Unravelling Robustness of Deep Face Recognition Networks Against Illicit Drug Abuse Images

Hruturaj Dhake and Akshay Agarwal
Data Science and Engineering
IISER Bhopal, India
{hruturaj19, akagarwal}@iiserb.ac.in

Abstract

Alteration in facial features can lead to a significant drop in recognition performance. These alterations can be due to several factors: one such prominent and less explored factor is illicit drug abuse. To advance the understanding of how drug abuse faces affect the performance of state-of-the-art deep face recognition (DFR) networks, in this study, we have utilized clean and illicit drug abuse faces. Extensive studies are performed on deep face recognition and soft biometric identification, such as gender, ethnicity, and expression recognition. It is observed that illicit drug abuse not only impacts the identity recognition performance but also degrades the soft biometrics identification accuracy. Therefore, to advance the integrity of DFR, we have performed the detection of illicit drug abuse as a potential solution to its mitigation. In the end, the robustness of the drug abuse face detector is evaluated under the prominent use of social-media filters on face images.

1. Introduction

Face recognition technology is widely used in modern society and a variety of industries, such as digital payment services, social networking platforms, and security systems. Its widespread adoption is attributed to its ability to swiftly authenticate and identify individuals, offering unparalleled convenience and security. However, external factors may compromise the effectiveness of face recognition systems, increasing the risk of inaccurate outcomes and denial of service [4, 14, 16, 20, 27]. One such prominent factor that can affect the performance of face recognition networks is illicit drug abuse. Drug abuse can cause asymmetry in lips and mandibular structure, increasing susceptibility to infections and face swelling. Other visible effects include dark circles, severe runny nose, irritations, swelling, sores, and itching, resembling acne and wounds [11]. Considering the increasing prevalence of face recognition technology in our daily lives, denying service based on facial alterations induced by illicit drug abuse, even to individuals undergoing rehabilitation, raises ethical and practical concerns. Although current facial recognition algorithms have achieved remarkable performance levels [28], even surpassing human-level performance [18], they remain highly vulnerable to such external factors [1, 19, 24, 31]. Given these factors, further research on the effects of drug abuse on facial recognition technologies is crucial which is still limited in the literature. Hence, this paper for the first time presents a thorough examination of the visible effects of drug abuse on facial features and their implications for face recognition and soft biometric attribute prediction. We assert that an evaluation can deepen our understanding of the vulnerabilities of the current deep face recognition (DFR) algorithms. Based on the analysis obtained on the sensitivity of DFRs, detection of drug abuse faces, and its challenging mitigation through traditional image enhancement methods, we demand sophisticated networks to address this challenge. This advancement can ensure the continued acceptability and reliability of facial recognition technology across diverse societal contexts.

2. Related Work

Interestingly, a limited amount of work has been done to understand the impact of illicit drug abuse on face recognition. The current literature is divided into two forms of studies: (i) understanding the vulnerability of traditional face recognition algorithms and (ii) classification of a subject as a drug addict or non-addict, *e.g.* Raghavendra *et al.* [24] have studied the impact of drug abuse on hand-crafted feature-based face recognition algorithms. The authors have used the Local Binary Patterns (LBP), and Binarised Statistical Image Features (BSIF) and the recognition has been performed using Sparse Representation Classifier (SRC). On a similar line, Yadav *et al.* [31] have performed a preliminary analysis on the effect of drug abuse on commercial and handcrafted feature-based face recognition algorithms. Pandey *et al.* [22] have proposed a transformation-based approach utilizing a scattering wavelet network to improve face recognition performance under the influence of drug abuse. Gnansekar *et al*. [12] and Gupta *et al*. [15] have proposed the framework for the classification of subjects as addict or nonaddict.

As can be seen, the existing literature does not adequately tackle the issue of illicit drug abuse including the vulnerability of state-of-the-art (SOTA) deep face recognition networks such as ArcFace [9] and FaceNet [25]. The prime reason for such is the unavailability of illicit drug abuse datasets. While few works performed the preliminary study in understanding the effect of illicit drug abuse on handcrafted feature-based face recognition algorithms and commercial algorithms, there is no study on whether the prediction of soft biometric attributes will also be affected due to the impact of illicit drug abuse. To address these limitations, we have performed a comprehensive study by analyzing state-of-the-art (SOTA) deep face recognition networks for face identification and soft attribute prediction. Subsequently, we assert that if we can identify whether the face image in question is clean or drug-addicted, we can apply an image enhancement to mitigate the impact of drugs on drug-addicted faces to boost recognition performance. In the end, we have proposed a drug abuse face detector using pre-trained convolutional neural networks (CNNs). The motivation of the above experiment can be observed from existing studies, *e.g*. [2, 6] tackle the issue of camera interoperability by first identifying the camera and later applying the camera-based selected enhancement on iris images to boost the performance. Inspired by this multi-level defense strategy, for the first time we have performed a preliminary study to evaluate whether traditional image enhancement techniques can mitigate the impact of illicit drug abuse. It is observed that general image enhancement can not mitigate the impact of drug abuse and hence demands the development of sophisticated enhancement techniques to mitigate the impact of illicit drug abuse.

3. Impact of Drug Abuse on Face Recognition

In this research, we studied the impact of illicit drug abuse on face identification and soft biometrics attribute prediction. First, we present a brief overview of the drug abuse face dataset. Later, the experimental results and analysis concerning face identification and soft attribute prediction have been provided.

3.1. Drug Abuse Faces (DAF)

To overcome the limitation of the illegal drug datasets, recently, Dhake and Agarwal [10] have proposed a dataset 'Drug Abuse Faces' containing two classes: 'after', featuring images of faces post-drug consumption, and 'before', consisting of images of faces before drug consumption, presenting a clean baseline unaffected by illicit drugs. Overall, the dataset comprised 230 images from 115 subjects. The illustration provided in Figure 1 showcases samples depicting faces affected by drug abuse.

Figure 1. Visualization of a few samples from Drug Abuse Faces showcasing 'before' and 'after' class images. The faces will be blurred and/or eye strips will be added later to protect the privacy of the users.

3.2. Effect of Deep Face Recognition

Our research endeavors to address a drawback of existing studies that took an effort to understand the vulnerability of traditional face recognition networks but left space to examine deep face recognition networks. Therefore, to evaluate deep face recognition under the influence of drugs, we have chosen several SOTA deep face recognition networks including VGG-Face (VGG) [23], ArcFace (AF) [9], FaceNet (F-Net) [25], OpenFace (OF) [7], FaceNet512 (F-Net512) [13], SFace (SF), and DeepID [21]. By scrutinizing their sensitivity, our study aims to identify vulnerabilities and performance variations across architectures, contributing to the overarching goal of enhancing the dependability and precision of facial recognition systems. Since the propossed dataset contains one before and after drug abuse image of each subject, we have utilized the Style-Based Regression Model (SAM) [5] to generate a third image from the clean image and use that image as a gallery image for face identification. The primary reason for such unavailability of images is the sensitivity of the topic and privacy of different individuals[1].

It is observed that the VGG-Face model yields the highest clean face recognition accuracy of 91.30% surpassing each of the deep face models used. However, the other state-of-the-art (SOTA) face recognition models such as ArcFace (AF) and FaceNet (F-Net) models also yield significantly high clean face accuracies of 89.57% and 84.35%, respectively. Interestingly, a discernible decline in accuracy is observed across each model when analyzing post-drug abuse images, *e.g*. VGG-Face which yields the best clean image face recognition accuracy suffers a drop of 17.39%. Similarly, FaceNet suffers a drop from 84.35% to 33.04%. This decline in recognition performance underscores the potential impact of drug abuse on facial recognition systems, emphasizing the need for adequate address of this underexplored topic. The detailed results highlighting the vulnerability of different deep face recognition models are reported in Table 1.

Table 1. Comparative evaluation of face recognition accuracy (%) of deep models using pre and post-drug abuse images.

Model	VGG	AF	F-Net	OF	F-Net512	SF	DeepID
Before	**91.3**	89.6	84.3	30.4	66.9	84.3	81.7
After	**73.9**	58.3	33.0	0.0	13.0	39.1	16.5

3.2.1 Robustness Under Social Media Filters

It is well known that in the social media age, the majority of facial images undergo social media filters to enhance their visual appearance[2]. Therefore, in the first-ever study we have studied the impact of multiple 'Instagram' filters on both clean and drug abuse faces to analyze the impact of social media filtering on deep face recognition networks. It is interesting to note that, filtering does not yield a significant impact on the clean face recognition performance on VGG-Face and ArcFace. The performance of these 2 networks degrades in the range of $0 - 2.6\%$ across different filter types. While F-Net also reflects the high robustness on Inkwell, Moon, and Xpro2 filter, its performance on clean images drops to 8.7%. Interestingly, each network is found highly sensitive when Instagram filters drug abuse faces come for recognition, *e.g*. the accuracy of SOTA VGG drops from 73.9% on original drug abuse images to 33.0% when the Xpro2 filter is applied to the drug abuse faces. This extensive experimental analysis shown in Table 2 demonstrates that not only the effect of the drug needs to be dealt with but also social media filtering can further increase the sensitivity of each deep face recognition network.

3.3. Soft Biometric Identification

Our above analysis demonstrates that the deep face models are highly sensitive to illicit drug abuse faces for identification. In this section, we moved one step further to evaluate their sensitivity to soft biometrics such as gender, ethnicity, and expression classification. First, we present the analysis of gender identification followed by the analysis of ethnicity identification. In the end, the results and sensitivity discussion around expression classification are provided.

3.3.1 Impact on Gender Detection

We have used the VGG-Face model which yields the SOTA performance above for face identification. It is observed that the VGG-Face is not only found effective for identification but also gender classification and yields an accuracy of 83% when clean images are provided for analysis. However, similar to identification, a significant drop in gender classification is observed. On drug abuse, a gender classification accuracy of 64% is observed which is 19% lower than on clean images. Table 3 provides detailed results of gender classification in terms of accuracy, precision, recall, and F-1 score for gender detection. The confusion matrix shown in Table 4 demonstrates an interesting fact that female images observed the misclassification while male gender remains intact, *reflecting the high chance of biases in gender detection under the influence of drug abuse.*

3.3.2 Impact on Ethnicity Detection

To further strengthen the quality of the research, we have now studied the impact of drugs on ethnicity detection. It is observed that the majority of the clean images are classified as 'White' race followed by the Latino Hispanic. Surprisingly a significant drop in the detection of the 'White' race is observed, and no impact of drug abuse on Latino Hispanics and Indian races is noticed. Interestingly, few races such as Asian, Middle Eastern, and Black have observed an increase in the number of images classified under the influence of drug abuse. The findings, detailed in Figure(s) 2 and 3, reveal variations in ethnicity data, with 24.35% of instances demonstrating changes.

3.3.3 Impact on Expression Detection

So far, it has been observed that deep face models are sensitive not only to face identification but also to gender and ethnicity/race classification. To further demonstrate whether

[1]In the future, we aim to utilize diffusion models to generate synthetic drug abuse faces.

[2]https://kaptur.co/statistics-how-filters-are-used-by-instagrams-most-successful-users/

Table 2. Sensitivity of deep face recognition models under the influence of social media filters. Both clean and drug abuse faces are filtered with several popular Instagram filters namely Inkwell, Moon, Xpro2, and Earlybird.

Filter	Original		Social Media Filters							
			Inkwell		Moon		xpro2		Earlybird	
Model	Before	After	Before	After	Before	After	Before	After	Before	After
VGG	**91.3**	73.9	90.4	64.3	91.3	68.7	91.3	**33.0**	89.6	66.1
AF	**89.6**	58.3	87.0	56.5	87.0	54.8	87.0	**13.0**	87.8	53.04
F-Net	**84.3**	33.0	82.6	35.6	83.5	31.3	80.0	**5.2**	75.6	24.3
OF	30.4	0.0	13.0	0.0	14.8	0.0	25.2	0.0	18.3	0.9
FNet512	67.0	13.0	57.4	12.2	58.3	12.2	61.7	**2.6**	53.0	10.4
SF	84.3	39.1	81.7	30.4	82.6	30.4	81.7	2.6	78.3	32.2
DeepID	**81.7**	16.5	42.6	**0.9**	60.9	5.22	51.3	20.0	67.8	20.9

Table 3. Gender classification performance on clean and drug abuse faces highlighting the sensitivity of deep face model in the soft biometrics classification.

Images	Precision	Recall	F-1 Score	Accuracy
Before	**0.83**	0.87	0.83	0.83
After	**0.73**	0.72	0.64	0.64

Table 4. Confusion matrix of gender classification for clean and drug abuse faces.

Predicted →	Before		After	
True ↓	Male	Female	Male	Female
Male	39	1	39	1
Female	18	57	40	35

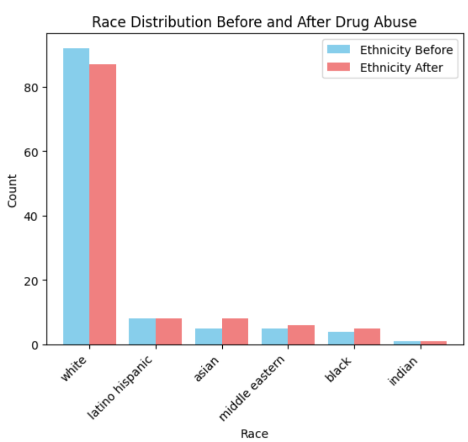

Figure 2. Impact of drug abuse on ethnicity recognition: Comparison of ethnicity detection results before and after substance abuse.

these models are sensitive in handling any other soft biometric modality, we performed experiments for expression detection. Figure 4 demonstrates a decline in occurrences of

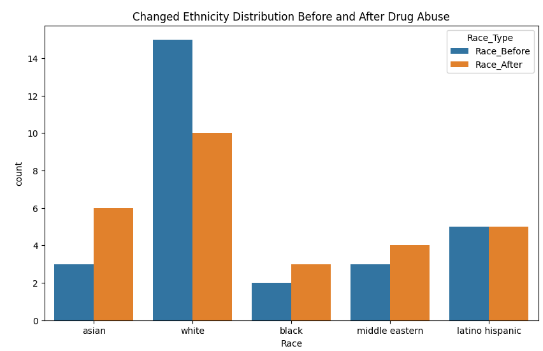

Figure 3. Variations in Ethnicity Detection Before and After Drug Abuse: An Analysis of Changes in Ethnicity Data

neutral facial expressions, accompanied by a significant increase in expressions of sadness and anger after drug abuse, suggesting a shift towards more negative emotional states. The increasing prevalence of the 'disgust' category highlights the intricate nature of emotional changes linked to drug abuse.

3.4. Mitigation of Drug Abuse and Noises

We assert that since impact of illicit drugs leaves a significant impact on the quality of faces which can be in the form of scars and dark spots. Further, the dataset used is collected from unconstrained (*'in-the-wild'*) online domains having issues with image quality such as noisy and blurred images. Therefore, one question that can be asked is whether image processing can mitigate the impact of drug abuse and help improve the performance of deep face recognition models. For that, we have applied several image enhancement techniques such as median blur, Gaussian blur, and morphological image opening, each followed by image sharpening. The results reported in Table 5 showcase that while VGG-Face yields the highest accuracy, it does not show any improvement in the recognition performance after image enhancement. Nevertheless, a few face recognition models, such as

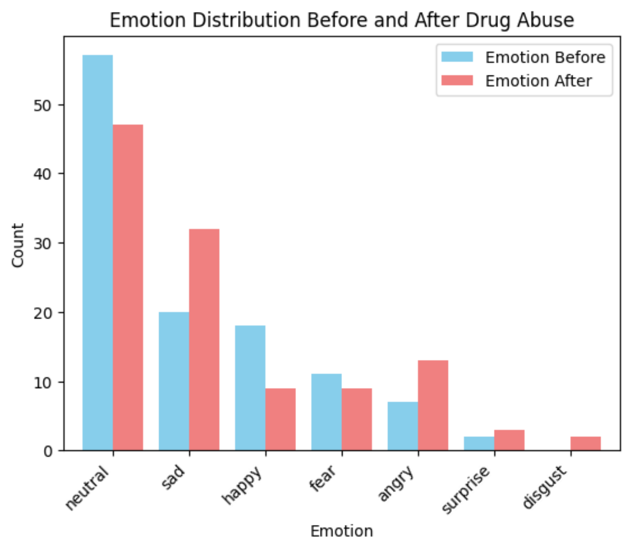

Figure 4. Detected emotion analysis which provides insights that drug abuse leads to increased sadness and angry emotions

Table 5. Impact of image processing techniques on mitigating the impact of drug abuse on face recognition.

Technique	VGG	AF	F-Net	OF	F-Net512	SF	DeepID
Original	68.7	50.4	30.4	0.9	10.4	35.6	13.9
Median Blur	65.2	53.9	28.7	1.7	11.3	35.6	16.5
Guassian Blur	68.7	55.6	29.6	1.7	11.3	39.1	17.4
Image Opening	65.2	53.0	27.8	0.9	13.9	23.5	11.3

ArcFace (AF) and SFace (SF) show a boost of 5.2% and 3.5%, respectively reflecting a potential that sophisticated image filtering can degrade the impact of drug abuse factors. Out of all the processing techniques, Gaussian blur is found the most effective processing medium to mitigate the impact of drug abuse. It is to be noted here that while existing studies have evaluated that drug abuse affects the recognition performance of handcrafted features and commercial face recognition, no study evaluated the impact on soft biometrics prediction and analyzed the robustness of image processing techniques to mitigate the impact of drug abuse face alterations.

4. Detection of Drug Abuse Face Identification

In another potential solution to mitigate the impact of drug abuse, we have performed experiments to detect whether a face in question is a drug abuse-altered face or a clean face. We assert that this detection can help in first segregating the clean faces from the drug-altered faces and later the detected altered images can only be used for specific enhancement to boost the recognition performance. It will save the cost of applying preprocessing by only enhancing the drug-altered faces and can act as a tradeoff between accuracy and computational cost.

Table 6. Performance of traditional machine learning classifiers on drug vs. clean image classification using image pixels.

Classifier	Precision	Recall	F1 Score	Accuracy	AUC-ROC
SVM	0.42	0.32	0.43	0.36	0.49
Decision Tree	0.58	0.28	0.38	0.54	0.54
KNN	0.33	0.32	0.33	0.34	0.30
Random Forest	0.52	0.44	0.48	0.52	0.46
Logistic Regression	0.42	0.32	0.36	0.44	0.49

We have utilized both traditional machine learning approaches and deep convolutional neural networks to perform binary classification: drug vs. non-drug (or clean). Since it is observed that the effect of illicit drugs brings distortion in facial features, therefore, we believe image pixels might have cues to filter out the drug abuse images. Based on this assertion, we have performed the binary classification utilizing the image pixels coupled with traditional but powerful machine learning classifiers such as support vector machine (SVM) and decision trees. The dataset is first divided into training and testing sets where the training set contains 180 images of both classes and the remaining 50 images are used for evaluation. For a thorough evaluation of the performance, we have used several metrics including accuracy, precision, recall, F1 score, and AUC-ROC score.

4.1. Detection using Traditional Methods

The results of binary classification using image pixels coupled with machine learning classifiers are reported in Table 6. Out of all the classifiers used, the decision tree yield the highest performance. However, it is observed that the classifier yields poor recall but has a higher precision value leading to an overall lower F-1 score. Whereas the random forest classifier shows slightly lower accuracy but yields comparable precision and recall values. Overall it can be mentioned that image pixels are not sufficient to detect drug abuse images. The prime reason can be understood that images are acquired from Internet sources (**in-the-wild**), and hence clean images can exhibit similar distortion present in drug abuse images. Further, environmental factors, including image acquisition, reduce the inter-class separation leading to poor classification performance.

To mitigate the impact of image pixels on drug abuse face classification, we employ a hybrid approach where we utilize state-of-the-art (SOTA) convolutional neural networks (CNN) for feature extraction and traditional classifiers for image classification. Further, it is observed that the dimensionality of extracted features is significantly high as compared to the number of training images; therefore, to avoid a potential issue of overfitting, we performed dimensionality reduction using principal component analysis (PCA). The number of components (n) in PCA is varied in the range from 50 to 150 to assess the impact on classification performance. Table 7 presents classification results for

Table 7. Effectiveness of CNN features coupled with machine learning classifiers for drug abuse face detection.

Model	PCA n=50	PCA n=100	PCA n=150
VGG 16 + SVM	0.74	0.70	0.72
Xception + SVM	0.70	0.68	0.70
DenseNet + SVM	0.66	0.64	0.66
InceptionV3 + SVM	0.70	0.78	0.78
VGG 16 + DT	0.66	0.58	0.70
Xception + DT	0.56	0.56	0.64
DenseNet + DT	0.68	0.62	0.50
InceptionV3 + DT	0.64	0.60	0.64
VGG 16 + RF	**0.78**	**0.72**	**0.76**
Xception + RF	0.68	0.68	0.64
DenseNet + RF	0.66	0.68	0.60
InceptionV3 + RF	0.76	0.66	0.62

Table 8. Drug abuse face classification using fine-tuned CNNs.

CNN	Precision	Recall	F1 Score	Accuracy	AUC-ROC
VGG 16	0.75	0.60	0.67	0.70	0.79
Xception	0.70	0.64	0.67	0.68	0.70
InceptionV3	0.79	0.76	0.78	0.78	0.88
DenseNet	**0.81**	**0.84**	**0.82**	**0.82**	**0.88**

various CNN models combined with machine learning classifiers and reduced PCA configurations. VGG16 combined with random forest achieved the highest accuracy across all classifiers and CNNs. It is to be noted that the effectiveness of VGG + random forest (RF) compared to other models is agnostic to different reduced feature dimensional space.

4.2. Drug Abuse Detection using CNN

Inspired by the success of CNNs as a feature extractor for drug abuse face detection, we have fine-tuned these pre-trained state-of-the-art CNNs for drug image classification. For that, a trainable dense layer is added and all other layers of the pre-trained CNN are kept frozen to balance the trade-off between accuracy and the computational cost. The results reported in Table 8 showcase that the DenseNet [17] yields state-of-the-art performance surpassing each network by a significant margin, *e.g*. the DenseNet model yields an accuracy of 82% in comparison to 70%, 68%, and 78% achieved by VGG16 [26], Xception [8], and Inception models [29], respectively. It is to be noted that the performance surpasses the accuracy obtained when the CNNs are used as a feature extractor. When we analyze the performance of individual networks on specific classes, it is observed that each classifier is effective in handling the before-drug abuse class (original/clean images) but is less effective in the after-drug abuse class except DenseNet.

Further, since face images are one of the most popular mediums of communication they have been extensively uploaded on social media platforms where enhancing the quality of face images is a popular practice. As shown earlier, these filter-based enhancements drastically impact the performance of face recognition networks [3, 30]. Therefore, before assuring the effectiveness of DenseNet in detecting drug abuse faces, we have evaluated its effectiveness against social-media-based filter enhancement. *For that, we applied* 4 *popular Instagram filters to the test set of the proposed dataset and generated* 200 *images in total. 'The classifier trained on original images is used to evaluate these unseen enhanced images'.*

Similar to face identification, it is observed that these Instagram filters including 'Inkwell', 'Moon', and 'Xpro2' drastically reduce the detection performance of DenseNet (the best-performing network). For instance, when the Moon filter-based enhancement is applied to face images; the accuracy of the model drops from 82% to 76%. However, interestingly, the filter does not impact the performance of the after-drug face class but affects the classification of the before-drug abuse class. In other words, it increases false negative errors. here we can observe that VGG16 with Random Forest outperformed followed by inceptionV3 with 78% and 76% accuracy.

5. Conclusion and Future Work

The proposed research rigorously addressed the impact of drug abuse-induced alterations on deep face recognition algorithms. Further, for the first time, the study extended beyond recognition failures to highlight the significant impact on gender and race classification. Additionally, we delved into the intricate dynamics of behavioral and emotional changes post-drug abuse. It is observed that drug alteration not only affects the identification performance but significantly affects the detection of several other soft biometrics attributes as well. It is demonstrated that simple image processing can not mitigate the impact of drug-induced alteration demanding the development of a sophisticated image enhancement technique. In the end, we have performed the drug vs. clean image classification; however, the most effective network is found vulnerable to social media filtering-based image enhancement. In the future, we aim to extend the dataset and develop an advanced image enhancement technique to mitigate the impact of drug abuse.

References

[1] Akshay Agarwal, Richa Singh, Mayank Vatsa, and Afzel Noore. Swapped! digital face presentation attack detection via weighted local magnitude pattern. In *IEEE International Joint Conference on Biometrics*, pages 659–665, 2017. 1

[2] Akshay Agarwal, Rohit Keshari, Manya Wadhwa, Mansi Vijh, Chandani Parmar, Richa Singh, and Mayank Vatsa. Iris sensor identification in multi-camera environment. *Information Fusion*, 45:333–345, 2019. 2

[3] Akshay Agarwal, Richa Singh, Mayank Vatsa, and Afzel Noore. Magnet: Detecting digital presentation attacks on

face recognition. *Frontiers in Artificial Intelligence*, 4: 643424, 2021. 6

[4] Akshay Agarwal, Richa Singh, Mayank Vatsa, and Afzel Noore. Boosting face presentation attack detection in multi-spectral videos through score fusion of wavelet partition images. *Frontiers in Big Data*, 5:836749, 2022. 1

[5] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Only a matter of style: Age transformation using a style-based regression model. *ACM Transactions on Graphics (TOG)*, 40 (4):1–12, 2021. 3

[6] Sunpreet S Arora, Mayank Vatsa, Richa Singh, and Anil Jain. On iris camera interoperability. In *IEEE International Conference on Biometrics: Theory, Applications and Systems*, pages 346–352, 2012. 2

[7] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: an open source facial behavior analysis toolkit. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1–10, 2016. 2

[8] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1251–1258, 2017. 6

[9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 2

[10] Hruturaj Dhake and Akshay Agarwal. On the robustness of drug abuse face classification. In *The Second Tiny Papers Track at ICLR 2024*. 2

[11] Thomas E Freese, Karen Miotto, and Cathy J Reback. The effects and consequences of selected club drugs. *Journal of substance abuse treatment*, 23(2):151–156, 2002. 1

[12] Sudarsini Tekkam Gnanasekar and Svetlana Yanushkevich. Face attributes and detection of drug addicts. In *IEEE International Conference on Emerging Security Technologies*, pages 1–6, 2019. 2

[13] Sixue Gong, Vishnu Naresh Boddeti, and Anil K Jain. On the intrinsic dimensionality of image representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3987–3996, 2019. 2

[14] Gaurav Goswami, Akshay Agarwal, Nalini Ratha, Richa Singh, and Mayank Vatsa. Detecting and mitigating adversarial perturbations for robust face recognition. *International Journal of Computer Vision*, 127:719–742, 2019. 1

[15] Tanu Gupta, Meghna Goyal, Goutam Kumar, and Balasubramanian Raman. Deep learning framework for detection of an illicit drug abuser using facial image. In *International Conference on Computer Vision and Image Processing*, pages 11–20, 2020. 2

[16] Mohamad Harastani, Amine Benterkia, Farnaz Majid Zadeh, and Amine Nait-Ali. Methamphetamine drug abuse and addiction: Effects on face asymmetry. *Computers in Biology and Medicine*, 116:103475, 2020. 1

[17] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708, 2017. 6

[18] Chaochao Lu and Xiaoou Tang. Surpassing human-level face verification performance on lfw with gaussianface. In *AAAI conference on artificial intelligence*, 2015. 1

[19] Puspita Majumdar, Akshay Agarwal, Richa Singh, and Mayank Vatsa. Evading face recognition via partial tampering of faces. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 11–20, 2019. 1

[20] Aman Mehra, Akshay Agarwal, Mayank Vatsa, and Richa Singh. Motion magnified 3-d residual-in-dense network for deepfake detection. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 5(1):39–52, 2023. 1

[21] Wanli Ouyang, Xiaogang Wang, Xingyu Zeng, Shi Qiu, Ping Luo, Yonglong Tian, Hongsheng Li, Shuo Yang, Zhe Wang, Chen-Change Loy, et al. Deepid-net: Deformable deep convolutional neural networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2403–2412, 2015. 2

[22] P Pandey, R Singh, and M Vatsa. Face recognition using scattering wavelet under illicit drug abuse variations, in 2016 int. conf. *Biometrics, ICB 2016*, 2016. 1

[23] Omkar Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 2

[24] Ramachandra Raghavendra, Kiran B Raja, and Christoph Busch. Impact of drug abuse on face recognition systems: a preliminary study. In *9th International Conference on Security of Information and Networks*, pages 24–27, 2016. 1

[25] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015. 2

[26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6

[27] Richa Singh, Akshay Agarwal, Maneet Singh, Shruti Nagpal, and Mayank Vatsa. On the robustness of face recognition algorithms against attacks and bias. In *AAAI Conference on Artificial Intelligence*, pages 13583–13589, 2020. 1

[28] Yi Sun, Ding Liang, Xiaogang Wang, and Xiaoou Tang. Deepid3: Face recognition with very deep neural networks. *arXiv preprint arXiv:1502.00873*, 2015. 1

[29] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. 6

[30] Zhe Wu, Zuxuan Wu, Bharat Singh, and Larry Davis. Recognizing instagram filtered images with feature de-stylization. In *AAAI Conference on Artificial Intelligence*, pages 12418–12425, 2020. 6

[31] Daksha Yadav, Naman Kohli, Prateekshit Pandey, Richa Singh, Mayank Vatsa, and Afzel Noore. Effect of illicit drug abuse on face recognition. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–7, 2016. 1