

Does a Face Mask Protect my Privacy?: Deep Learning to Predict Protected Attributes from Masked Face Images

Sachith Seneviratne¹, Nuran Kasthuriarachchi², Sanka Rasnayaka⁴, Danula Hettiachchi³, and Ridwan Shariffdeen⁴

¹ University of Melbourne, Australia - sachith.seneviratne@unimelb.edu.au

² University of Moratuwa, Sri Lanka - nuran.11@cse.mrt.ac.lk

³ RMIT University, Australia - danula.hettiachchi@rmit.edu.au

⁴ National University of Singapore - {sanka,ridwan}@u.nus.edu

Abstract. Contactless and efficient systems are implemented rapidly to advocate preventive methods in the fight against the COVID-19 pandemic. Despite the positive benefits of such systems, there is potential for exploitation by invading user privacy. In this work, we analyse the privacy invasiveness of face biometric systems by predicting privacy-sensitive soft-biometrics using masked face images. We train and apply a CNN based on the ResNet-50 architecture with 20,003 synthetic masked images and measure the privacy invasiveness. Despite the popular belief of the privacy benefits of wearing a mask among people, we show that there is no significant difference to privacy invasiveness when a mask is worn. In our experiments we were able to accurately predict sex (94.7%), race (83.1%) and age (MAE 6.21 and RMSE 8.33) from masked face images. Our proposed approach can serve as a baseline utility to evaluate the privacy-invasiveness of artificial intelligence systems that make use of privacy-sensitive information. We open-source all contributions for reproducibility and broader use by the research community.

Keywords: COVID-19, Masked Faces, Privacy, Computer Vision

1 Introduction

Since the outbreak of SARS-CoV-2 (COVID-19), the use of face masks has become ubiquitous around the world and has been identified as an important public health response to fight against the ongoing pandemic. The mass shift to wearing masks during the COVID-19 pandemic has radically changed the way in which many of our mundane activities are carried out. This situation demands the enablement of contactless and efficient operations, especially in retail services. Contactless technologies like face and iris based detection systems are pushed to reach newer heights, in contrast applications that rely on fingerprint recognition modalities suffer a significant loss due to the emerging requirements as an after-effect of the COVID-19 pandemic [7]. In particular, face recognition is praised as

one of the efficient and contactless means of verifying identity and prior research has studied the impact and techniques to improve face-recognition systems to further advance contactless operations [6, 16]. Using computer vision to enhance contactless and efficient operations has shown promise in various applications (i.e. public compliance monitoring [19]). In this work, we investigate the impact of using computer vision, specifically in face authentication systems for contactless identification and the possible implications on privacy. Despite the scalable automation it provides, face-recognition technology needs to adhere to the privacy regulations such as the General Data Protection Regulation (GDPR) and improve the perception of users to increase trust and acceptance.

Considering the advancements in surveillance and monitoring technologies in response to COVID-19, the norms of acceptable information flow may shift. For instance, users’ perspectives on the use of location information (which is privacy-sensitive information), has drastically changed in times of crisis [15]. However, such temporary measures during a crisis may not prevail as a permanent and long-term acceptance because it would unnecessarily reduce a person’s privacy. Although a wider acceptance of surveillance systems can be seen in the current situation, we argue that a popular misconception of, “wearing face masks will increase privacy protection” exists among most people. Therefore, we first investigate the perception of users with respect to face biometric solutions and their understanding of privacy protection. We conduct an online survey to extract the opinions of users on their privacy with face masks and we learn that generally users have higher confidence of privacy protection when using a face mask. In our study, we find that perceived privacy of wearing a mask is higher with a statistical significance ($P=0.00964 < 0.05$).

Systems that use face biometrics could potentially reveal privacy-sensitive information such as soft-biometrics, which includes but are not limited to age, sex and race. Many of the artificial intelligence systems use such privacy-sensitive information but are restricted for the intended purposes. We evaluate the possible violations of privacy-protection in such systems that use face biometrics, with the use of masked face images and quantify the privacy invasiveness of such implementations. We implement several techniques to predict privacy-sensitive soft-biometrics such as age, sex and race, and we were able to achieve an accuracy of 94.7% and 83.1% in correctly classifying the sex and race, respectively. We were also able to accurately predict the age with an RMSE score of 8.33 and MAE score of 6.21. We then analysed the privacy invasiveness in our implementation for images with mask and without mask, to understand the privacy preservation when using a face mask. We show that there is no significant difference in privacy protection by quantifying the privacy invasiveness using the Privacy Vulnerability Index (PVI) [25] for both settings, which recorded **only a 2.9% difference** that implies no significance in wearing a mask.

In this work, we make three contributions: **(1)** Quantitative analysis on privacy invasion on masked face images. To the best of our knowledge, we are the first to study the predictability of age, sex and race using masked face images. **(2)** Study the perception of privacy protection on wearing a face mask. Our

results indicate that people consider masked faces to be less privacy invasive. In support of reproducible research, we open-source our model weights and scripts for the benefit of the research community. These models will enable future study on masked face biometric systems related to privacy protection.

2 Related Work

2.1 Biometrics and Privacy

The use of biometrics have raised various privacy concerns due to the possibility of predicting protected attributes. Many studies have evaluated the predictability of soft-biometric attributes such as age, gender and race using common biometrics such as face [14], iris [28], fingerprint [3], voice [10] and gait [25]. In this work, we go beyond than prediction and provide means of quantifying the privacy invasiveness in systems that use soft-biometric.

2.2 Face Biometric and Masks

While computer vision research has examined face recognition methods robust to partial occlusions [18,30], with increased global mask use due to COVID-19, there is a renewed interest in masked face recognition. Recent work shows that current state of the art face recognition methods trained with full face images fail in accurately recognising masked faces [12]. Although researchers have created real-world masked face datasets [6,29], there is limited work on developing specific machine learning models trained with masked images. In addition, face masks have also introduced a family of computer vision challenges such as mask detection [21]. While some prior work has implications on masked biometric analysis using masks [2], they have not used masked facial images for analysis. In particular, while biometric analysis focused around the ocular region can provide useful insights into masked facial analysis, we argue that only actual analysis on masked facial imagery provides realistic insights into masked biometric analysis. This arises from the fact that, based on the masking process used, considerable portions of the ocular region may be occluded as well. Therefore, performing end-to-end evaluation of masked facial images provides a more realistic picture of the situation corresponding to real world usage.

2.3 User Perception

User perceptions towards biometric modalities tend evolve with time [24]. State of the art face recognition methods can achieve high levels of accuracy and are widely used for different applications, including authentication and surveillance. While users are generally more familiar and comfortable with face biometric solutions [5], users also tend to resist face recognition based solutions due to privacy concerns [20]. Furthermore, wearing masks can limit the face area exposed to face recognition systems. However, there is limited literature on how people

perceive the difference between masked and non-masked face recognition. This understanding could be potentially influenced by other challenges humans face when wearing masks. For instance, research shows that people find it challenging to match familiar faces, match unfamiliar faces and recognise emotions when faces are occluded with objects such as masks and sunglasses [23].

3 Methodology

Our evaluation consists of three components. First, we conducted a survey to understand the perception of users on privacy protection while wearing a mask, from the increased surveillance systems due to COVID-19. Second, we generate a synthetic face mask dataset, predict protected attributes from masked face images and compare our results with prior work that use non-masked faces. Third, we show how unmasked face images invade privacy and analyse the impact of image attributes on our predictions.

3.1 User Perception Survey

The main objective of this survey was to study the perceptions of people towards the privacy invasiveness of masked faces in comparison to unmasked face images. We aim to answer the following questions: “*Do people feel that wearing a face mask will protect their privacy?*” and “*Which is considered more private among Age, Race and Sex*”. An online survey was designed to collect this information with *Yes/No questions* comparing the privacy invasiveness of masked and unmasked face images, *Three point Likert scale questions* evaluating perceived privacy invasiveness of masked images and unmasked images and a *Sorting Activity* to sort Age, Race and Sex based on importance. The survey is conducted anonymously on a voluntary basis in June 2021. The relative ordering of the sorting activity will be used to measure the Relative Importance Index (RII) value for each of the three attributes.

3.2 Dataset and Synthetic Mask Generation

There is no openly-available large-scale mask dataset with soft-biometric labels for age, gender and race. Therefore, we select UTK faces dataset, the most commonly cited face dataset in the literature and generate a masked dataset by digitally painting a mask on top of the face image. We follow the process outlined in [22] to generate synthetic masks on the face images. This process is depicted in Fig. 1. We open-source the scripts used for this process.

UTK faces dataset has 23,542 face images with labels for age, gender and race. Following the masking process [22] we create a data-set of 23,002 masked images. We show the distributions of the attributes in Figure 2. We bin the ages as follows, baby: 0-3 years, child: 4-12 years, teenagers: 13-19 years, young: 20-30 years, adult: 31-45 years, middle aged: 46-60 years and senior: 61 years and above, in line with the analysis in [13].

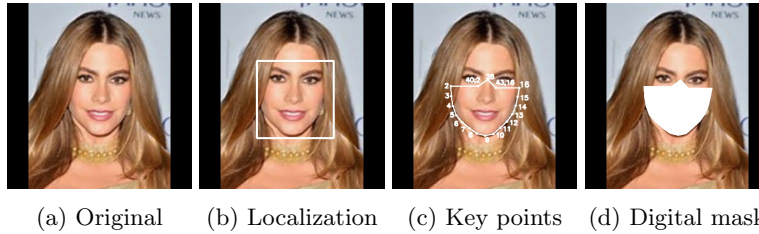


Fig. 1: Synthetic mask creation pipeline (an example from UTK faces)

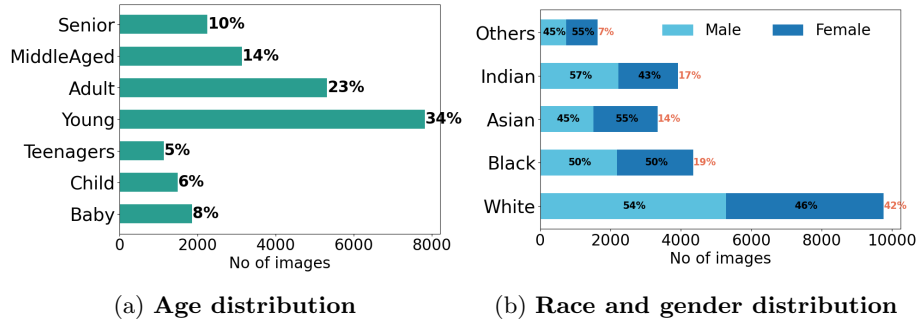


Fig. 2: UTK faces dataset summary

3.3 Computer Vision Workflow

We use a computer vision based method with convolutional neural networks (CNN) to build models for 3 different tasks - age, sex and race prediction. Rather than building individual CNNs from scratch for each task, we build an initial facial representation within the neural network by using only the UTK facial dataset. We use a ResNet50 architecture and pre-train our representation for 3038 epochs in an unsupervised manner using the framework introduced in [17] and the projection head and augmentations from [8] using the parameters and settings described in [9]. Pre-training is carried out on a 4 GPU node with batch size set to 128. This pre-trained representation is then fine-tuned end-to-end for each specific task with a new fully connected layer incorporated past the final bottleneck layer of the pre-trained ResNet50 architecture. The size of this final output layer depends on the task. For regression the layer has a single node - as the output is a single continuous variable, while for classification we incorporate a number of nodes equal the number of classes in the problem (for instance, 2 for sex, 5 for race and 7 for age - corresponding to categorical labels). Training is carried out for 3500 epochs for each task with Stochastic Gradient Descent and learning rate $1e-3$, and we isolate the single checkpoint with the best validation performance to evaluate on a holdout dataset. We evaluate the impact of additional image augmentations using RandAugment [11] with default ImageNet

parameters. We open-source all contributions⁵, including trained models on the different splits and dataset splits for full reproducibility.

We evaluate against an open-sourced masked-facial representation (MUFM) [27] released as part of the Masked Facial Recognition Competition 2021 [4], which claims to be a generic masked-face representation adaptable to any task on masked faces. We run this evaluation for the task of sex classification on a random split of 70% training 20% validation and 10% testing on the UTK dataset. Based on results, we extend our analysis using the best performing combination of representation and technique for the tasks of race classification and age regression on similar random splits of UTK dataset (see Table 1). As a follow up experiment we build models on a new split of UTK data that ensures a uniform split as discussed in Section 3.2. We train models as before (discarding previously trained models), but change the age regression to an age bracket based classification following other work in the literature [13]. By doing so we compare against multiple existing state of the art techniques for age prediction (see Table 2). Note that our models are at a disadvantage due to roughly half of the face being absent/occluded in the image.

3.4 Privacy Vulnerability Index(PVI)

The Privacy Vulnerability Index [25] is used to quantify the privacy invasiveness of a biometric modality. We use this measure to compare the privacy invasiveness of face images and masked face images. The PVI of a biometric depends on two factors, **(1) Predictability (p_i)**: how well can protected attributes be predicted using the biometric modality, measured by classification accuracy. **(2) Importance (s_i)** : how important is each personal attribute, measured using the RII calculated from the user perception study. The PVI value for masked and unmasked images is calculated as a weighted sum of these two values using the equation, $PVI = (\sum_i s_i * p_i) / \sum_i s_i$.

4 Evaluation Results

4.1 User Perception Study

The survey resulted in 60 complete responses. The users' responses to if the face image and masked face image could lead to privacy invasiveness is used to examine if there is a statistically significant difference in the perception towards the two modalities. We perform a the Mann-Whitney U test with a single-tail, to show that the perceived privacy of wearing a mask is higher with statistical significance ($P = 0.00964 < 0.05$). Figure. 3 show the distribution of user responses.

Survey participants were asked if wearing a mask preserves privacy (compared to not wearing a mask). 50% of the participants said yes while 40% felt

⁵ <https://github.com/sachith500/MaskedFaceRepresentation>

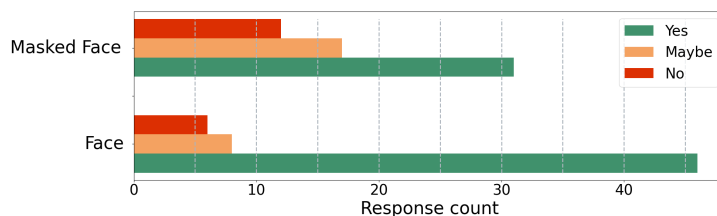


Fig. 3: User perception on privacy invasion with masked face images compared to face images.

both violated privacy equally and 10% said masked face images violates privacy more.

Relative Importance: The resulting relative ordering with the Relative Importance Index (RII) values within brackets is; 1) Age [0.3765] 2) Race [0.3353] 3) Gender [0.2882].

4.2 Prediction Accuracy

Table 1 presents the overall accuracy for models built for masked face images. Table 2 compares the results of masked and unmasked faces.

Table 1: Attribute prediction using masked face images. The first experiment (sex) was used to verify that our model performance was superior to existing masked representations.

Method	Sex	Race	Age	
	Accuracy	Accuracy	MAE	RMSE
Using representation [27] + transforms [11]	0.9374	-	-	-
Our method with transforms from [11]	0.9401	0.8220	6.2788	8.4836
Our method without complex transforms	0.9361	0.8134	6.2168	8.3372

4.3 Impact of Image Attributes

We examine whether the original user attributes (i.e., Sex, Race, Age category) influence our prediction outcome of masked faces. To this end, we select the best performing model from initial evaluation (Table 1) and evaluate sex, race and age prediction models using a new uniform test split with a balanced attribute composition. Fig. 4 presents confusion matrices for each model outcome. We obtained an overall accuracy of 94.65% for sex, 83.12% for race and 67.94% for age category.

Table 2: Overall result comparison with SOTA for each protected-attribute. Models are retrained for the uniform split using optimal parameters from experiments in Table 1.

	Unmasked Face - SOTA	Masked Face (Random Split)	Masked Face (Uniform Split)
Sex	[13] 98.23%	94.01%	94.65%
Race	[1] 91.23%	82.20%	83.12%
Age (MAE) - Regression	[26] 5.44	6.21	-
Age - Classification	[13] 70.1%	-	67.94%

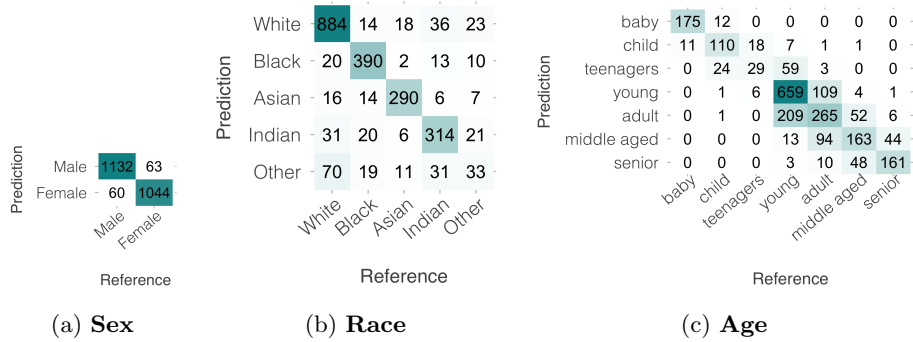


Fig. 4: Confusion matrices for Sex, Race and Age prediction using masked images.

We conduct chi-square tests of independence to examine the relationship between different image attributes and the ability to accurately predict the them. When considering the image attribute sex there is no significant difference between prediction outcomes of sex ($\chi^2(1) = 0.006$, $p = 0.936$) and race ($\chi^2(1) = 0.578$, $p = 0.447$). However, a significant difference is noted for age category ($\chi^2(1) = 4.019$, $p = 0.045 < 0.05$). Furthermore, when considering race, there is a significant difference in prediction outcomes for sex ($\chi^2(4) = 12.53$, $p = 0.014 < 0.05$), race ($\chi^2(4) = 523.07$, $p < 0.001$) and age ($\chi^2(4) = 49.951$, $p < 0.001$) prediction. Similarly, for image attribute age, there is a significant difference for outcomes of sex ($\chi^2(6) = 164.57$, $p < 0.001$), race ($\chi^2(6) = 13.449$, $0.036 < 0.05$) and age ($\chi^2(6) = 374.08$, $p < 0.001$) prediction. In summary, we note that image attributes race and age category having a significant impact on all the prediction outcomes while sex only influence age prediction. In addition, Fig. 5 provides the accuracy for each subgroup of images based on sex, race, age category of the person appearing in the image.

4.4 Privacy Invasiveness

We use the PVI equation with the SOTA for face images and our best results for masked face images to quantify the level of privacy invasiveness of both.

Age	Age Prediction					Race Prediction					Sex Prediction				
	White	Black	Asian	Indian	Other	White	Black	Asian	Indian	Other	White	Black	Asian	Indian	Other
baby	0.87	1	0.98	0.94	1	0.91	0.67	0.98	0.53	0.42	0.65	1	0.83	0.76	0.83
child	0.75	0.6	0.62	0.83	1	0.95	0.8	0.88	1	0	0.92	0.8	0.62	0.83	0.67
teenagers	0.22	0	0	0.2	0.25	0.97	0.86	0.33	0.6	0.5	0.91	1	1	1	1
young	0.64	0.82	0.85	0.75	0.93	0.83	0.94	0.9	0.75	0.18	0.98	0.95	0.92	0.99	0.82
adult	0.61	0.41	0.52	0.69	0.47	0.87	0.97	0.81	0.83	0	0.99	0.99	1	0.98	0.89
middle aged	0.61	0.42	0.54	0.57	0.57	0.9	0.81	0.69	0.8	0.14	0.97	1	1	0.98	1
senior	0.73	0.5	1	0.67	1	0.95	0.85	0.6	0.75	0	0.97	0.95	0.93	0.92	1
baby	0.88	1	1	0.88	0.86	0.96	0.5	0.97	0.94	0.29	0.64	0.75	0.71	0.82	0.86
child	0.84	0.83	0.64	0.64	0.58	0.95	0.67	0.71	0.73	0.17	0.88	1	0.93	0.91	0.92
teenagers	0.46	0	0	0.14	0.25	0.97	0.5	0.75	0.57	0.38	1	0.75	1	1	1
young	0.8	0.9	0.99	0.97	0.85	0.84	0.89	0.91	0.89	0.2	0.99	0.97	0.99	1	0.98
adult	0.52	0.29	0.21	0.42	0.38	0.9	0.91	0.89	0.77	0.23	0.96	0.97	1	1	0.92
middle aged	0.45	0.24	0.67	0.33	0	0.97	0.88	0.33	0.67	0	0.97	0.82	1	1	1
senior	0.76	0.62	0.86	0.2		0.96	1	1	1		0.93	0.75	0.86	1	

Fig. 5: Attribute prediction accuracy for each sub group

$PVI_f = 0.828$, $PVI_{mf} = 0.853$. The privacy invasiveness reduction by wearing a mask is only 2.9%. Which is very low compared to the 50% of people who thought that masked faces to be more private.

5 Discussion and Conclusion

5.1 Predicting Protected Attributes

Our study shows it is possible to predict sex, race, and age with a high degree of accuracy. When compared to the state of the art methods that predict these attributes for non-masked face images, we only note absolute accuracy differences of 3.58% for sex, 8.11% for race and 2.16% for age categories with nearly half the face (with key features like facial hair and lips for sex, wrinkles for age) occluded by a mask. Based on results in Table 2, we find that incorporating augmentations during training can improve sex and race prediction. During age prediction it slightly lowers accuracy. This likely stems from how predicting age is harder than race or sex (even for humans) and augmentations may create discrepancies between fine-grained features such as wrinkles which has less impact on predicting race or sex.

5.2 Biases from Image Attributes

While our models achieved high levels of overall accuracy for masked images, we observed that image attributes race and age can influence the prediction outcomes. For instance, age category prediction accuracy for teenagers (25.21%) and adults (49.71%) is low compared to the overall accuracy (67.94%). As shown in Figure 5, prediction accuracies are consistently low across different sex and

race categories as well. This is inline with results reported in prior work on biases in user attribute classification using regular face images [13]. Furthermore, when considering race prediction, we note that race prediction accuracy is lower for Other category (20.12%) with limited samples, when compared to the overall value (83.12%). We argue that biases from user attributes can greatly influence the model outcomes. Therefore, appropriate measure should be taken to account for sampling biases particularly for commercial applications of face-recognition technology.

5.3 Privacy Preservation

Our study highlights a mismatch between user perception and the reality regarding privacy preservation through face masks. Compared to regular face exposure, users perceive a significantly higher level of privacy when wearing face masks. However, we show that the ability to predict protected attributes from masked face images is not largely different from face images (Table 2) and the privacy invasiveness reduction by wearing a mask is only 2.9%. This inaccurate perceived privacy could lead to a false sense of safety for masked users, and therefore users could be targets for exploitation by malicious applications. In addition, distinct characteristics of face masks could contribute to more robust surveillance applications that users are not aware of. In the light of increased make usage, we argue that it is essential to raise user awareness and research privacy protection methods concerning face masks.

5.4 Limitations

We note a few limitations in our study. First, as there is no masked image dataset available with attributes such as sex, race and age, our evaluation is based on a synthetic mask generation process. Second, our user study is limited to 60 participants and we did not collect demographic information which may reveal interesting insights. Third, our source dataset has imbalances among classes which is reflected in our analysis.

5.5 Conclusion and Future Work

In this paper, we predict sex (94.7%), race (83.1%) and age (68.0%) on masked face images using a computer vision approach. Despite the popular belief that masks protect user privacy, we show that masks only reduce privacy invasiveness by 2.9% when compared to state of the art face recognition approaches. We further analyse the impact of image labels on the prediction ability and provide a baseline for future research by open-sourcing our models. Our research paves the way for future work that aim to study how to preserve user privacy when wearing masks while maintaining utility as a biometric modality. We open-source our contributions, including masking and inference scripts, trained models and data splits for reproducibility and broader use for both privacy and mask related research.

References

1. Ahmed, M.A., Choudhury, R.D., Kashyap, K.: Race estimation with deep networks. *Journal of King Saud University-Computer and Information Sciences* (2020)
2. Alonso-Fernandez, F., Diaz, K.H., Ramis, S., Perales, F.J., Bigun, J.: Soft-biometrics estimation in the era of facial masks. In: 2020 International Conference of the Biometrics Special Interest Group (BIOSIG). pp. 1–6 (2020)
3. Badawi, A.M., Mahfouz, M., Tadross, R., Jantz, R.: Fingerprint-based gender classification. *IPCV* **6**, 41–46 (2006)
4. Boutros, F., Damer, N., Kolf, J.N., Raja, K., Kirchbuchner, F., Ramachandra, R., Kuijper, A., Fang, P., Zhang, C., Wang, F., et al.: Mfr 2021: Masked face recognition competition. In: 2021 IEEE International Joint Conference on Biometrics (IJCB). pp. 1–10. IEEE (2021)
5. Buckley, O., Nurse, J.R.: The language of biometrics: Analysing public perceptions. *Journal of Information Security and Applications* **47**, 112–119 (2019). <https://doi.org/10.1016/j.jisa.2019.05.001>, <https://doi.org/10.1016/j.jisa.2019.05.001>
6. Cabani, A., Hammoudi, K., Benhabiles, H., Melkemi, M.: MaskedFaceNet – A dataset of correctly/incorrectly masked face images in the context of COVID-19. *Smart Health* **19**(October 2020), 100144 (2021). <https://doi.org/10.1016/j.smhl.2020.100144>, <https://doi.org/10.1016/j.smhl.2020.100144>
7. Carlaw, S.: Impact on biometrics of covid-19. *Biometric Technology Today* **2020**(4), 8–9 (2020). [https://doi.org/https://doi.org/10.1016/S0969-4765\(20\)30050-3](https://doi.org/https://doi.org/10.1016/S0969-4765(20)30050-3), <https://www.sciencedirect.com/science/article/pii/S0969476520300503>
8. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations (2020)
9. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning (2020)
10. Childers, D.G., Wu, K.: Gender recognition from speech. part ii: Fine analysis. *The Journal of the Acoustical society of America* **90**(4), 1841–1856 (1991)
11. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space (2019)
12. Damer, N., Grebe, J.H., Chen, C., Boutros, F., Kirchbuchner, F., Kuijper, A.: The Effect of Wearing a Mask on Face Recognition Performance: An Exploratory Study. *BIOSIG 2020 - Proceedings of the 19th International Conference of the Biometrics Special Interest Group* (2020)
13. Das, A., Dantcheva, A., Bremond, F.: Mitigating bias in gender, age and ethnicity classification: a multi-task convolution neural network approach. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. pp. 0–0 (2018)
14. Guo, G., Fu, Y., Huang, T.S., Dyer, C.R.: Locally adjusted robust regression for human age estimation. In: 2008 Ieee Workshop on Applications of Computer Vision. pp. 1–6. IEEE (2008)
15. Haggag, O., Haggag, S., Grundy, J., Abdelrazek, M.: Covid-19 vs social media apps: Does privacy really matter? In: 2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS). pp. 48–57 (2021). <https://doi.org/10.1109/ICSE-SEIS52602.2021.00014>
16. Hariri, W.: Efficient Masked Face Recognition Method during the COVID-19 Pandemic (2021). <https://doi.org/10.21203/rs.3.rs-39289/v1>, <http://arxiv.org/abs/2105.03026>

17. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning (2020)
18. Kim, J., Choi, J., Yi, J., Turk, M.: Effective representation using ICA for face recognition robust to local distortion and partial occlusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(12), 1977–1981 (2005). <https://doi.org/10.1109/TPAMI.2005.242>
19. Lanang Afkaar Ar, M., Muzakki Adytia S, S., Nugraha, Y., Rizka R, F., Ernesto, A., Intan Kanggrawan, J., Suherman, A.L.: A computer vision-based object detection and counting for covid-19 protocol compliance: a case study of jakarta. In: 2020 International Conference on ICT for Smart Society (ICISS). vol. CFP2013V-ART, pp. 1–5 (2020). <https://doi.org/10.1109/ICISS50791.2020.9307594>
20. Liu, Y.L., Yan, W., Hu, B.: Resistance to facial recognition payment in China: The influence of privacy-related factors. *Telecommunications Policy* **45**(5) (2021). <https://doi.org/10.1016/j.telpol.2021.102155>
21. Loey, M., Manogaran, G., Taha, M.H.N., Khalifa, N.E.M.: A hybrid deep transfer learning model with machine learning methods for face mask detection in the era of the COVID-19 pandemic. *Measurement: Journal of the International Measurement Confederation* **167**(May 2020), 108288 (2021). <https://doi.org/10.1016/j.measurement.2020.108288>, <https://doi.org/10.1016/j.measurement.2020.108288>
22. Ngan, M.L., Grother, P.J., Hanaoka, K.K.: Ongoing face recognition vendor test (frvt) part 6b: Face recognition accuracy with face masks using post-covid-19 algorithms (2020)
23. Noyes, E., Davis, J.P., Petrov, N., Gray, K.L., Ritchie, K.L.: The effect of face masks and sunglasses on identity and expression recognition with super-recognizers and typical observers. *Royal Society Open Science* **8**(3) (2021). <https://doi.org/10.1098/rsos.201169>
24. Rasnayaka, S., Sim, T.: Who wants continuous authentication on mobile devices? In: 2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS). pp. 1–9 (2018). <https://doi.org/10.1109/BTAS.2018.8698599>
25. Rasnayaka, S., Sim, T.: Your tattletale gait privacy invasiveness of imu gait data. In: 2020 IEEE International Joint Conference on Biometrics (IJCB). pp. 1–10 (2020). <https://doi.org/10.1109/IJCB48548.2020.9304922>
26. Savchenko, A.V.: Efficient facial representations for age, gender and identity recognition in organizing photo albums using multi-output convnet. *PeerJ Computer Science* **5**, e197 (2019)
27. Seneviratne, S., Kasthuriarachchi, N., Rasnayaka, S.: Multi-dataset benchmarks for masked identification using contrastive representation learning. In: 2021 Digital Image Computing: Techniques and Applications (DICTA). pp. 01–08 (2021). <https://doi.org/10.1109/DICTA52665.2021.9647194>
28. Thomas, V., Chawla, N.V., Bowyer, K.W., Flynn, P.J.: Learning to predict gender from iris images. In: 2007 First IEEE International Conference on Biometrics: Theory, Applications, and Systems. pp. 1–5. IEEE (2007)
29. Wang, Z., Wang, G., Huang, B., Xiong, Z., Hong, Q., Wu, H., Yi, P., Jiang, K., Wang, N., Pei, Y., Chen, H., Miao, Y., Huang, Z., Liang, J.: Masked Face Recognition Dataset and Application pp. 1–3 (2020), <http://arxiv.org/abs/2003.09093>
30. Weng, R., Lu, J., Tan, Y.P.: Robust Point Set Matching for Partial Face Recognition. *IEEE Transactions on Image Processing* **25**(3), 1163–1176 (3 2016). <https://doi.org/10.1109/TIP.2016.2515987>, <http://ieeexplore.ieee.org/document/7377089/>