Evaluating Trade-offs in Computer Vision Between Attribute Privacy, Fairness and Utility

William Paul
JHU/APL

Philip Mathew
JHU/APL

Fady Alajaji
Queens University

Philippe Burlina
JHU/APL

Abstract

This paper investigates to what degree and magnitude tradeoffs exist between utility, fairness and attribute privacy in computer vision. Regarding privacy, we look at this important problem specifically in the context of attribute inference attacks, a less addressed form of privacy. To create a variety of models with different preferences, we use adversarial methods to intervene on attributes relating to fairness and privacy. We see that that certain tradeoffs exist between fairness and utility, privacy and utility, and between privacy and fairness. The results also show that those tradeoffs and interactions are more complex and nonlinear between the three goals than intuition would suggest.

1 Introduction

Despite recent successes in developing artificial intelligence (AI) and deep learning (DL) systems [22] for a range of applications including image interpretation, vehicular navigation, robotics, and medicine [17, 29, 28, 31, 9, 25, 33], ensuring these systems are reliable enough to build trust is still an open problem. Trust and ethical concerns about AI systems such as explainability, adversarial vulnerabilities (adversarial machine learning or AML) [16, 10] and importantly, fairness [6, 7] and privacy [32, 26, 18] – the two main foci of this work – have recently put into question the deployment of certain autonomous systems.

A common theme underpinning a number of these concerns is being able to discern what information the model is using to make its decision. Fairness across groups, which is trying to ensure performance of the system is equal in some respect between different subpopulations denoted by some sensitive attribute, is typically viewed as the model learning correlations between the sensitive attribute and the task in an excessive manner. Privacy has a number of different forms, but commonly focuses on cases either where the model may memorize the training data used in some manner or, what we address in this work, where features are used as a surrogate for the original data to infer private attributes about individuals, known as attribute privacy. For attribute privacy, the model effectively carries information about the private attribute into the features, either inadvertently or for use at the task at hand, acting similarly to concerns raised in addressing fairness. However, if we want to address these concerns simultaneously, there can be a balancing act for determining what information is used for the task, leading to potential trade-offs between how well the task is performed (utility) as well as the fairness and attribute privacy of the model. Though such trade-offs have been shown for differential privacy [20], to the best of our knowledge, we have not seen such an analysis of trade-offs between all three areas for attribute privacy or computer vision related tasks such as facial recognition.

Consequently, in this work, we seek to explore and evaluate how these trade-off present themselves, how severe they are, as well as to what degree we can control them. We leverage adversarial methods as one of the most prominent methods to address group fairness, also using these methods for attribute privacy due to similar motivations. We evaluate on a variety of image datasets, focusing on the tasks of classification and facial recognition, and train a diverse family of models with different preferences of each area of concern to better understand their corresponding metrics are affected.
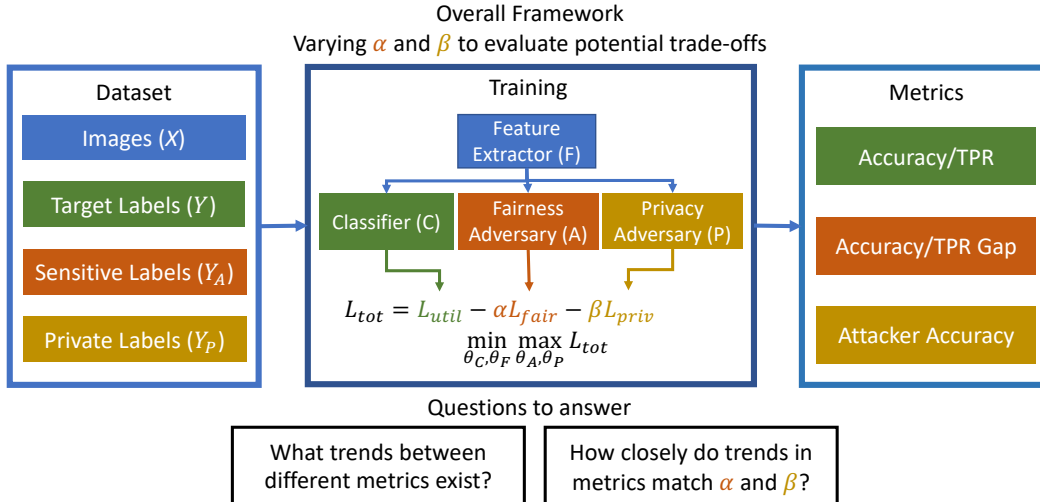
Overall Framework
Varying α and β to evaluate potential trade-offs

Dataset

Images (X)

Target Labels (Y)

Sensitive Labels (Y_A)

Private Labels (Y_P)

Training

Feature Extractor (F)

Classifier (C) Fairness Adversary (A) Privacy Adversary (P)

$$L_{tot} = L_{util} - \alpha L_{fair} - \beta L_{priv}$$
$$\min_{\theta_C, \theta_F} \max_{\theta_A, \theta_P} L_{tot}$$

Metrics

Accuracy/TPR

Accuracy/TPR Gap

Attacker Accuracy

Questions to answer

What trends between different metrics exist?

How closely do trends in metrics match α and β?

Figure 1: Diagram depicting our overall methodology and framework. In order to create classifiers with different characteristics, we modulate the amount of information related to the sensitive and private attributes. Finally, we evaluate utility, fairness, and privacy holistically, answering what interactions exist between them as well as how they interact with the chosen coefficients.

2 Prior Work

There has been increased interest in looking at aspects of assured AI now that deep learning has demonstrated it can operate on par or beyond human capabilities for tasks including classification, detection in natural images and that machine learning model allow for having similar performance to that of clinicians for diagnostics for pathologies like skin lesions or retinopathies [8]. There are a number of different areas within assuring AI, such as being able to provide explanations, robust to imperceptual attacks [10], as well as our focus, ensuring fair and private operation.

Fairness is typically viewed disparities between subpopulations characterized by some attribute, otherwise known as group fairness, and has been investigated in a number of prior studies ranging from natural language processing to medical imagery [36, 5, 36, 21]. A broad taxonomy of methods [11] is organized along the line of approaches that intervene in different parts of the training process, from data preprocessing, calibrating the predictions, and changing the model, including modifying the loss function used which is our approach in this work. Model interventions are a common way to address fairness [15, 4, 1], typically focusing on reducing the undue influence of an attribute within some internal features of the model. This is done via reducing some notion of distance between features of different subpopulatons, such as maximum mean discrepancy, or other form of discrimination such as those induced by adversaries [35, 37, 34].

Privacy has a number of different instantiations, such as protecting against membership inference attacks, where an attacker is trying to infer whether a data point was used for training, or achieving some level of differential privacy, controlling how much influence the training dataset has on the final weights. However, our focus in this work is on a third form called attribute inference attacks, where an attacker is attempting to predict a private attribute about the original data from a transformed or masked version acting as a surrogate. In information theory, attempting to ensure robustness to this attack while maintaining some level of utility of the transformed data is known as the privacy funnel [24, 2]. More specifically, the goal of the privacy funnel is to maximize the mutual information between the transformed data and the original data, while keeping the mutual information between transformed data and the private attribute below a certain threshold. Unlike differential privacy, attribute privacy requires the network to actively prune out information about the private attribute in order to defend against an adversarial threat model. Consequently, there is significant overlap between addressing attribute privacy and group fairness, as both arguably only differ in terms of how the included information in the feature is presented downstream [3].

Addressing both group fairness and attribute privacy simultaneously is less studied however. For differential privacy, [20] appears to be the most prominent approach towards addressing both, though this work focuses on differential privacy. [30, 12] both evaluated how well adversarial learning in facial recognition can help ensure group fairness and attribute privacy. Motivating our work however, they did not investigate what trends exist between fairness, privacy, and utility.

3 Methodology

We detail notation and defintions used, how we construct our framework for experiments, as well as what metrics are used.

Notation and Definitions: In terms of notation used, we consider the dataset to be comprised of images X, target labels Y, predicted labels $\hat{Y}$, sensitive labels Y_A, and private labels Y_P. The models we use consist of: the feature extractor F, converting images to features with weights θ_F; the classifier (C) using the features to predict Y with weights θ_C; the fairness adversary A attempting to predict Y_A from the features using weights θ_A; and the privacy adversary (P) attempting to predict Y_P from the features using weights θ_P. α and β are linear coefficients scaling the loss for the fairness and privacy adversaries.

We review the fairness criteria we focus on below:

Definition 1 (Fairness Criteria). *Given the sensitive label A, target label Y, and the classifier predictions $\hat{Y}$, the classifier is said to satisfy:*

- *accuracy parity if $P(\hat{Y} = Y|Y_A) = P(\hat{Y} = Y)$, i.e., the predictions matching Y are independent of A.*

- *equality of opportunity if for a given Y=y, $P(\hat{Y}|Y_A, Y = y) = P(\hat{Y}|Y = y)$, i.e., the predictions for the class is y, commonly taken to be the positive class, are independent of A.*

Demographic parity and equality of odds are two notable criteria we do not evaluate on. For demographic parity, all of the tasks we consider in this work focus on more descriptive attributes which are less likely to have allocative bias that demographic parity addressed, where the actual labeling assignment is not fair. Similarly, accuracy parity is focused on as a weaker form of equality of odds, not penalizing the model trading false negatives for false positives.

For attribute privacy, we consider the following threat model for attribute privacy:

Definition 2. *(Attribute Privacy) Given oracle access to $F(X)$ for input data X, producing features associated with X, access to labeled public data $\{(X^i, Y_P^i)\}_{i=1,...,n}$, and access to $\{F(\tilde{X}^i)\}_{i=1,...,\tilde{n}}$, the features corresponding to hidden data $\tilde{X}$, an adversary is trying to infer on the hidden data the unknown $\tilde{Y}_P$ corresponding to $\tilde{X}$ better than chance, i.e., achieving $P(\tilde{Y}_P|F(X)) > P(\tilde{Y}_P)$.*

The adversary's objective is to generalize from the features they have corresponding private attributes for to the features. This form of privacy is not addressed by usual forms of differential privacy, which try to lessen the influence of the training data used, as the feature extractor must activity prune information about Y_P from X.

3.1 Investigating Trade-Offs between Fairness and Privacy

There are a number of potential methods that can remove information about an attribute from features, but a consistently used technique is using adversaries to remove information. Consequently, when training models, the corresponding optimization is:

$$\min_{\theta_C, \theta_F} \max_{\theta_A, \theta_P} \mathbb{CE}(Y, C(F(X))) - \alpha\mathbb{CE}(Y_A, A(F(X), Y)) - \beta\mathbb{CE}(Y_P, P(F(X), Y)) \quad (1)$$

where $\mathbb{CE}$ denotes the cross entropy and both adversaries are conditional on Y as we are targeting fairness criteria that are conditional on Y and for privacy as the attacker can easily acquire Y from the features.

For each dataset, as shown in Figure 1, we then perform a grid search over α and β, with each hyperparameter either zero or going from 10^{-2} to 10 logarithmically in 10 steps, and evaluate the corresponding effects on utility, fairness and privacy.

Table 1: Results on single metrics on different datasets. We show the baseline metrics without any intervention, where the specific metric for utility and fairness is shown in parentheses for those columns. For privacy, the percentage in parentheses is when the attack accuracy is no better than chance.

Dataset	Baseline Utility	Baseline Fairness	Baseline Privacy	Best Utility	Best Fairness	Best Privacy
CelebA	74.26% (Acc.)	9.27% (Acc. Gap)	82.54% (50%)	76.26%	4.11%	70.33%
CelebA FR	85.69% (TPR)	7.53% (TPR Gap)	69.92% (50%)	85.69%	1.87%	65.38%
EyePACs	63.33% (TPR)	21.00% (TPR Gap)	79.83% (33%)	70.00%	8.00%	67.77%
CheXpert	62.69% (TPR)	24.95% (TPR Gap)	38.69% (25%)	67.16%	14.57%	33.47%

Table 2: Metrics incorporating multiple metrics. We negate utility when computing correlations on utility to match direction of improvement. For CSR metrics, letters in parantheses denote which grouping of (α, β) the ranking belongs to.

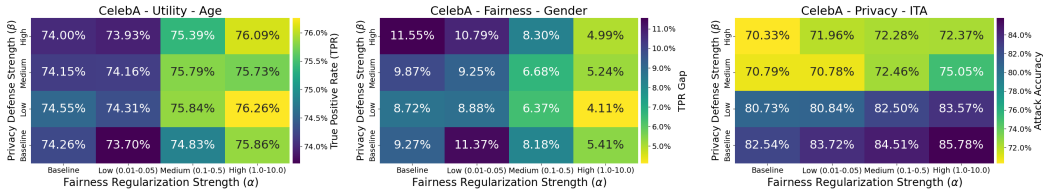| Dataset	U./F. Corr.	U./P. Corr.	F./P. Corr.	CSR$(0.6, 0.2, 0.2)$	CSR$(0.2, 0.6, 0.2)$	CSR$(0.2, 0.2, 0.6)$
CelebA	0.50	-0.01	-0.14	91.04% (H., H.)	88.98% (H., H.)	88.39% (H., H.)
CelebA FR	-0.92	-0.08	0.15	77.67% (L., M.)	75.11% (H., M.)	77.64% (M., M.)
EyePACs	0.58	-0.25	-0.19	76.00% (H., M.)	92.00% (H., M.)	92.00% (H., M.)
CheXpert	-0.02	0.07	0.03	87.71% (M., M.)	82.42% (H., L.)	87.71% (M., M.)

Figure 2: Heatmaps on CelebA over the different grouped regularization strengths for α and β.

Metrics: We utilize several metrics in this study, primarily for evaluating area of concern as well as evaluating a combination of these metrics. For utility, we use either the overall accuracy or true positive rate. For fairness, to expand into cases where Y_A is not binary, we take fairness to be the maximum pairwise absolute difference of the utility metric chosen over the subpopulations defined by Y_A. For privacy, to match the threat model, we train an separate adversary, a linear model in this work, on features extracted from the validation dataset, mimicking the access to the labeled public data. The linear model is trained using loss re-weighting to ensure the adversary does not be a constant prediction. The hidden data is consequently the test dataset used, and the metric for privacy is the balanced accuracy for reasons noted in the previous sentence.

For metrics evaluating the combination of metrics, we use pairwise correlation to measure the linear relationship between different metrics. To incorporate preferences that are not strictly pairwise, we also introduce a metric called Conjunctive Soft Ranking (CSR) to rank the models. CSR is effectively the convex combination between the normalized metrics for each area, normalized so that the worst model over α and β is 0% and the best is 100%.

$$N(M_{\alpha,\beta}) = \frac{M_{\alpha,\beta} - \min_{\alpha',\beta'} M_{\alpha',\beta'}}{\max_{\alpha',\beta'} M_{\alpha',\beta'} - \min_{\alpha',\beta'} M_{\alpha',\beta'}} \tag{2}$$

$$CSR_{\alpha,\beta}(\gamma_U, \gamma_A, \gamma_P) = 100(\gamma_U N(M_{\alpha,\beta}^U) + \gamma_A(1 - N(M_{\alpha,\beta}^A)) + \gamma_P(1 - N(M_{\alpha,\beta}^P))) \tag{3}$$

where γ_U, γ_A, and γ_P sum to one, and $M_{\alpha,\beta}$ denotes the metric for the model trained using α and β, with the superscript denoting which area the metric is measuring. To probe these rankings and see how the best α and β changes as the preference changes, we chose preferences focusing on a primary metric (using a weight of 0.6) and weighing the other two equally (using a weight of 0.2 each).

4 Experiments

We evaluate on CelebA [23] on both age classification and facial recognition tasks as well as EyePACs [13] and CheXpert [19] on disease classification. Age classification on CelebA uses accuracy and
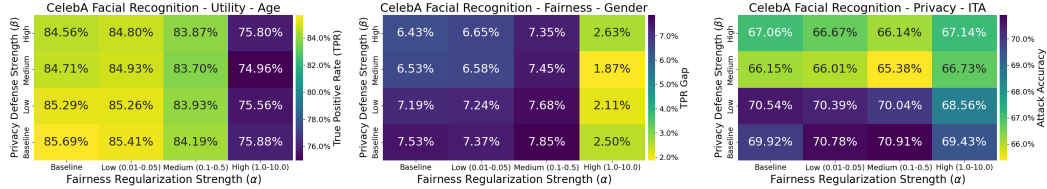
Figure 3: Heatmaps on CelebA for facial recognition over the different grouped regularization strengths for α and β.

accuracy gap as its utility and fairness metrics respectively, while true positive rate and true positive rate gap are used for the remaining tasks. For facial recognition, the true positive rate is calibrated so that there is a false positive rate of 10^{-3} for both subpopulations. For classification tasks, we use a ResNet50 pretrained on Imagenet as our model, and adversaries are three layer multi-layer perceptrons MLP with ReLU activations. For facial recognition, we instead use the same architecture for both the model and adversaries as [12] in using a frozen pretrained ArcFace model and appending a learnable MLP. The model takes in a image of 224 by 224 pixels for all datasets except CheXpert, which matches [19] in using a resolution of 320 by 320 pixels. For training, we use alternating minimization, switching the optimization every batch for classification and every seventy batches for facial recognition. For model selection, we used the model with the best validation loss. In the grid search we perform, we train each set of α and β with three different random seedings, and for visualization, we create a heatmap for each metric and dataset, grouping α or β into the categories of Baseline (0, B.), Low ([0.01, 0.05], L.), Medium ([0.1-0.5], M.), and High ([1.0-10.0], H.) and taking the median in each group for use in the heatmap.

Datasets: CelebA is a large facial imagery datasets consisting of celebrities and providing a large number of potential attributes ranging from age to gender. We evaluate fairness with respect to gender, and use a heuristic called Individual Typology Angle (ITA), thresholded to produce a binary attribute, as a surrogate for skin color for privacy. For age classification, the task is to predict age, and for facial recognition, the task is match the identity from a database of stored features, a particular notable case where attribute privacy is important. EyePACs is a retinal imagery dataset where the task is to determine if a particular fundus photo is referable for diabetic retinopathy, where fairness is evaluated with respect to ITA again, and privacy is with respect to the quality of the fundus photo as described in [14], a surrogate for locations where imaging capabilities are less developed. CheXpert is a dataset of chest X-rays with a number of different disease classifications, of which we take predicting pleural effusion as the task. and fairness is with respect to age and privacy with respect to race.

ITA for both datasets is computed as in the procedure detailed in [27], and both datasets for classification use training splits that exacerbate fairness by balance the task label while undersampling the positive target and sensitive attributes, while the testing split are balanced across the trio of target, sensitive, and private attributes. Splits for facial recognition are partitioned by the identity without controlling for balancing, while CheXpert reuses the provided splits.

Discussion: Tables 1 and 2 and Figures 2, 3, 4, and 5 detail our results. Starting with Table 1, we see that we are able to successfully improve both fairness and privacy taken over our sweep. Improvements in utility are more muted, with medical datasets seeing larger improvements compared with the tasks on CelebA. For Table 2, the correlations taken between pairs of metrics are typically strongest between utility (U.) and fairness (U.), which is not inherently surprising given that they both deal with performance. More interesting is that only facial recognition has a strong negative correlation, where improved fairness means decreased overall performance, while classification tasks are either moderately positive or near independent. Part of this may be due to test balancing, but even for the unbalanced CheXpert test split, none of the models had worse utility than the baseline, as seen in the heatmap. Correlation between utility and privacy (P.) was more muted, with EyePACs having the highest magnitude. One caveat here, underscoring the importance of looking at all three metrics concurrently, is that for EyePACs the utility never decreased when intervening on privacy, instead increasing more with a higher α compared with a higher β. For the correlations between fairness and privacy, the intuition regarding the sign of the correlation matches more closely with how the trends behave, with datasets having negative correlations focusing the most on privacy typically has the worst fairness and vice versa. Finally, for the CSR metrics, we see that typically the same model is
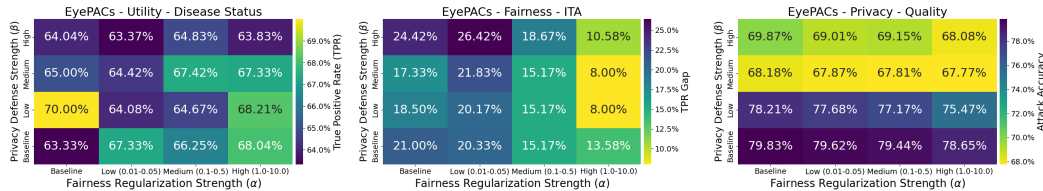
5

Figure 4: Heatmaps on EyePACs over the different grouped regularization strengths for α and β.
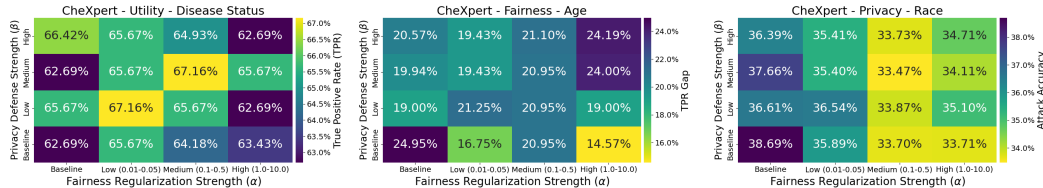

Figure 5: Heatmaps on CheXpert over the different grouped regularization strengths for α and β.

usually sufficient across the chosen preferences, with the facial recognition having the most variation, though better individual metrics can be obtained.

Regarding the figures, looking at the utility heatmap, we see that the facial recognition is the standout, with utility slightly decreasing with higher β, but much more significantly with higher α as noted by the correlation above. The rest of the datasets typically see little change (CelebA classification) or improvements in utility. Part of this may be due to the slightly different training procedure, though both gender and skin color are commonly considered significantly identifying features and in turn may have a more fundamental relationship to the task at hand compared to the classification based tasks. For the fairness heatmap, as one would expect we see improvements in fairness most strongly with higher levels of α. Interesting items of note are how, without intervening on fairness directly, fairness is worsened at the highest levels of β compared to the baseline for CelebA and EyePACs while CelebA FR and CheXpert have slightly improved fairness solely intervening on privacy, though EyePACs and CelebA see improvements at more moderate levels of β. Additionally, though intervening on both fairness and privacy for CheXpert is worse than only intervening on fairness, we see the opposite for the other datasets with the best fairness model using both interventions, suggesting that intervening on multiple attributes can be beneficial for fairness. Finally for the privacy heatmaps, we see the strongest improvements by increasing β, though, as with other cases, CheXpert is an outlier in that intervening on fairness typically produces stronger effects than intervening on privacy. We also see that the best privacy is usually acquired on a combination of both fairness and privacy, except for CelebA where the influence of the fairness regularization worsens privacy, forming a gradient from the bottom right most cell to the top left most cell.

5   Conclusion

This study investigates for the first time the interaction and tradeoffs between attribute privacy, fairness, and utility in computer vision and show that these interactions are more subtle and complex than straight inverse tradeoffs. The tradeoffs depend heavily on the data, with fairness and privacy at times being at odds with each other or mutually cooperative. Also, our experiments includes examples where models with the best fairness usually are some combination of addressing both fairness and privacy, suggesting to practitioners that intervening on multiple attributes can be beneficial. Potential next steps are to more tightly address these separate concerns, as well as more explicitly include given preferences in the training process.

References

[1] Alvi, M., Zisserman, A., and Nellåker, C. (2018). Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*.

[2] Asoodeh, S., Alajaji, F., and Linder, T. (2015). Notes on information-theoretic privacy.

[3] Barbano, C. A., Tartaglione, E., and Grangetto, M. (2021). Bridging the gap between debiasing and privacy for deep learning. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 3799–3808.

[4] Beutel, A., Chen, J., Zhao, Z., and Chi, E. H. (2017). Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*.

[5] Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing systems*, pages 4349–4357.

[6] Burlina, P., Joshi, N., Paul, W., Pacheco, K. D., and Bressler, N. M. (2020a). Addressing artificial intelligence bias in retinal disease diagnostics. *arXiv preprint arXiv:2004.13515 and Translational Vision Science and Technology (accepted)*.

[7] Burlina, P., Paul, W., Mathew, P., Joshi, N., Pacheco, K. D., and Bressler, N. M. (2020b). Low-shot deep learning of diabetic retinopathy with potential applications to address artificial intelligence bias in retinal diagnostics and rare ophthalmic diseases. *JAMA Ophthalmology*.

[8] Burlina, P. M., Joshi, N., Pacheco, K. D., Liu, T. A., and Bressler, N. M. (2019a). Assessment of deep generative models for high-resolution synthetic retinal image generation of age-related macular degeneration. *JAMA Ophthalmology*, 137(3):258–264.

[9] Burlina, P. M., Joshi, N. J., Ng, E., Billings, S. D., Rebman, A. W., and Aucott, J. N. (2019b). Automated detection of erythema migrans and other confounding skin lesions via deep learning. *Computers in biology and medicine*, 105:151–156.

[10] Carlini, N. and Wagner, D. (2017). Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 3–14.

[11] Caton, S. and Haas, C. (2020). Fairness in machine learning: A survey.

[12] Dhar, P., Gleason, J., Roy, A., Castillo, C. D., and Chellappa, R. (2021). Pass: Protected attribute suppression system for mitigating bias in face recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15087–15096.

[13] EyePACS (2015). Diabetic retinopathy detection. data retrieved from Kaggle, https://www.kaggle.com/c/diabetic-retinopathy-detection.

[14] Fu, H., Wang, B., Shen, J., Cui, S., Xu, Y., Liu, J., and Shao, L. (2019). Evaluation of retinal image quality assessment networks in different color-spaces. In *Lecture Notes in Computer Science*, pages 48–56. Springer International Publishing.

[15] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014a). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680.

[16] Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014b). Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.

[17] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

[18] Hui, B., Yang, Y., Yuan, H., Burlina, P., Gong, N. Z., and Cao, Y. (2021). Practical blind membership inference attack via differential comparisons. *arXiv preprint arXiv:2101.01341*.

[19] Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., Seekins, J., Mong, D. A., Halabi, S. S., Sandberg, J. K., Jones, R., Larson, D. B., Langlotz, C. P., Patel, B. N., Lungren, M. P., and Ng, A. Y. (2019). Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison.

[20] Jagielski, M., Kearns, M., Mao, J., Oprea, A., Roth, A., Sharifi-Malvajerdi, S., and Ullman, J. (2019). Differentially private fair learning. In *International Conference on Machine Learning*, pages 3000–3008. PMLR.

[21] Kinyanjui, N. M., Odonga, T., Cintas, C., Codella, N. C., Panda, R., Sattigeri, P., and Varshney, K. R. (2020). Fairness of classifiers across skin tones in dermatology. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 320–329. Springer.

[22] LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436.

[23] Liu, Z., Luo, P., Wang, X., and Tang, X. (2015). Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.

[24] Makhdoumi, A., Salamatian, S., Fawaz, N., and Médard, M. (2014). From the information bottleneck to the privacy funnel. In *2014 IEEE Information Theory Workshop (ITW 2014)*, pages 501–505. IEEE.

[25] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529.

[26] Paul, W., Cao, Y., Zhang, M., and Burlina, P. (2021). Defending medical image diagnostics against privacy attacks using generative methods. *arXiv preprint arXiv:2103.03078*.

[27] Paul, W., Hadzic, A., Joshi, N., Alajaji, F., and Burlina, P. (2020). Tara: Training and representation alteration for ai fairness and domain generalization. *arXiv preprint arXiv:2012.06387*.

[28] Redmon, J., Divvala, S. K., Girshick, R. B., and Farhadi, A. (2015). You only look once: Unified, real-time object detection. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788.

[29] Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:1137–1149.

[30] Robinson, J. P., Qin, C., Henon, Y., Timoner, S., and Fu, Y. (2021). Balancing biases and preserving privacy on balanced faces in the wild.

[31] Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.

[32] Shokri, R., Stronati, M., Song, C., and Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE.

[33] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *nature*, 550(7676):354–359.

[34] Song, J., Kalluri, P., Grover, A., Zhao, S., and Ermon, S. (2019). Learning controllable fair representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2164–2173.

[35] Wadsworth, C., Vera, F., and Piech, C. (2018). Achieving fairness through adversarial learning: an application to recidivism prediction. In *Conference on Fairness, Accountability, and Transparency in Machine Learning (FATML)*.

[36] Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. (2013). Learning fair representations. In *International Conference on Machine Learning*, pages 325–333.

[37] Zhang, B. H., Lemoine, B., and Mitchell, M. (2018). Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340.