in the classification confidence for both humans, $F(6, 567)$
 $= 49.70$, $p < .001$, $\eta^2 = 0.35$ and computer systems, $F(6, 203) = 205.75$, $p <$
 $.001$, $\eta^2 = 0.86$ ($\alpha = 0.0083$) across different emotion categories with masks
 on the face. Post hoc two-sample t-tests for the differences in classification
 confidence of each category for the masked emotional images by both humans
 895 and computer systems showed a significant difference ($p < .001$) for most of the
 emotion categories (see Tables and 4d and 4e).

Fig. 5c presents the confidence obtained from categorizing images by hu-
 mans and computer systems when the partial mask was used as the type of
 covering. As can be seen from the figure, the confidence of computer systems
 900 in categorizing happy and surprise expressions increased when the partial mask
 was used as the type of covering. However, the classification confidence of hu-
 mans in categorizing happiness and surprise expression decreased when partial
 masks were used compared to when there were fully covering masks on the face.
 Nevertheless, post hoc tests showed a significant difference in the classification
 905 confidence for both humans, $F(6, 567) = 68.13$, $p < .001$, $\eta^2 = 0.42$ and com-
 puter systems, $F(6, 203) = 300.06$, $p < .001$, $\eta^2 = 0.90$ ($\alpha = 0.0083$) across
 different emotion categories when the partial masks was used. Post hoc two-
 sample t-tests for the differences in classification confidence of each category for

(a) unmasked

(b) masked

(c) partial mask

(d) sunglasses

Figure 5: Violin plots showing confidence obtained from categorizing a) unmasked, b) masked, c) partial mask, and d) sunglasses emotions by humans and computer systems.

the partial mask emotional images (by both humans and computer systems) also
 910 showed a significant difference for most of the emotion categories (see Tables 5a
 and 5b).

Fig. 5d presents the confidence obtained by humans and computer systems
 from categorizing sunglasses images. Here, the confidence achieved by both
 humans and computer systems has decreased in all classes (except for surprise
 915 in humans) when there was sunglasses on the face compared with when there
 were no coverings added to the face. The computer system has experienced its
 largest decrease in neutral, sad, and surprise expressions (all $> 28\%$ decrease)
 whereas the largest decrease was seen in disgust, fear, and sad expressions (all
 $> 17\%$ decrease) for humans. Post hoc tests showed that there was a significant
 920 difference in the classification confidence for both humans, $F(6, 567) = 82.47$,
 $p < .001$, $\eta^2 = 0.47$ and computer systems, $F(6, 203) = 12.58$, $p < .001$, η^2
 $= 0.27$ ($\alpha = 0.0083$) across different emotion categories when the sunglasses
 was used as the type of covering. Post hoc two-sample t-tests for differences in
 classification confidence of each category for the sunglasses emotional images by
 925 both humans and computer systems showed a significant difference for most of
 the emotion categories (see Tables 5c and 5d).

7.2. Comparison of recent work carried out to categorize emotion from faces with different coverings.

Table 6 presents results from recent studies carried out to categorize emotion
 930 from faces with different coverings.

It is worth noting that the results presented are not a like-for-like comparison
 as the compared studies used different datasets and presented a different number
 of images to the machine learning classifiers and/or the human participants.
 Besides, different stimulus sets might have varying difficulties, and therefore,
 935 these results may not be directly compared.

However, what is consistent is that, across all of these studies, the accuracy
 achieved by humans and the machine learning classifiers reduced from when
 there was no covering to when any type of covering (either the fully covering

	Sum Square	df	F	<i>p</i>	Eta Squared
Intercept	9.08E+04	1.0	243.79	<.001	0.05
C(Group)	5.21E+04	1.0	139.85	<.001	0.03
C(Covering)	4.41E+04	3.0	39.49	<.001	0.07
C(Emotion)	1.32E+05	6.0	58.96	<.001	0.02
C(Group): C(Covering)	2.02E+04	3.0	18.04	<.001	0.09
C(Group): C(Emotion)	1.72E+05	6.0	76.87	<.001	0.01
C(Covering): C(Emotion)	1.28E+05	18.0	19.03	<.001	0.06
C(Group): C(Covering):C(Emotion)	1.94E+05	18.0	28.86	<.001	0.1
Residual	1.15E+06	3080.0	-	-	-

(a) ANOVA comparing humans and computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.159 (-0.37)	0.257 (0.30)	<.001 (-1.67)	<.001 (-1.03)	<.001 (-1.03)	<.001 (-1.21)
Disgust	0.159 (-0.37)	-	0.019 (0.63)	<.001 (-1.10)	0.038 (-0.55)	0.027 (-0.59)	0.008 (-0.72)
Fear	0.257 (0.30)	0.019 (0.63)	-	<.001 (-1.83)	<.001 (-1.25)	<.001 (-1.24)	<.001 (-1.41)
Happy	<.001 (-1.67)	<.001 (-1.10)	<.001 (-1.83)	-	0.017 (0.64)	0.086 (0.45)	0.129 (0.40)
Neutral	<.001 (-1.03)	0.038 (-0.55)	<.001 (-1.25)	0.017 (0.64)	-	0.727 (-0.09)	0.429 (-0.21)
Sad	<.001 (-1.03)	0.027 (-0.59)	<.001 (-1.24)	0.086 (0.45)	0.727 (-0.09)	-	0.718 (-0.09)
Surprise	<.001 (-1.21)	0.008 (-0.72)	<.001 (-1.41)	0.129 (0.40)	0.429 (-0.21)	0.718 (-0.09)	-

(b) unmasked by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.887 (0.04)	0.036 (-0.55)	<.001 (-1.69)	<.001 (2.93)	<.001 (1.64)	<.001 (-1.16)
Disgust	0.887 (0.04)	-	0.122 (-0.40)	0.003 (-0.82)	<.001 (2.93)	<.001 (1.61)	0.025 (-0.61)
Fear	0.036 (-0.55)	0.122 (-0.40)	-	0.062 (-0.51)	<.001 (2.96)	<.001 (1.72)	0.496 (-0.20)
Happy	<.001 (-1.69)	0.003 (-0.82)	0.062 (-0.51)	-	<.001 (2.99)	<.001 (1.80)	<.001 (1.00)
Neutral	<.001 (2.93)	<.001 (2.93)	<.001 (2.96)	<.001 (2.99)	-	<.001 (-2.28)	<.001 (-2.97)
Sad	<.001 (1.64)	<.001 (1.61)	<.001 (1.72)	<.001 (1.80)	<.001 (-2.28)	-	<.001 (-1.76)
Surprise	<.001 (-1.16)	0.025 (-0.61)	0.496 (-0.20)	<.001 (1.00)	<.001 (-2.97)	<.001 (-1.76)	-

(c) unmasked by computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.049 (0.52)	0.157 (-0.37)	<.001 (-1.32)	<.001 (-1.50)	0.365 (0.24)	0.002 (-0.86)
Disgust	0.049 (0.52)	-	0.002 (-0.86)	<.001 (-1.98)	<.001 (-2.10)	0.334 (-0.25)	<.001 (-1.54)
Fear	0.157 (-0.37)	0.002 (-0.86)	-	0.003 (-0.81)	<.001 (-1.03)	0.029 (0.58)	0.171 (-0.36)
Happy	<.001 (-1.32)	<.001 (-1.98)	0.003 (-0.81)	-	0.245 (-0.31)	<.001 (1.54)	0.029 (0.58)
Neutral	<.001 (-1.50)	<.001 (-2.10)	<.001 (-1.03)	0.245 (-0.31)	-	<.001 (1.70)	0.002 (0.84)
Sad	0.365 (0.24)	0.334 (-0.25)	0.029 (0.58)	<.001 (1.54)	<.001 (1.70)	-	<.001 (-1.10)
Surprise	0.002 (-0.86)	<.001 (-1.54)	0.171 (-0.36)	0.029 (0.58)	0.002 (0.84)	<.001 (-1.10)	-

(d) masked by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	<.001 (5.84)	<.001 (2.88)	<.001 (6.79)	<.001 (4.78)	<.001 (5.53)	<.001 (4.76)
Disgust	<.001 (5.84)	-	<.001 (-2.70)	<.001 (1.18)	0.005 (-0.76)	0.082 (-0.46)	<.001 (-1.81)
Fear	<.001 (2.88)	<.001 (-2.70)	-	<.001 (3.57)	<.001 (1.81)	<.001 (2.37)	<.001 (1.51)
Happy	<.001 (6.79)	<.001 (1.18)	<.001 (3.57)	-	<.001 (-1.63)	<.001 (-1.74)	<.001 (-3.52)
Neutral	<.001 (4.78)	0.005 (-0.76)	<.001 (1.82)	<.001 (-1.63)	-	0.137 (0.40)	0.020 (-0.63)
Sad	<.001 (5.53)	0.082 (-0.46)	<.001 (2.37)	<.001 (-1.74)	0.137 (0.40)	-	<.001 (-1.31)
Surprise	<.001 (4.76)	<.001 (-1.81)	<.001 (1.51)	<.001 (-3.52)	0.020 (-0.63)	<.001 (-1.31)	-

(e) masked by computer systems

Table 4: Statistical test results presenting a) ANOVA obtained from comparing humans and computer systems based on their confidence of categorizing people's emotions with different types of covering. b) Post hoc test (*t*-test) results showing the *p*-value (and Cohen's *d* effect size) obtained from comparing unmasked images classification confidence by humans and c) computer systems. d) Post hoc test (*t*-test) results showing the *p*-value (and Cohen's *d* effect size) obtained from comparing masked images classification confidence by humans and e) computer systems. Note that the α level in all the post hoc comparisons has been adjusted to 0.0014 using *Bonferroni correction*.

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.036 (-0.56)	0.014 (-0.66)	<.001 (-2.70)	<.001 (-1.98)	<.001 (-1.21)	<.001 (-2.34)
Disgust	0.036 (-0.56)	-	0.838 (-0.05)	<.001 (-1.76)	<.001 (-1.21)	0.039 (-0.55)	<.001 (-1.50)
Fear	0.014 (-0.66)	0.838 (-0.05)	-	<.001 (-1.88)	<.001 (-1.25)	0.044 (-0.53)	<.001 (-1.58)
Happy	<.001 (-2.70)	<.001 (-1.76)	<.001 (-1.88)	-	0.036 (0.56)	<.001 (1.27)	0.370 (0.24)
Neutral	<.001 (-1.98)	<.001 (-1.21)	<.001 (-1.25)	0.036 (0.56)	-	0.010 (0.69)	0.237 (-0.31)
Sad	<.001 (-1.21)	0.039 (-0.55)	0.044 (-0.53)	<.001 (1.27)	0.010 (0.69)	-	<.001 (-1.00)
Surprise	<.001 (-2.34)	<.001 (-1.50)	<.001 (-1.58)	0.370 (0.24)	0.237 (-0.31)	<.001 (-1.00)	-

(a) partial mask by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	<.001 (8.02)	<.001 (8.07)	<.001 (-2.49)	0.613 (0.13)	<.001 (9.91)	<.001 (-2.93)
Disgust	<.001 (8.02)	-	0.287 (0.28)	<.001 (-6.62)	<.001 (-3.49)	0.007 (0.73)	<.001 (-6.75)
Fear	<.001 (8.07)	0.287 (0.28)	-	<.001 (-6.69)	<.001 (-3.57)	0.265 (0.30)	<.001 (-6.82)
Happy	<.001 (-2.49)	<.001 (-6.62)	<.001 (-6.69)	-	<.001 (1.92)	<.001 (7.06)	<.001 (-0.49)
Neutral	0.613 (0.13)	<.001 (-3.49)	<.001 (-3.57)	<.001 (1.92)	-	<.001 (3.77)	<.001 (-2.31)
Sad	<.001 (9.91)	0.007 (0.73)	0.265 (0.30)	<.001 (7.06)	<.001 (3.77)	-	<.001 (-7.15)
Surprise	<.001 (-2.93)	<.001 (-6.75)	<.001 (-6.82)	<.001 (-0.49)	<.001 (-2.31)	<.001 (-7.15)	-

(b) partial mask by computer systems

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.394 (0.22)	<.001 (0.88)	<.001 (-1.85)	<.001 (-0.97)	0.144 (0.38)	<.001 (-1.51)
Disgust	0.394 (0.22)	-	0.021 (0.62)	<.001 (-1.96)	<.001 (-1.14)	0.619 (0.13)	<.001 (-1.65)
Fear	<.001 (0.88)	0.021 (0.62)	-	<.001 (-2.76)	<.001 (-1.83)	0.041 (-0.54)	<.001 (-2.41)
Happy	<.001 (-1.85)	<.001 (-1.96)	<.001 (-2.76)	-	0.005 (0.77)	<.001 (2.41)	0.229 (0.32)
Neutral	<.001 (-0.97)	<.001 (-1.14)	<.001 (-1.83)	0.005 (0.77)	-	<.001 (1.41)	0.079 (-0.47)
Sad	0.144 (0.38)	0.619 (0.13)	0.041 (-0.54)	<.001 (2.41)	<.001 (1.41)	-	<.001 (-2.02)
Surprise	<.001 (-1.51)	<.001 (-1.65)	<.001 (-2.41)	0.229 (0.32)	0.079 (-0.47)	<.001 (-2.02)	-

(c) sunglasses by humans

	Anger	Disgust	Fear	Happy	Neutral	Sad	Surprise
Anger	-	0.548 (0.16)	0.017 (0.64)	0.771 (0.08)	<.001 (1.42)	<.001 (1.53)	<.001 (1.43)
Disgust	0.548 (0.16)	-	0.202 (0.34)	0.810 (-0.06)	<.001 (1.18)	<.001 (1.04)	<.001 (1.09)
Fear	0.017 (0.64)	0.202 (0.34)	-	0.127 (-0.41)	<.001 (1.00)	0.004 (0.78)	0.002 (0.87)
Happy	0.771 (0.08)	0.810 (-0.06)	0.127 (-0.41)	-	<.001 (1.23)	<.001 (1.10)	<.001 (1.15)
Neutral	<.001 (1.42)	<.001 (1.18)	<.001 (1.00)	<.001 (1.23)	-	0.080 (-0.47)	0.333 (-0.26)
Sad	<.001 (1.53)	<.001 (1.04)	0.004 (0.78)	<.001 (1.10)	0.080 (-0.47)	-	0.394 (0.23)
Surprise	<.001 (1.43)	<.001 (1.09)	0.002 (0.87)	<.001 (1.15)	0.333 (-0.26)	0.394 (0.23)	-

(d) sunglasses by computer systems

Table 5: Statistical test results for post hoc tests: a) p -value (and Cohen's d effect size) comparing partial mask images classification confidence by humans and b) computer systems. c) p -value (and Cohen's d effect size) for comparing sunglasses images classification confidence by humans and d) computer systems. Note that the α level has been adjusted to 0.0014 using *Bonferroni correction*.

Table 6: Table presenting the results of recent work carried out to categorize emotion from faces with different coverings. Note that results do not represent a like-for-like comparison.

Reference	Covering Type	Classifier	Accuracy
Carbon et al. [1]	None	Humans	89.67%
	Fully Covering Mask		72.77%
Noyes et al. [5]	None	Humans	80.48%
	Fully Covering Mask		61.51%
	Sunglasses		76.63%
Marini et al. [15]	None	Humans	93.00%
	Fully Covering Mask		93.00%
	Partial Mask		81.00%
Shehu et al. [63]	None	Machine Learning	100.0%
	Fully Covering Mask		57.71%
	Partial Mask		76.86%
	Sunglasses		86.57%
Jiang et al. [64]	None	Machine Learning	75.18%
	Sunglasses		69.09%
This work	None	Humans	82.72%
	Fully Covering Mask		57.49%
	Partial Mask		72.93%
	Sunglasses		72.93%
	None	Machine Learning	98.48%
	Fully Covering Mask		24.00%
	Partial Mask		32.48%
	Sunglasses		80.00%

mask, partial mask, or the sunglasses) was added to the face. Consistent with
940 the current study, these results (see Table 6) suggest that partial face coverings
impair emotional facial expression for both humans and computer systems by a
large amount.