

# TOWARDS EXPLAINABLE FACE AGING WITH GENERATIVE ADVERSARIAL NETWORKS

Angelo Genovese, Vincenzo Piuri, and Fabio Scotti

Department of Computer Science, Università degli Studi di Milano, Italy

## ABSTRACT

Generative Adversarial Networks (GAN) are being increasingly used to perform face aging due to their capabilities of automatically generating highly-realistic synthetic images by using an adversarial model often based on Convolutional Neural Networks (CNN). However, GANs currently represent black box models since it is not known how the CNNs store and process the information learned from data. In this paper, we propose the first method that deals with explaining GANs, by introducing a novel qualitative and quantitative analysis of the inner structure of the model. Similarly to analyzing the common genes in two DNA sequences, we analyze the common filters in two CNNs. We show that the GANs for face aging partially share their parameters with GANs trained for heterogeneous applications and that the aging transformation can be learned using general purpose image databases and a fine-tuning step. Results on public databases confirm the validity of our approach, also enabling future studies on similar models.

**Index Terms**— GAN, Face aging, CNN, Deep Learning


## 1. INTRODUCTION

Face aging (or age progression) techniques consist in modifying face images to simulate the effect of aging. Aging is especially useful when it is not possible to capture an up-to-date image of the individual, for example in cross-age face verification or in finding lost children [1], or to perform data augmentation [2].

Traditionally, methods for face aging are divided in two categories [3]: *i*) prototype approaches and *ii*) physical model approaches. Methods belonging to the first category aim at simulating the aging by computing an average face (prototype) for each age group, extracting a texture difference between the prototypes, and applying the texture difference to the test image to simulate aging [4, 5]. Methods belonging to the second category try to generate realistic aged images by determining a parametric model of the aging process that considers the physical structure of the skin, muscles, and bones [6, 7].

Recently, computational intelligence methods based on Deep Learning (DL) have been increasingly used due to their capabilities of automatically learning data representations [11, 12, 13, 14]. Among DL methods, several approaches to face aging have considered Generative Adversarial Networks (GAN) [15], which have shown high accuracy in automatically generating realistic synthetic images. In particular, GANs represent a type of generative model composed by a generator  $G$ , that generates synthetic data, and a discriminator  $D$ , that classifies the data generated by  $G$  as real or

This work was supported in part by the Italian Ministry of Research within the PRIN project COSMOS (201548C5NT). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research, within the project “Deep Learning and CUDA for advanced and less-constrained biometric systems”.



**Fig. 1.** Result of applying the DEX age estimator [8] on the CIFAR-10 object dataset [9] and the CelebA face dataset [10]. In both cases the age estimator correctly processes the images, since image characteristics influencing apparent age are general and not exclusive to face images.

synthetic. In most recent cases, both  $G$  and  $D$  are implemented as Convolutional Neural Networks (CNN).

It is possible to divide GAN-based methods for face aging in three main categories: *i*) methods based on conditional GANs, *ii*) methods based on conditional GANs and autoencoders, and *iii*) methods based on cycle GANs. Methods in the first category focus on generating a realistic facial image of a particular person while imposing the conditions of preserving the identity of the individual and having the desired age [2, 16, 17, 18, 19]. Methods belonging to the second category propose the combination of a conditional GAN with an autoencoder, with the purpose of performing age progression as well as age regression (rejuvenating) [3, 15, 20]. Methods in the third category perform facial aging by learning the image style represented by age characteristics and then transferring the age style on the target image [21].

The use of GANs for image synthesis and facial aging has recently enabled to automatically generate highly-realistic images [22], with a simpler procedure with respect to handcrafted approaches [23]. As a drawback, such approaches are often theoretically intractable and in many cases represent black box models that generate data in an unpredictable way. However, in some applications such as medical or forensic analyses, it is not possible to rely on unexplainable decisions. For this reason, Explainable Artificial Intelligence (XAI) techniques are emerging to obtain powerful yet interpretable machine learning models [24, 25].

In the context of XAI, this paper introduces a novel qualitative and quantitative analysis of the similarity of the inner structure of the CNN-based generators in the GANs used for face aging<sup>1</sup>. By performing the proposed analysis on different GANs trained for het-

<sup>1</sup><http://iebil.di.unimi.it/gansXaiAge/index.htm>

erogeneous application scenarios, we observed that GANs for face aging include general purpose image processing operators and partially share their parameters with the other GANs. Fig. 1 describes the application of a state-of-the-art age estimator on both object and face images, showing that image characteristics influencing different apparent ages are general and observable in many objects. We evaluated the proposed method with transfer learning, by training GANs on heterogeneous databases and fine-tuning them for face aging, with results confirming our observations, showing that the aging transformation can be learned using general purpose image databases and a fine-tuning step.

This is the first method in the literature that deals with explaining GAN models, by analyzing the similarity of the filters of the CNNs, which are the main processing module and comprise the large majority of the parameters used in GANs. The paper is structured as follows. Section 2 introduces GANs for face aging, Section 3 describes the methodology, and Section 4 presents the experimental results. Finally, Section 5 concludes the work.

## 2. GANS FOR FACE AGING

Face aging based on GANs can be represented as a particular case of style transfer, where the GAN learns the style that transforms images representing young people ( $X$ ) into images of older individuals ( $Y$ ) [21]. In this work, we consider the cycleGAN (cyGAN) model [26], which has been successfully applied for image style transfer by learning the transformation between images in the origin training set ( $X$ ) and images in the destination training set ( $Y$ ). In particular, a cyGAN transforms an image from domain  $X$  to domain  $Y$  by simultaneously training two GANs:  $GAN_{forward} = (G, D_X)$  and  $GAN_{backward} = (F, D_Y)$ , where  $G$  transforms  $X \rightarrow Y$ ,  $F$  transforms  $Y \rightarrow X$ ,  $D_X$  discriminates between  $\{x\} \subseteq X$  and  $\{F(x)\}$ ,  $D_Y$  discriminates between  $\{y\} \subseteq Y$  and  $\{G(y)\}$ .

CyGANs have the advantages of training both the age progression ( $GAN_{forward}$ ) and age regression ( $GAN_{backward}$ ) GANs at the same time, while not requiring an age estimator during the training process, not requiring multiple images of the same person, and not requiring paired samples (e.g., a different number of images can be present in the origin and destination sets).

## 3. METHODOLOGY


In this section, we describe the proposed novel methodology for explaining the use of GANs for face aging, based on both qualitative and quantitative analyses. The approach consists of three steps: *i*) similarity analysis of the CNNs in GANs. Similarly to analyzing the common genes in two DNA sequences, we analyze the common filters in two CNNs; *ii*) validation of the aging process. We use a CNN-based age estimator to validate the accuracy of the GANs for face aging; *iii*) fine-tuning. We train GANs on heterogeneous application scenarios and fine-tune them to perform face aging.

### 3.1. Similarity Analysis of the CNNs in GANs

The proposed similarity analysis increases the explainability of face aging by performing both a qualitative and quantitative examination of the similarity of the filters in the CNNs of the generators  $G$  and  $F$  of different cyGANs. In particular, our similarity analysis methodology is composed of three steps: *i*) filter extraction, *ii*) filter classification, and *iii*) Cross-GAN Filter Similarity Index (CGFSI).

#### 3.1.1. Filter Extraction

In this step, we extract the two-dimensional filters from the CNN-based generators  $G$  and  $F$ , represented by the weights of all the convolutional layers present in the network architecture. In this work,



**Fig. 2.** Example of the qualitative similarity analysis on one layer of the generator  $G$  of two cyGANs trained on heterogeneous application scenarios: (a) *novel* filters; (b) *common* filters. The majority of the filters belong to the *common* category and consist in general purpose image processing operators.

we focus only on convolutional layers since they comprise the large majority of the parameters of  $G$  and  $F$ .

We analyze the set of filters  $\{H\}$  for each network layer separately, since in many cases the different layers in a CNN process images at different levels of abstraction [27]:

$$H(l) = \{h_1, h_2, \dots, h_{n(l)}\}, \quad 1 \leq l \leq M, \quad (1)$$

where  $H(l)$  is the ensemble of filters for the  $l$ -th layer,  $h_i(x, y)$  is the  $i$ -th filter,  $n(l)$  is the number of filters in the  $l$ -th layer, and  $M$  is the number of layers of the CNN. We extract the filters from both  $G$  and  $F$ , obtaining  $H_G(l)$ ,  $H_F(l)$ .

#### 3.1.2. Filter Classification

To perform a qualitative analysis of the filters in the CNNs of two different GANs, in this step we propose a method for classifying and visually analyzing them based on their level of similarity. In particular, we compare the filters  $\{H\}$  of the generators belonging to two different GANs:  $cyGAN_1$  and  $cyGAN_2$ . The analysis determines at each level  $l$  and for each filter  $h_i \in cyGAN_1$  the corresponding  $h_j \in cyGAN_2$  with the greatest level of similarity  $s$ .

First, we extract the filters from the CNNs of  $cyGAN_1$ ,  $cyGAN_2$ , obtaining  $H_1(l)$ ,  $H_2(l)$ . Second, we analyze the filters separately for each layer  $l$ , to determine the sets of matching filters  $C(l)$ :

$$C(l) = \{(h_i, h_j, s) : h_i \in H_1(l), h_j \in H_2(l), s = c(h_i, h_j), j = \arg \max_j (s)\}, \quad \forall l, \quad (2)$$

where  $c(\cdot)$  represents the similarity function. In this work, we used the two-dimensional cross-correlation function. We perform the same analyses on the filters extracted from the generators  $G$  and  $F$ , obtaining  $C_G(l)$ ,  $C_F(l)$ . Third, we apply a threshold  $th_s$  on the similarity values  $s$  to separate the filters of  $C_G(l)$ ,  $C_F(l)$  in two categories:

$$(h_i, h_j, s) \in \begin{cases} novel & \text{if } s < th_s \\ common & \text{if } s \geq th_s \end{cases}, \quad \forall (h_i, h_j, s) \in C_G, C_F. \quad (3)$$

Fig. 2 describes the qualitative analysis of the filters  $C_G(l)$ , by showing the *novel* filters and the *common* filters obtained by analyzing a specific layer of the generator  $G$  of two cyGANs trained on heterogeneous application scenarios. In the figure, we considered  $th_s = 0.5$ . In this case, the majority (85%) of the filters belong to the *common* category, therefore the CNNs in the two GANs share 85% of filters with similarity  $s \geq 0.5$ . In most cases, the filters consist of general purpose image processing operators. It is also possible to observe a form of filter redundancy since some of the filters in  $G$

are either repeated or are highly similar to other filters in the same cyGAN.

### 3.1.3. Cross-GAN Filter Similarity Index

To perform a quantitative analysis of the similarity of the filters in the CNNs of two different GANs, in this step we propose the Cross-GAN Filter Similarity Index (CGFSI) as the average percentage of the filters in the *common* category for each layer  $l$ :

$$CGFSI = \frac{1}{M} \sum_{l=1}^M \frac{|C_G(l) \cup C_F(l)\{(h_i, h_j, s) \in \text{common}\}|}{|C_G(l)| + |C_F(l)|}, \quad (4)$$

where  $|\cdot|$  indicates the cardinality of the set. We apply the proposed  $CGFSI(cyGAN_1, cyGAN_2)$  to explain how much the filters of the CNNs in  $cyGAN_1$  are similar to the filters in  $cyGAN_2$ .

## 3.2. Validation of the Aging Process

To validate the accuracy of the aging process performed using cyGANs in a quantitative way, contrarily to the method in the literature that only proposes a qualitative evaluation of cyGANs for face aging [21], in this step we use a CNN-based age estimator to estimate the apparent age of the generated face images. In this work, we used the DEX age estimator proposed in [8].

First, we perform the face aging by applying the GAN:  $I_{aged} = cyGAN_{(aging,A)}(I_{orig})$ , where  $cyGAN_{(aging,A)}$  represents the GAN that transforms the image  $I_{orig}$  to the target age  $A$ . Second, we estimate the age of the generated image  $\hat{a}_{aged}$ . Third, we evaluate the accuracy of the aging transformation by computing the mean of the Absolute Age Difference (AAD) between the target age  $A$  and the estimated age  $\hat{a}_{aged}$ .

## 3.3. Fine-Tuning

To perform a qualitative and quantitative evaluation of how the similarity of the filters in the CNNs enables transfer learning in two GANs, in this step we consider a cyGAN trained on a different application scenario and fine-tune it to perform age progression.

First, we consider  $cyGAN_{aging}$ , trained for age progression, and  $cyGAN_{other}$ , trained on a different application scenario (e.g., trained on the general purpose object database CIFAR-10). Second, we perform a fine-tuning of  $cyGAN_{other}$  on the translation domains  $X, Y$  of the face aging problem, obtaining  $cyGAN_{cross}$ . Third, we compare the results of face aging performed using  $cyGAN_{cross}$  with the ones obtained using  $cyGAN_{aging}$ , both in a qualitative way by comparing the apparent age of the images transformed by the two GANs, and in a quantitative way by using the procedure described in Section 3.2.

## 4. EXPERIMENTAL RESULTS

In this section, we introduce the databases used in our study and describe the result of the proposed approach. In particular, we present the obtained CGFSI values, the validation of the aging process, and the results of the face aging obtained by performing a fine-tuning of GANs trained in different application scenarios.

### 4.1. Databases

To evaluate the proposed method, we considered two face databases captured in the wild. First, we used the AgeDB face dataset [28], composed of 16,487 images with different sizes, each associated with the verified real age of the person. Second, we used the CelebA face dataset [10], composed of 202,599 images with different sizes, with no associated age. We used the method described in [29] to detect the faces in the images and cropped all the images to the same average size of  $175 \times 240$  pixels. To compute the CGFSI, we used the general purpose image translation datasets *horse2zebra*

**Table 1.** Cross-GAN Filter Similarity Indexes (CGFSI): quantitative results of the similarity between different GANs. The CGFSI between  $cyGAN_{aging}$  and  $cyGAN_{other}$  are always higher than the CGFSI between  $cyGAN_{blank}$  and the other GANs, indicating that the training process learns similar filters in the CNNs of the GANs, even on heterogeneous application scenarios.

GAN name	$cyGAN_{aging}$		$cyGAN_{other}$		$cyGAN_{blank}$		
	20+20	20+40	40+20	<i>horse2zebra</i>	<i>sat2map</i>	<i>blank</i>	
$cyGAN_{aging}$	20+20	-	0.63	0.62	0.61	0.59	0.47
	20+40	0.64	-	0.63	0.61	0.60	0.47
	40+20	0.60	0.60	-	0.59	0.58	0.47
$cyGAN_{other}$	<i>horse2zebra</i>	0.56	0.55	0.56	-	0.56	0.47
	<i>sat2map</i>	0.53	0.53	0.54	0.54	-	0.47
$cyGAN_{blank}$		0.33	0.32	0.35	0.37	0.38	-

**Table 2.** Validation of the aging process: quantitative results of the accuracy of the cyGANs for face aging. The  $AAD_{cyGAN}$  values describe the difference between the estimated age and the destination aging class. The transformed images feature a high variation in their estimated age.

GAN name	Age class		$AAD_{cyGAN}$ (MeanStd)
	Orig.	Dest.	
$cyGAN_{aging,20+20}$	(20, 40)	(40, 60)	8.03 <sub>6.65</sub>
$cyGAN_{aging,20+40}$	(20, 40)	(60, 80)	18.42 <sub>13.26</sub>
$cyGAN_{aging,40+20}$	(40, 60)	(60, 80)	10.95 <sub>9.21</sub>

Notes. AAD: Absolute Age Difference.


and *sat2map* [30]. To perform the fine-tuning, we considered the CIFAR-10 object database [9], composed of 60,000 images with size  $32 \times 32$  pixels, depicting everyday life objects and animals.

We used the DEX age estimator [8] to estimate the ages of the CelebA database and we divided the dataset in 5 classes based on the age groups [21]: (0, 20), (20, 40), (40, 60), (60, 80), (80, 100). We discarded the class (80, 100) since it contains a limited number of images. Also, we observed that performing face aging on individuals in the class (0, 20) often results in unrealistic images because the proportions of the face attributes are still in their development phase and are not realistically aged by the cyGAN. Therefore, we focused on the three middle classes (20, 40), (40, 60), (60, 80). For each class, we randomly selected 90% images for training and 10% for testing, obtaining in total  $\approx 100,000$  images for training and  $\approx 10,000$  for testing. Similarly, we applied the DEX age estimator on the CIFAR-10 database and performed the same subdivision based on their apparent age, obtaining in total  $\approx 40,000$  images for training and  $\approx 4,000$  for testing.

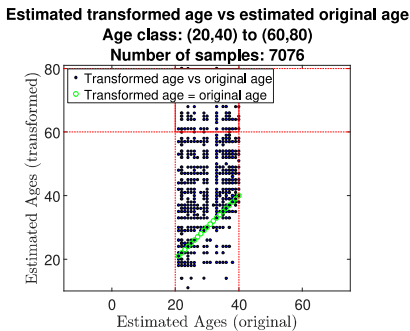
## 4.2. Cross-GAN Filter Similarity Index

We considered the cyGAN implementation with the parameters available at [30] and trained a different  $cyGAN_{aging}$ , for 5 epochs each, for each pair of age classes of the CelebA database. The training process took approximately 24 hours for each cyGAN. Similarly, we trained a different  $cyGAN_{other}$  for each of the general purpose image translation datasets *horse2zebra* and *sat2map*. As a comparison, we considered also  $cyGAN_{blank}$ , which is not trained on any database and in which the filters are initialized with random numbers following a normal distribution.

For each pair of GANs, we performed a quantitative evaluation of their similarity by computing the CGFSI using the threshold value  $th_s = 0.9$ . Table 1 shows the values of the CGFSI for the different pairs of GANs. It is possible to observe that the CGFSI between  $cyGAN_{aging}$  and  $cyGAN_{other}$  is always higher than the CGFSI between  $cyGAN_{blank}$  and the other GANs. This indicates that the training process learns similar filters in the CNNs of GANs and that



**Fig. 3.** Validation of the aging process: quantitative results of the DEX age estimator on the AgeDB dataset. The estimated age  $\hat{a}$  is highly correlated with the real age  $a$ .



**Fig. 4.** Validation of the aging process: quantitative results of the age estimation after applying  $cyGAN_{aging,20+40}$  to transform the images of CelebA from the age class (20, 40) to (60, 80). The aging process increased the apparent age of the majority of the images, however the transformed images feature a high variation in the estimated age.

different GANs could share their parameters even if trained on heterogeneous imaging applications.


#### 4.3. Validation of the Aging Process

To validate the aging process, first we performed a quantitative evaluation of the accuracy of the DEX age estimator using the AgeDB dataset, by comparing the real ages with the ones obtained by DEX. We obtained  $AAD_{DEX} = 6.5$  years, with a standard deviation of 5.7 years, similar to the results published in [8]. The estimated age  $\hat{a}$  is highly correlated with the real age  $a$ , with the correlation coefficient  $r = 0.80$ . Fig. 3 shows the distribution of estimated ages  $\hat{a}$  against the real ages  $a$ .

Second, we performed a quantitative evaluation of the accuracy of the aging process performed using cyGANs, by applying the  $cyGAN_{aging}$  on the testing subset of the CelebA database and computing the  $AAD_{aging}$  between the age estimated by DEX and the target aging class. Table 2 shows the  $AAD_{aging}$  for all the  $cyGAN_{aging}$ . From the table, it is possible to observe that the transformed images feature a high variation in their estimated age  $\hat{a}$ , with  $AAD_{cyGAN} > AAD_{DEX}$ . Fig. 4 shows the distribution of the estimated ages of the original images  $\in (20, 40)$  and the transformed images  $cyGAN_{aging,20+40}(20, 40) \approx (60, 80)$ . From the figure, it is possible to observe that the aging process increased the apparent age differently for every image.

#### 4.4. Fine-Tuning

We trained a different  $cyGAN_{other}$ , for 5 epochs each, for each pair of age classes of the CIFAR-10 database, obtaining  $cyGAN$



**Fig. 5.** Fine-tuning: qualitative example of the results of the aging progress obtained using  $cyGAN_{aging}$  and  $cyGAN_{cross}$ . First row: images obtained using  $cyGAN_{aging}$  trained on CelebA database; second row: images obtained using  $cyGAN_{cross}$  trained on CIFAR-10 database and fine-tuned on CelebA database. The images obtained using  $cyGAN_{cross}$  exhibit the same realistic face aging as the ones obtained with  $cyGAN_{aging}$ .

**Table 3.** Fine-tuning: quantitative results, expressed in terms of AAD, of the accuracy of  $cyGAN_{aging}$ , trained on the CelebA database, compared with the accuracy of  $cyGAN_{cross}$ , trained on the CIFAR-10 database and fine-tuned on the CelebA database.

GAN name	Age class		$AAD_{cyGAN}$ (MeanStd)
	Orig.	Dest.	
$cyGAN_{aging,20+20}$	(20, 40)	(40, 60)	8.03 <sub>6.65</sub>
$cyGAN_{aging,20+40}$	(20, 40)	(60, 80)	18.42 <sub>13.26</sub>
$cyGAN_{aging,40+20}$	(40, 60)	(60, 80)	10.95 <sub>9.21</sub>
$cyGAN_{cross,20+20}$	(20, 40)	(40, 60)	7.34 <sub>6.52</sub>
$cyGAN_{cross,20+40}$	(20, 40)	(60, 80)	17.95 <sub>13.72</sub>
$cyGAN_{cross,40+20}$	(40, 60)	(60, 80)	8.66 <sub>8.68</sub>

Notes. AAD: Absolute Age Difference.

$other,20+20, \dots, cyGAN_{other,40+20}$ . Then, we performed a fine-tuning, for 2 epochs each, of each cyGAN on the corresponding age classes of the CelebA database, obtaining  $cyGAN_{cross,20+20}, \dots, cyGAN_{cross,40+20}$ .

Fig. 5 shows a qualitative evaluation of the fine-tuning process, by comparing the images generated by  $cyGAN_{aging}$  and by  $cyGAN_{cross}$ . Table 3 shows the quantitative evaluation by comparing the results of  $AAD_{aging}$  and  $AAD_{cross}$ , obtained by comparing the difference in the desired and estimated age both on the images transformed using  $cyGAN_{aging}$  and using the corresponding  $cyGAN_{cross}$ . From the figure and the table, it is possible to observe that the  $cyGAN_{cross}$  obtained results comparable to  $cyGAN_{aging}$ , showing that the aging transformation is not exclusively dependent on attributes of the face and can be learned using general purpose object image databases and a fine-tuning step. The results indicate that the training process of GANs produces similar filters even if applied on heterogeneous application scenarios and that is possible to perform transfer learning procedures to obtain similar transformed images even by training on different databases.

## 5. CONCLUSIONS

In this paper, we proposed a novel approach for increasing the explainability of current GAN models. Our method is based on an innovative qualitative and quantitative analysis of the inner structure of GANs for face aging, which shows that GANs partially share their parameters also when trained in heterogeneous application scenarios. To evaluate our approach, we trained GANs on general purpose object databases and performed a fine-tuning on the face database, with results showing that the aging transformation does not depend entirely on face-specific features. Future works should consider different transformations and generative models.

## 6. REFERENCES

- [1] Y. Fu, G. Guo, and T. S. Huang, "Age synthesis and estimation via faces: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 11, pp. 1955–1976, Nov. 2010.
- [2] X. Tang, Z. Wang, W. Luo, and S. Gao, "Face aging with identity-preserved conditional generative adversarial networks," in *Proc. of CVPR*, Jun. 2018, pp. 7939–7947.
- [3] L. Chen, X. Hu, and Z. Zhang, "Face aging with boundary equilibrium conditional autoencoder," *IEEE Access*, vol. 6, pp. 54834–54843, 2018.
- [4] X. Shu, J. Tang, Z. Li, H. Lai, L. Zhang, and S. Yan, "Personalized age progression with bi-level aging dictionary learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 905–917, Apr. 2018.
- [5] I. Kemelmacher-Shlizerman, S. Suwajanakorn, and S. M. Seitz, "Illumination-aware age progression," in *Proc. of CVPR*, Jun. 2014, pp. 3334–3341.
- [6] W. Wang, Y. Yan, Z. Cui, J. Feng, S. Yan, and N. Sebe, "Recurrent face aging with hierarchical autoregressive memory," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018.
- [7] W. Wang, Z. Cui, Y. Yan, J. Feng, S. Yan, X. Shu, and N. Sebe, "Recurrent face aging," in *Proc. of CVPR*, Jun. 2016, pp. 2378–2386.
- [8] R. Rothe, R. Timofte, and L. Van Gool, "Deep expectation of real and apparent age from a single image without facial landmarks," *Int. J. Comput. Vis.*, vol. 126, no. 2, pp. 144–157, Apr. 2018.
- [9] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009, <http://www.cs.toronto.edu/~kriz/cifar.html>.
- [10] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. of ICCV*, Dec. 2015.
- [11] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [12] R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, and F. Scotti, "A novel pore extraction method for heterogeneous fingerprint images using Convolutional Neural Networks," *Pattern Recognit. Lett.*, vol. 113, no. 1, pp. 58–66, Oct. 2018.
- [13] A. Anand, R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, and F. Scotti, "Age estimation based on face images and pre-trained Convolutional Neural Networks," in *Proc. of CISDA*, Nov. 2017, pp. 1–7.
- [14] A. Genovese, V. Piuri, K. N. Plataniotis, and F. Scotti, "Palmnet: Gabor-PCA convolutional networks for touchless palmprint recognition," *IEEE Trans. Inf. Forensics Security*, 2019.
- [15] J. Zeng, X. Ma, and K. Zhou, "CAAE++: Improved CAEE for age progression/regression," *IEEE Access*, vol. 6, pp. 66715–66722, 2018.
- [16] G. Antipov, M. Baccouche, and J. Dugelay, "Face aging with conditional generative adversarial networks," in *Proc. of ICIP*, Sep. 2017, pp. 2089–2093.
- [17] G. Antipov, M. Baccouche, and J. Dugelay, "Boosting cross-age face verification via generative age normalization," in *Proc. of IJCB*, Oct. 2017, pp. 191–199.
- [18] S. Liu, Y. Sun, D. Zhu, R. Bao, W. Wang, X. Shu, and S. Yan, "Face aging with contextual generative adversarial nets," in *Proc. of MM*, 2017, pp. 82–90.
- [19] H. Yang, D. Huang, Y. Wang, and A. K. Jain, "Learning face age progression: A pyramid architecture of GANs," in *Proc. of CVPR*, Jun. 2018, pp. 31–39.
- [20] Z. Zhang, Y. Song, and H. Qi, "Age progression/regression by conditional adversarial autoencoder," in *Proc. of CVPR*, Jul. 2017, pp. 4352–4360.
- [21] S. Palsson, E. Agustsson, R. Timofte, and L. Van Gool, "Generative adversarial style transfer networks for face aging," in *Proc. of CVPRW*, Jun. 2018, pp. 21650–21658.
- [22] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," *CoRR*, vol. abs/1812.04948, 2018.
- [23] Y. Hong, U. Hwang, J. Yoo, and S. Yoon, "How generative adversarial nets and its variants work: An overview of GAN," *ACM Comput. Surv.*, 2017.
- [24] J. Choo and S. Liu, "Visual analytics for explainable deep learning," *IEEE Comput. Graph. Appl.*, vol. 38, no. 4, pp. 84–92, Jul. 2018.
- [25] J. M. Alonso, C. Castiello, and C. Mencar, "A bibliometric analysis of the explainable artificial intelligence research field," in *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Theory and Foundations*, J. Medina et al., Ed., pp. 3–15. Springer, 2018.
- [26] J. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. of ICCV*, Oct. 2017, pp. 2242–2251.
- [27] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [28] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotzia, and S. Zafeiriou, "AgeDB: the first manually collected, in-the-wild age database," in *Proc. of CVPRW*, 2017.
- [29] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Proc. Lett.*, vol. 23, no. 10, pp. 1499–1503, Oct. 2016.
- [30] "Unpaired image-to-image translation using cycle-consistent adversarial networks," 2017, <https://junyanz.github.io/CycleGAN/>.