

# Disentangled Face Identity Representations for Joint 3D Face Recognition and Neutralisation

Anis Kacem  
SnT  
University of Luxembourg  
Luxembourg  
anis.kacem@uni.lu

Kseniya Cherenkova  
SnT, University of Luxembourg  
Artec3D  
Luxembourg  
kseniya.cherenkova@uni.lu

Djamila Aouada  
SnT  
University of Luxembourg  
Luxembourg  
djamila.aouada@uni.lu

**Abstract**—In this paper, we propose a new deep learning based approach for disentangling face identity representations from expressive 3D faces. Given a 3D face, our approach not only extracts a disentangled identity representation, but also generates a realistic 3D face with a neutral expression while predicting its identity. The proposed network consists of three components; (1) a Graph Convolutional Autoencoder (GCA) to encode the 3D faces into latent representations, (2) a Generative Adversarial Network (GAN) that translates the latent representations of expressive faces into those of neutral faces, (3) and an identity recognition sub-network taking advantage of the neutralized latent representations for 3D face recognition. The whole network is trained in an end-to-end manner. Experiments are conducted on three publicly available datasets showing the effectiveness of the proposed approach.

**Keywords**— 3D Face Recognition, 3D Shape Modelling, Generative Adversarial Networks, Graph Convolutions

## I. INTRODUCTION

Automatic recognition of faces is a non-intrusive technology that, if performed accurately, can open the door to many innovative applications and revolutionize the interactions of humans with infrastructures and services. With the recent advances in acquisition systems, impressive face recognition systems have been deployed in real-world applications [1], [10], [14]–[16], [18], [21]. In particular, 3D sensors played an important role in making these advances thanks to the rich geometrical information and their robustness to multiple real-world constraints (e.g., pose variations, illumination changes, etc) [19].

Another important constraint in real-world scenarios consists of recognizing faces under different facial deformations usually termed as facial expressions [2], [10]. Indeed, the captured information when acquiring faces encodes different attributes, e.g., identity, expression, pose, etc [2], [3], [13]. Recognizing faces usually requires decoupling identity attributes from other attributes. Following this direction, multiple works tried to build 3D face parametric models disentangling identity attributes from expression ones. In particular, 3D Morphable Models (3DMM) have been widely used in different forms to achieve this goal [2]–[4]. Such models usually assume linear combination of the different attributes which limits their disentanglement abilities due to the natural non-linearity of face shape variations. Aware of such limitations, recent works proposed non-linear models to model 3D face shapes [5], [8], [13], [20]. [20] proposed a Graph Convolutional Autoencoder (GCA) making use of spectral graph convolutions [6] to encode 3D face shapes into non-linear latent representations. [8] presented an intrinsic Generative Adversarial

Network (GAN) architecture named as MeshGAN operating directly on 3D face meshes using a similar strategy as [20]. Their proposed method allows the generation of new identities and expressions. Similarly, [5] introduced spiral convolutions on meshes and used them within a GAN-based architecture to synthesise face meshes. Taking another direction, the authors in [13] decoupled identity and expression representations by building two separate branches of networks, one for identity and the other for expression. Their proposed networks were based on spectral graph convolutions and were applied on deformation representations computed on the face meshes in advance. In order to control the distribution in latent space, they used a variational strategy when training each branch. More recently, the authors in [23] proposed a variational auto-encoder for disentangling identity information from expressive 3D faces and added an additional regulariser on the identity part to solve the issue of imbalanced information over the expressive input faces and the reconstructed neutral faces. While obtaining impressive performance on two datasets, the reported results are not applicable in real-world settings since the method has not been tested on unseen identities.

The main focus of this paper is to effectively disentangle identity attributes from expression ones and use them within a face recognition scenario. Accordingly, we learn a non-linear mapping translating expressive 3D faces to neutral ones. This is achieved at the level of latent representations obtained using a Graph Convolutional Autoencoder (GCA) [20]. The latent space translation is learned using a Generative Adversarial Network (GAN) applied on pairs of expressive and neutral latent representations [12]. Moreover, additional modules and constraints are proposed to better preserve identity features when neutralizing expressions. The whole pipeline is trained on 3D faces of specific subjects (identities) then tested on unseen ones. Figure 1 shows an overview of the proposed approach. The main contributions of this paper are:

- An end-to-end network allowing joint face expression neutralization and face recognition. Given an expressive 3D face, the proposed network can also synthesise the corresponding realistic neutral 3D face;
- A novel data-driven disentangled identity representation obtained by a realistic translation from expressive faces to neutral ones;
- An experimental validation of the proposed approach which is consistent with real-world settings on three publicly available datasets. The identities used for training are not used for testing.

The rest of the paper is organized as follows: Section II describes the proposed approach. In Section III, we provide the network architecture and details on the training. The experimental evaluation is reported in Section IV. Finally, Section V concludes the paper and draws some perspectives.

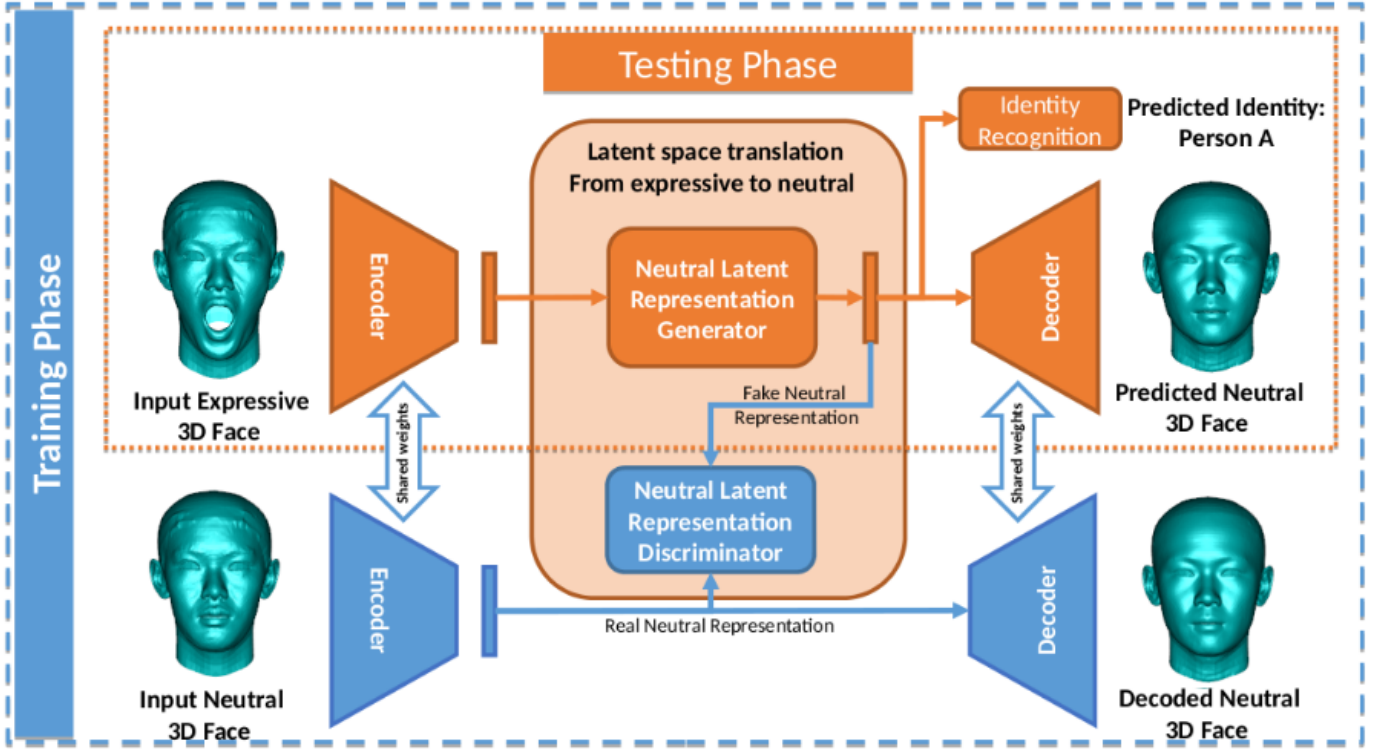


Fig. 1. Approach Overview. In training phase: pairs of expressive and neutral 3D faces of the same identity are collected. These pairs are fed to a two-branch autoencoder sharing the same weights. A GAN-based module consisting of a generator and a discriminator is introduced in the latent space to learn a valid translation from expressive to neutral representation. An identity recognition module is adopted to learn better identity representations. In testing phase: An expressive 3D face is given as input to the network. The encoder encodes the 3D face into a latent representation. Then, the generator of the GAN module translates this latent representation into neutral one which is then decoded into a neutral 3D face. The output features of the face recognition module can be also used to conduct face recognition. The figure is best viewed in color.

## II. PROPOSED APPROACH

Let us consider a set of  $m$  face meshes  $\{F_j\}_{j=1}^m$ . A face mesh is defined as a graph structure,  $F = (V, A)$ . Here,  $V = [v_1, v_2, \dots, v_n]$  denotes the 3D coordinates of  $n$  vertices and  $A$  represents an  $n \times n$  adjacency matrix with  $A_{ij} = 1$  if an edge connection exists between vertex  $v_i$  and  $v_j$ , and  $A_{ij} = 0$  otherwise.

A set of pairs of expressive and neutral 3D faces are collected such that each pair  $(F_i^e, F_i^n)$  involves an expressive 3D face  $F_i^e$  and its neutral counterpart  $F_i^n$  (i.e., a neutral 3D face of the same identity). The main objective of our approach is to translate the expressive face  $F_i^e$  into a neutral face  $\tilde{F}_i^n$  that approximates well the ground truth neutral face  $F_i^n$ , while allowing the prediction of an identity code  $\tilde{y}_i$  that should match the ground truth identity label  $y_i$  of both  $F_i^e$  and  $F_i^n$ .

### A. Graph Convolutional Autoencoder

The first step to our approach is to embed the input 3D face meshes into a low-dimensional latent space that encodes the expression and identity attributes of the input face meshes. To do so, we use a mesh-based autoencoder consisting of an encoder  $Enc$  and a decoder  $Dec$ . The encoder  $Enc$  maps the input face pair  $(F_i^e, F_i^n)$  into latent representations  $(z_i^e, z_i^n)$ . The decoder  $Dec$  directly maps back the neutral latent representation  $z_i^n$  into a neutral face mesh  $\tilde{F}_i^n$  and also maps the neutralised version of the expressive latent representation  $z_i^e$  into the same neutral face mesh  $\tilde{F}_i^n$ . The neutralization of the latent representation will be explained in the next section. From an implementation point of view, the proposed autoencoder can be seen as two branches of Encoder-Decoder sharing the same weights.

Due to the graph structure of the input face meshes, standard convolutions cannot be used in  $Enc$  and  $Dec$ . As a solution, we use the spectral graph convolutions proposed in [6]. In what follows, we recall the background related to these convolutions.

**Spectral Graph Convolutions:** These convolutions are defined on the graph Fourier transform domain. The Fourier transform of the vertices  $V$  is given by  $V_w = U^t V$ , where  $U \in \mathbb{R}^{n \times n}$  whose columns are the orthogonal eigenvectors of the associated Laplacian matrix  $L$ , while the inverse Fourier transform is defined as  $V = UV_w$ . The associated Laplacian to a graph  $F$  is defined as  $L = \Lambda - A$ , where  $\Lambda$  is a diagonal matrix representing the degree of each vertex in  $V$  with diagonal elements  $\Lambda_{ii} = \sum_{j=1}^n A_{ij}$ . The convolution operator in Fourier space can be defined as a Hadamard product of the Fourier transforms of the graph and the filter. For efficiency matters, it is common to use a recursive Chebyshev polynomial of order  $P$  [9], [20] to approximate a kernel  $k_\theta(L) = \sum_{p=0}^{P-1} \theta_p T_p(\hat{L})$ , where  $\hat{L} = 2L/e_{max} - I_n$  is the scaled Laplacian matrix,  $e_{max}$  maximal eigenvalue of  $L$ ,  $\theta \in \mathbb{R}^P$  is a vector of Chebyshev coefficients that should be learned, and  $T_p \in \mathbb{R}^{n \times n}$  is the Chebyshev polynomial of order  $p$  that can be computed recursively [9]. The spectral convolution is finally defined as,

$$c_j = \sum_{i=1}^{n_f} k_{\theta_{i,j}}(L) V_{:,i} \in \mathbb{R}^n, \quad (1)$$

where  $c_j$  is the  $j$ -th feature of the output  $c \in \mathbb{R}^{n \times n_{of}}$  ( $n_{of}$  is the number of output features),  $V_{:,i}$  is the  $i$ -th feature of the input  $V \in \mathbb{R}^{n \times n_f}$  ( $n_f$  is the number of input features), and  $\theta_{i,j} \in \mathbb{R}^P$  are the Chebyshev trainable coefficients. For more details about spectral graph convolutions, readers are referred to [9], [20].

For notation simplicity, we will use in the following  $z_i^e = Enc(F_i^e)$  to denote the expressive latent representation and  $z_i^n = Enc(F_i^n)$  to represent the neutral latent representation.

### B. Latent Space Translation

Given a set of pairs of latent representations  $(z_i^e, z_i^n)$  learned from the expressive and neutral face meshes  $(F_i^e, F_i^n)$ , we aim to find a mapping function  $\mathcal{G} : z_i^e \rightarrow z_i^n$  that translates the expressive latent representation into its corresponding neutral latent representation. To approximate this function we use a neural network  $\mathcal{G}$  on the expressive latent representation  $z_i^e$  and constrain its output to be close to the neutral latent representation  $z_i^n$  by minimizing the L1 distance:

$$\mathcal{L}_{L1} = \|\mathcal{G}(z_i^e) - z_i^n\|_1. \quad (2)$$

For a more realistic translation, we further constrain the distribution of expressive latent representations to match that of neutral latent representations. This is achieved by employing a conditional Generative Adversarial Network (cGAN) [12], [17] on these latent representations. The generator of this cGAN is given by the network  $\mathcal{G}$  that takes the expressive latent representations as input and tries to generate the corresponding neutral latent representations. The discriminator  $\mathcal{D}$  is defined as another network that considers the neutral latent representations  $\mathcal{G}(z_i^e)$  as fake examples, while taking the neutral latent representations  $z_i^n$  as real examples. Furthermore, the discriminator  $\mathcal{D}$  is conditioned using the expressive latent representation  $z_i^e$  as proposed in [12] for image-to-image translation. This condition is established by concatenating the condition  $z_i^e$  to the inputs of the discriminator (*i.e.*,  $z_i^n$  or  $\mathcal{G}(z_i^e)$ ). The adversarial loss to be minimized is then defined as,

$$\mathcal{L}_{GAN} = \mathbb{E}_{z_i^e, z_i^n} [\log \mathcal{D}(z_i^n, z_i^e)] + \mathbb{E}_{z_i^e, z_i^n} [\log(1 - \mathcal{D}(\mathcal{G}(z_i^e), z_i^e))]. \quad (3)$$

The total loss function for latent space translation is given by,

$$\mathcal{L}_{l2l} = \lambda_{L1} \mathcal{L}_{L1} + \lambda_{GAN} \mathcal{L}_{GAN}, \quad (4)$$

where  $\lambda_{L1}$  and  $\lambda_{GAN}$  denote the weights of  $\mathcal{L}_{L1}$  and  $\mathcal{L}_{GAN}$ , respectively.

### C. Face Recognition Module

In order to conduct face recognition, we take advantage of the generated neutral latent representation for an expression-robust face recognition. Specifically, we use a neural network  $\mathcal{R}$  on the generated neutral latent representation  $\mathcal{G}(z_i^e)$  to predict the identity  $\tilde{y}_i = \mathcal{R}(\mathcal{G}(z_i^e))$  of the expressive face  $F_i^e$ . The same network is also used on the real neutral latent representation  $z_i^n$  to obtain a prediction  $\tilde{y}_i = \mathcal{R}(z_i^n)$ . The standard cross-entropy loss is used for training over the identities as follows,

$$\mathcal{L}_{id} = - \sum_{j=1}^s y_i^j \log \tilde{y}_i^j, \quad (5)$$

where  $s$  denotes the number of subjects (*i.e.*, classes) and  $y_i^j$  is the  $j$ -th value of the ground truth label  $y_i$ .

It is important to mention that the face recognition module is not only used for predicting identities but also for a better preservation of the identity information when neutralizing the expression. We experimentally validate this claim in Section IV-A.

## III. FINAL OBJECTIVE AND NETWORK ARCHITECTURE

In addition to the loss functions of latent space translation and face recognition module defined in Eq. 4 and Eq. 5, respectively, we further define two loss functions at the end of the decoder  $Dec$ . The first one enforces the autoencoder to reconstruct an input neutral face  $F_i^n$ , while the second constrains the whole network to reconstruct, starting from an expressive face  $F_i^e$ , an output neutral face that is close to the input neutral face  $F_i^n$ . These reconstruction losses are formulated using an  $L_1$  loss as follows,

$$\mathcal{L}_{rec} = \|Dec(Enc(F_i^n)) - F_i^n\|_1 + \|Dec(\mathcal{G}(Enc(F_i^e))) - F_i^n\|_1, \quad (6)$$

The final objective to be minimised is finally defined as,

$$\mathcal{L}_{tot} = \lambda_{l2l} \mathcal{L}_{l2l} + \lambda_{id} \mathcal{L}_{id} + \lambda_{rec} \mathcal{L}_{rec}, \quad (7)$$

where  $\lambda_{l2l}$ ,  $\lambda_{id}$ , and  $\lambda_{rec}$  represent the weights of the loss functions  $\mathcal{L}_{l2l}$ ,  $\mathcal{L}_{id}$ , and  $\mathcal{L}_{rec}$ , respectively.

**Inference details:** As mentioned in Section II, it is crucial for the training phase of our approach to collect pairs of expressive face meshes and their corresponding neutral faces. These pairs are used to train the latent space translator introduced in Section II-B. During testing phase, the only required input is an expressive face mesh  $F_i^e$  that is encoded using  $Enc$  to obtain an expressive latent representation  $z_i^e$ . The latter is then fed to the generator  $\mathcal{G}$  to generate a neutral latent representation  $z_i^n$ . The generated neutral representation is finally passed to the face recognition module  $\mathcal{R}$  to predict the identity of the input expressive face and also decoded to produce a neutralized face mesh.

**Network architecture:** The encoder  $Enc$  has four Chebyshev convolutional layers. The first three layers have 16 output features per vertex and the last layer has 32 output features. The output of the last convolutional layer is fed into a (Fully Connected) FC layer to obtain a vector latent representation of size 25. The decoder  $Dec$  consists of a FC layer that transforms back the latent representation to a graph representation of  $n$  vertices with 32 features per vertex. The output of this FC is then passed to four Chebyshev convolutional layers with a number of output features 32, 16, 16, and 3, respectively. Each Chebyshev convolutional layer in  $Enc$  and  $Dec$  is approximated using  $P = 6$  polynomials and is followed by a biased ReLU [11]. Regarding the latent representation generator  $\mathcal{G}$ , we use a succession of four FC layers with a number of output features 100, 200, 50, and 25, respectively. Each layer is followed by a ReLU activation. A similar architecture is employed for the discriminator  $\mathcal{D}$  which consists of four FC layers of output size 100, 200, 50, and 1, respectively. The first three FC layers are followed by Leaky ReLU activations and the last layer is followed by a Sigmoid activation. Finally, the face recognition module  $\mathcal{R}$  consists of two FC layers, the first one has an output size 100 and is followed by a ReLU activation, and the second one has an output size equal to the number of training subjects and is followed by a Softmax layer.

In all our experiments we empirically set the weights of the loss function terms as follows:  $\lambda_{l2l} = 0.5$ ,  $\lambda_{L1} = 0.4$ ,  $\lambda_{GAN} = 1$ ,  $\lambda_{id} = 0.05$ , and  $\lambda_{rec} = 2$

## IV. EXPERIMENTS

We experimentally validate the proposed approach on three publicly available datasets. On one hand, we evaluate the face neutralization and compare it with some state-of-art methods. On the other hand, we conduct face recognition based on the disentangled identity representations and compare our results with baselines. The used datasets are the following:

**COMA dataset [20]:** is a dataset designed for dynamic 3D face analysis. It contains more than 20k 3D faces collected from 12 identities. Each subject performed 12 facial expressions showing different challenges.

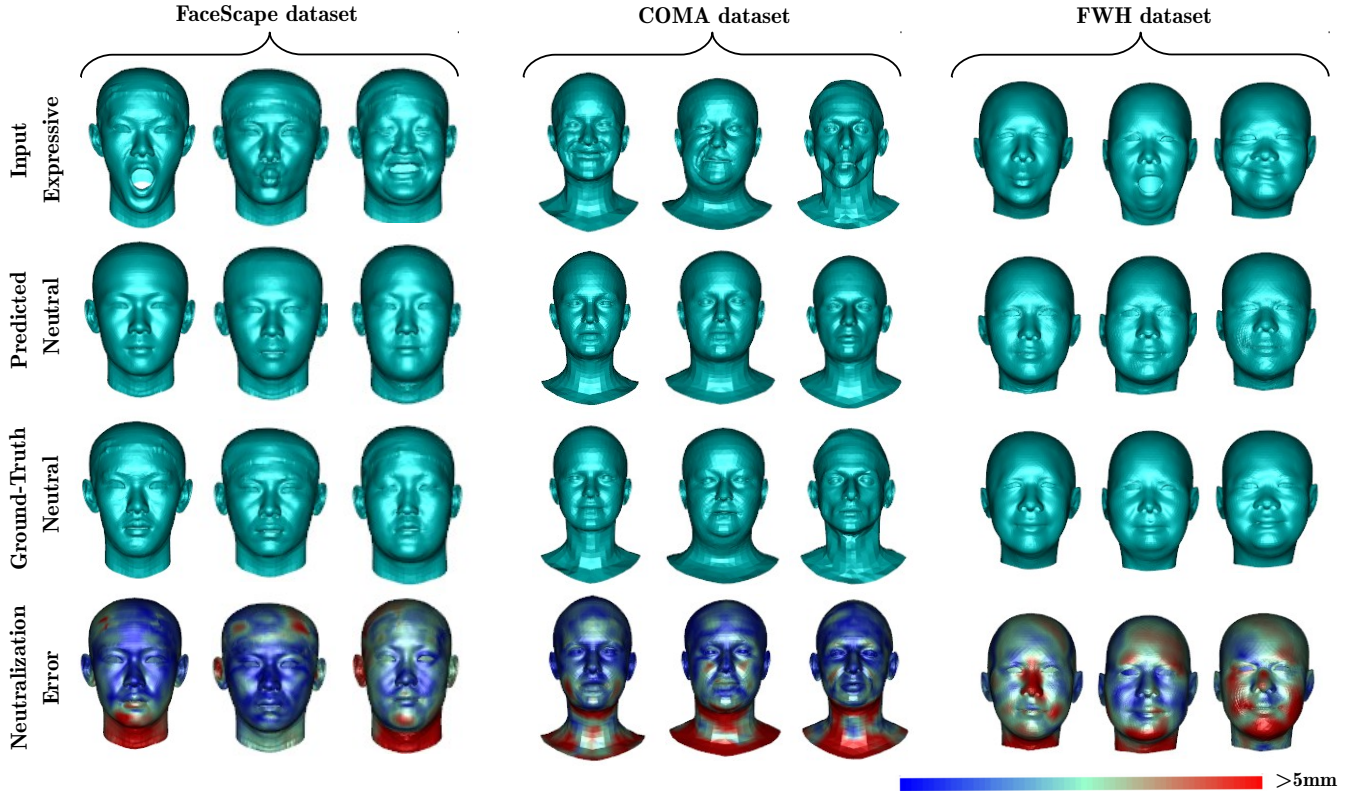


Fig. 2. Face expression neutralization results on FaceScape, COMA, and FWH datasets. The first row shows different input expressive 3D faces, the second row presents the predicted neutralized 3D faces, the third row depicts the corresponding ground-truth neutral 3D faces, and the last row depicts the errors of the neutralization computed using per-vertex Euclidean distances between predicted neutral (second row) and ground-truth neutral (third row) 3D faces. Best viewed in color.

**FaceWareHouse (FWH) dataset [7]:** is a 3D face shape dataset including 47 expressions (including neutral) performed by 150 different identities. The intensity of the expressions involved in this dataset varies from very weak to very strong expressions.

**FaceScape dataset [22]:** consists of a dataset of 938 subjects, each one performed 21 expressions inspired from the FWH dataset (including neutral).

#### A. Face Expression Neutralization

We evaluate the effectiveness of the expression neutralization by comparing the neutralized expressive face to its ground-truth neutral counterpart.

**Experimental protocol:** We split the three datasets into training and testing set following a cross-subject protocol. 70% of the subjects were used for training and the rest for testing. The fitted base models given by the considered datasets have been employed since the proposed approach assumes registered 3D face shapes. The pairs of expressive and neutral 3D faces are directly given by the FWH and FaceScape datasets since both involve static 3D face shapes with different expressions (including neutral). In the dynamic COMA dataset, we take the first frame of each 3D face sequence as neutral face and select 12 expressive faces from other frames.

**Results and discussions:** In Figure 2, we report qualitative results of the face expression neutralization process. In particular, given some input expressive faces selected from the three considered datasets, we present the predicted neutralized faces and compare them to the corresponding ground-truth neutral faces using a per-vertex Euclidean distance. Overall, our approach can perfectly neutralize weak and strong expressions on all the datasets. However, we can observe that the neutralization on FaceScape dataset preserves better the

identity features. For example, the facial wrinkles were not properly reconstructed on some examples of COMA and FWH datasets. This can be explained by the larger number of identities included in the FaceScape dataset compared to FWH and COMA datasets. This suggests that the generalization ability of the proposed network to unseen subjects can be impacted by the number of identities considered in the training.

In Table IV-A, we quantitatively evaluate the expression neutralization results by computing the mean errors on the testing set of all datasets between predicted neutral 3D faces and ground-truth ones. In accordance with the observed qualitative results, the lowest error was recorded on FaceScape dataset reaching a mean error of 2.02 millimeters (mm). The highest error was obtained the FWH dataset corresponding to 3.07 mm. On COMA dataset, we obtained an error of 2.73 mm.

We also compare our results to an adapted version of the Graph Convolutional Autoencoder (GCA) method [20]. In particular, we trained the GCA network by considering the prepared pairs of expressive and neutral faces as input and ground-truth faces, respectively. The mean errors of the adapted GCA were clearly outperformed by our approach on all the datasets. A further comparison has been conducted with the Disentangled Representation Learning (DRL) approach [13]. Specifically, we considered the output of the identity branch of the DRL network as predicted neutral face and evaluated the mean errors following the same protocol as our approach. In Table IV-A, we can note that the DRL approach was more accurate than ours by more than 1 mm on the FWH dataset. It is important to note that DRL method relies on an intermediate representation called *Deformation Representation (DR)* which needs to be computed before training and/or inference. Another important step in DRL approach

is the data augmentation performed on the DR representations to increase the number of identities in the dataset. In particular, 10k faces were added on top of the original FWH dataset used in our approach. Note that we did not compare our results with [23] since their reported results on COMA dataset were obtained by training and testing on the same subjects. The protocol followed by [23] consisted of sorting all videos in alphabetical order, and then taking 10 frames for every 100 frames as the test set and training on the rest. The used protocol to validate our approach does not include the same identities in the training and testing sets.

Finally, we conduct an ablation study of the two loss function terms  $\mathcal{L}_{12l}$  and  $\mathcal{L}_{id}$  defined in Eq. 4 and Eq. 5, respectively. In Table. IV-A, we can observe that dropping one of these terms results in a lower performance showing their importance in the training.

TABLE I  
COMPARISON OF MEAN ERROR OF EXPRESSION NEUTRALIZATION ON DIFFERENT DATASETS. ALL ERRORS ARE IN MILLIMETERS.

Method	COMA	FWH	FaceScape
GCA [20]	3.28	3.56	2.88
DLR [13]	-	1.83	-
Ours w/o $\mathcal{L}_{12l}$	3.45	3.92	3.29
Ours w/o $\mathcal{L}_{id}$	3.06	3.35	2.56
<b>Ours</b>	2.73	3.07	2.02

### B. Face Recognition

A direct application of the disentangled identity representation obtained using our method consists of face recognition. As mentioned in Section II-C, we introduced a face recognition module that improved the face neutralization by enhancing identity preservation as shown in Section IV-A. Such module can also be used to predict the identity of an input face mesh under different facial expressions. To validate this claim, we conducted face identification experiments based on the embedding of the face recognition module  $\mathcal{R}$ . In particular, we consider the identity features obtained by the last FC layer of the face recognition module (before applying softmax) as an identity representation of an input 3D face. These representations are then compared using cosine similarity measure to achieve face identification.

**Experimental protocol:** For face identification, only FWH and FaceScape datasets were retained since COMA dataset consists of only 12 subjects making not useful to evaluate face identification. As done for expression neutralization experiments, the two considered datasets were split into training and testing sets following a cross-subject protocol. The training of the whole network was achieved on the training set and the testing set of unseen subjects was considered to evaluate face identification. We considered neutral 3D faces of the testing set as gallery and the remaining expressions as probe for the two datasets. The evaluation metric was the rank-1 accuracy.

TABLE II  
COMPARISON OF RANK-1 ACCURACIES OF FACE IDENTIFICATION ON FWH AND FACEscape DATASETS.

Method	FWH (%)	FaceScape (%)
GCA [20]	90.48	95.64
Ours w/o $\mathcal{L}_{12l}$	88.75	97.83
<b>Ours</b>	92.34	99.88

**Results and discussions:** In Table II, we report the rank-1 accuracies obtained using our method on FaceScape and FWH datasets. On

the FaceScape dataset which includes a large number of identities, we achieved a rank-1 accuracy of 99.88%. A lower performance consisting of 92.34% was obtained on the FWH dataset which involves lower number of identities than the FaceScape dataset. As noticed for expression neutralization in Section IV-A, face identification results suggest that the performance of the proposed method is better when a large number of identities are involved in the training. Similarly to expression neutralization experiments in Section IV-A, we compare our face identification results with the adapted GCA method [20]. The considered identity features to conduct face identification with GCA were the latent representations of the input expressive faces (i.e., bottleneck of the autoencoder). Our approach outperformed the GCA method by around 4% and 2% on FaceScape and FWH datasets, respectively.

We also evaluate our results when discarding the  $\mathcal{L}_{12l}$  loss term. This means that expression neutralization is only ensured by the reconstruction losses defined in Eq. 6. Unsurprisingly, face identification accuracies were lower than those achieved by the retained approach on the two considered datasets.

### V. CONCLUSION AND FUTURE WORKS

In this paper, we proposed a neural network method for learning disentangled identity representations. Given an input expressive 3D face, the proposed network neutralizes the facial expression and also provides relevant identity features for face identification. This is achieved by translating expressive latent space representations of 3D faces obtained using spectral graph convolutions to neutral ones. The obtained results of face expression neutralization and face identification on three publicly available datasets showed the effectiveness of the proposed approach. However, some limitations were also noticed when the number of identities is not sufficient which resulted in less generalization abilities to unseen subjects. A data augmentation step increasing face shape variations before training can be introduced in the future to overcome this problem. As additional future work, the collection of expressive and neutral pairs can be avoided by considering an unpaired translation of the latent space such as in [24].

### ACKNOWLEDGMENT

This work was funded by the National Research Fund (FNR), Luxembourg, under the project reference CPPP17/IS/11643091/IDform/Aouada

### REFERENCES

- [1] Face ID apple inc. <https://support.apple.com/en-us/HT208108>. Accessed: 2022-03-29.
- [2] Brian Amberg, Reinhard Knothe, and Thomas Vetter. Expression invariant 3d face recognition with a morphable model. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 1–6. IEEE, 2008.
- [3] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999.
- [4] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5543–5552, 2016.
- [5] Giorgos Bouritsas, Sergiy Bokhnyak, Stylianos Ploumpis, Michael Bronstein, and Stefanos Zafeiriou. Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7213–7222, 2019.
- [6] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [7] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2013.

- [8] Shiyang Cheng, Michael Bronstein, Yuxiang Zhou, Irene Kotsia, Maja Pantic, and Stefanos Zafeiriou. Meshgan: Non-linear 3d morphable models of faces. *arXiv preprint arXiv:1903.10384*, 2019.
- [9] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852, 2016.
- [10] Hassen Drira, Boulbaba Ben Amor, Anuj Srivastava, Mohamed Daoudi, and Rim Slama. 3d face recognition under expressions, occlusions, and pose variations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2270–2283, 2013.
- [11] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323, 2011.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [13] Zi-Hang Jiang, Qianyi Wu, Keyu Chen, and Juyong Zhang. Disentangled representation learning for 3d face shape. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11957–11966, 2019.
- [14] Yinjie Lei, Mohammed Bennamoun, Munawar Hayat, and Yulan Guo. An efficient 3d face recognition approach using local geometrical signatures. *Pattern Recognition*, 47(2):509–524, 2014.
- [15] Huibin Li, Di Huang, Jean-Marie Morvan, Yunhong Wang, and Liming Chen. Towards 3d face recognition in the real: a registration-free approach using fine-grained matching of 3d keypoint descriptors. *International Journal of Computer Vision*, 113(2):128–142, 2015.
- [16] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017.
- [17] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [18] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015.
- [19] Dijana Petrovska-Delacrétaz, Gérard Chollet, and Bernadette Dorizzi. *Guide to biometric reference systems and performance evaluation*. Springer, 2009.
- [20] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018.
- [21] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [22] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: A large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [23] Sun, H., Pears, N. & Gu, Y. Information Bottlenecked Variational Autoencoder for Disentangled 3D Facial Expression Modelling. *Proceedings Of The IEEE/CVF Winter Conference On Applications Of Computer Vision*. pp. 157-166 (2022)
- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.