

The Florence 4D Facial Expression Dataset

Filippo Principi, Stefano Berrettini, Claudio Ferrari, Naima Otberdout, Mohamed Daoudi, Alberto del Bimbo

▶ To cite this version:

Filippo Principi, Stefano Berrettini, Claudio Ferrari, Naima Otberdout, Mohamed Daoudi, et al.. The Florence 4D Facial Expression Dataset. IEEE conference series on Automatic Face and Gesture Recognition, Jan 2023, Waikoloa, Hawaii, United States. hal-03835780

HAL Id: hal-03835780 https://hal.science/hal-03835780

Submitted on 1 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Florence 4D Facial Expression Dataset

Filippo Principi¹, Stefano Berretti¹, Claudio Ferrari², Naima Otberdout³,

Mohamed Daoudi⁴, Alberto Del Bimbo¹

¹ Media Integration and Communication Center, University of Florence, Italy

² Department of Engineering and Architecture, University of Parma, Italy

³ Ai movement - University Mohammed IV Polytechnic, Rabat, Morocco

⁴ Univ. Lille, CNRS, Centrale Lille, Institut Mines-Télécom, UMR 9189 CRIStAL, F-59000 Lille, France IMT Nord Europe, Institut Mines-Télécom, Univ. Lille, Centre for Digital Systems, F-59000 Lille, France

Abstract—Human facial expressions change dynamically, so their recognition / analysis should be conducted by accounting for the temporal evolution of face deformations either in 2D or 3D. While abundant 2D video data do exist, this is not the case in 3D, where few 3D dynamic (4D) datasets were released for public use. The negative consequence of this scarcity of data is amplified by current deep learning based-methods for facial expression analysis that require large quantities of variegate samples to be effectively trained. With the aim of smoothing such limitations, in this paper we propose a large dataset, named Florence 4D, composed of dynamic sequences of 3D face models, where a combination of synthetic and real identities exhibit an unprecedented variety of 4D facial expressions, with variations that include the classical neutralapex transition, but generalize to expression-to-expression. All these characteristics are not exposed by any of the existing 4D datasets and they cannot even be obtained by combining more than one dataset. We strongly believe that making such a data corpora publicly available to the community will allow designing and experimenting new applications that were not possible to investigate till now. To show at some extent the difficulty of our data in terms of different identities and varying expressions, we also report a baseline experimentation on the proposed dataset that can be used as baseline.

I. INTRODUCTION

Facial expressions play a primary role in interpersonal relations and are one fundamental way to convey our emotional state [7]. The automatic analysis of facial expressions focused first on images and videos, with rare examples using 3D data [1], [11], [26], [32]. These initial studies focused more on datasets with posed expressions [18], impersonated by actors. The trend is now moving towards spontaneous (not posed) datasets, with some examples of in-the-wild acquisitions [9], [14]. For a summary description of 2D datasets for macro and micro facial expression analysis, we refer to the survey in [13].

Thanks to the rise of powerful deep learning based solutions, the interest in applying expression recognition/generation on 3D and 4D data is growing rapidly. However, with such paradigm, the variety of applications that can be designed and their effectiveness are mainly bounded by the volume and variety of available data to train the models. Thus, it is evident the importance of collecting sufficiently large and variegate datasets, which need to be designed for each specific task. The greater cost and challenges in acquiring dynamic 3D (4D) data is at the base of their limited availability. 3D data are also intrinsically more difficult to process. For example, reducing 3D scans to a same number of vertices connected by a same topology is by itself a task which is complicated to solve in general, while a dense correspondence between scans is, in the majority of cases, a prerequisite to make learning methods to work properly. In addition to being difficult to acquire, 3D data are also difficult to annotate automatically, while manual intervention is impractical and error prone for large volumes of data. A possible workaround is constituted by the creation of synthetic data. For example, this direction has been successfully taken in the case of human body, with a relevant example given by the (not public) dataset used for training the 3D position detector of the skeleton joints in the Kinect for Xbox 360 [28]. Also, if properly designed, the domain gap between real and synthetic 3D data can be made significantly small [29], reducing the performance drop that is usually observed in these cases.

In this paper, we propose the "Florence 4D Facial Expression Dataset" (Florence 4D for short), the first collection of 3D dynamic sequences of facial expressions that includes many identities and long sequences with composed expressions. The peculiar characteristics of our dataset are:

- a large number of identities, both real and synthetic, is included, mostly balanced between male and female;
- all the 3D models share the same number of vertices and connectivity among them, *i.e.*, same topology:
- large number of expressions (70 expressions in total), most of them being variations or combinations of 8 prototypical ones, so to cover a much larger spectrum;
- expression sequences with a large variability: sequences with neutral to apex transitions, as well as with transitions between two apex expressions are included;
- the temporal evolution of the sequences is generated with randomized velocity for improved realism.

All the above characteristics cannot be found in currently available datasets, and can likely open the way to the exploration of completely new tasks. To show that the synthetic data we have generated can effectively complement real data, we compare it with other common benchmarks of real scans in the task of landmark-based 3D face fitting.

The dataset is available at the following link

TABLE I

Most popular HR 3D/4D facial expression datasets. Florence 4D Facial Expressions dataset is unique in providing both real and synthetic 3D models, with an unprecedented number of sequences, also including expression-to-expression transitions

Datasets	#IDs	Pose	Expressions	# sequences
FRGC v2.0 [22]	466	slight	small cooperative expressions-disgust, happiness, sadness,	static
			surprise,-and large uncooperative expression, not categorized	
Bosphorus [27]	105	13 yaw and pitch systematic	neutral plus 6 basis expressions, selected AUs	static
		head rotations)		
BU-3DFE [31]	100	frontal	neutral plus 6 basis expressions	static
Florence 2D/3D [2]	61	2 frontal plus 2 side	neutral	static
BU-4DFE [30]	101	frontal	neutral plus 6 basis expressions	606
CoMA [25]	12	frontal	12 extreme and asymmetric expressions	144
D3DFACS [5]	10	frontal	Action Units	519
VOCASET [6]	12	frontal	Speech	480
Florence 4D (Ours)	95	frontal	70 expressions as variations/combinations of 8 categories:	205,200
			anger, fear, sadness, disgust, surprise, anticipation, trust, joy	

www.micc.unifi.it/resources/datasets/
florence-4d-facial-expression/

II. RELATED WORK

In the following. we summarize the existing 3D dynamic datasets for facial expression analysis. We also refer to the 3D static datasets because they can be used in combination with the dynamic ones in some specific application. We restrict our analysis to high-resolution (HR) 3D datasets, while there are also some low resolution datasets either static or dynamic that have been captured with low-cost, low-resolution cameras like Kinect. Notable examples are the Eurecom dataset [19], the IIIT-D RGB-D face database [12], and FaceWarehouse [4]. Table I compares the main characteristics of the HR datasets discussed below.

a) Static 3D datasets: The Face Recognition Grand Challenge dataset (FRGC v2.0) [21], [22] includes 3D face scans partitioned in three sets, namely, Spring2003 (943 scans of 277 individuals), Fall2003, and Spring2004 (4,007 scans of 466 subjects in total). Individuals have been acquired with frontal view from the shoulder level, with very small pose variations. About 60% of the faces have neutral expression, and the others show expressions of disgust, happiness, sadness, and surprise. The Bosphorus dataset [27] comprises 4,666 high-resolution scans of 105 individuals. There are up to 54 scans per subject, which include prototypical expressions and facial Action Unit (AU) activation. The raw scans of Bosphorus have an average of 30K vertices on the face region. This dataset also contains rotated and occluded scans. The Bimghamton University 3D facial expression dataset [31] comprises 100 subjects (56 females and 44 males) in the six prototypical expressions (angry, disgust, fear, happy, sad, surprise) plus neutral, each expression being reproduced at four intensity levels, from low to exaggerated. Therefore, there are 25 3D expression scans for each subject, resulting in 2,500 3D facial expression scans in total. The subjects vary in gender, ethnicity (White, Black, East-Asian, Middle-east Asian, Indian, and Hispanic Latino) and age (from 18 to 70). The Florence 2D/3D hybrid face dataset [2] includes 3D face scans and video acquisitions of 61 subjects. For each subject, there are two neutral and two side scans plus RGB video recordings



Fig. 1. The overall data pre-processing pipeline used for generating topologically consistent 4D expression morphs. First, the models/scans from the three sources were mapped to the DAZ Studio's Genesis 8 Female (G8-F) topology using the Wrap 3 software (note this is not necessary for the 33 female models that are already in the G8-F topology); then, Daz 3D was used to generate these models as morphing from G8-F, thus making the subsequent animation steps easier (*i.e.*, they can be applied to G8-F), while keeping consistent topology across all the models.

both posed indoor and unconstrained outdoor. Being one of the few datasets including both 2D videos and 3D models of same subjects, it has found large use in evaluating methods that reconstruct the 3D geometry of the face from 2D data.

b) Dynamic 3D (4D) datasets: Given the difficulty of collecting dynamic 3D data, in the literature only a few datasets do exist, which are described in the following. The Bimghamton University 3D plus time (4D) facial expression dataset [30] includes 3D facial expression scans captured at 25 frames per second. For each subject, there are six model sequences showing the six prototypical facial expressions (anger, disgust, happiness, fear, sadness, and surprise), respectively. Each expression sequence contains about 100 frames. The database includes 606 3D facial expression sequences captured from 101 subjects, with a total of approximately 60,600 frame models. The resulting database consists of 58 female and 43 male subjects, with a variety of ethnic/racial ancestries, including Asian, Black, Hispanic/Latino, and White. The CoMA dataset [25] consists of 12 subjects, each one performing 12 extreme and asymmetric expressions. Each expression comes as a sequence of fully-registered meshes with 5,023 vertices. Each sequence is composed of 140 meshes on average, for a total of 20,466 scans. The D3DFACS dataset [5] is a collection of dynamic 3D facial expressions, annotated following the Facial Action Coding System. It contains AU sequences from 10 people, with 519 sequences in total. A version of the dataset with scans registered to a known topology (the same of CoMA) is also available [15]. Finally, we mention VOCASET [6], a 4D face dataset with about 29 minutes of 4D scans captured at 60fps and synchronized audio from 12 speakers (4 males, and 4 females).

Different from all the above, our dataset is a mixture of real and synthetic identities, and the synthesized expressions include the standard 6 prototypical ones as well as mixed expressions, with 4D sequences including also expressionto-expression transitions, the latter feature making it unique.

c) Synthetic datasets: The work of Wood et al. [29] proposed to use a procedurally-generated parametric 3D face model in combination with a comprehensive library of handcrafted assets to render training images with realism and diversity. Machine learning systems for face-related tasks such as landmark localization and face parsing were trained with this data showing that synthetic data can both match real data in accuracy as well as open up new approaches, where manual labeling would be impossible. This work demonstrated it is possible to perform face-related computer vision in the wild using synthetic data alone. In particular, in support of our proposed Florence 4D, authors in [29] showed that it is possible to synthesize data with minimal domain gap, so that models trained on synthetic data generalize to real in-the-wild datasets.

III. PROPOSED DATASET

We identified a key missing aspect in the current literature of 4D face analysis, that is the ability of modeling complex, non-standard expressions and transitions between them. Indeed, current models and datasets are limited to the case, where a facial expression is performed assuming a neutral-apex-neutral transition. This does not hold in the real world, where people continuously switch between one facial expression to another. These observations motivated us to generate the proposed Florence 4D dataset, which is described in the following sections.

A. Source identities

Florence 4D includes real and synthetic identities from different sources: (a) CoMA identities; (b) high-resolution 3D face scans of real identities; (c) synthetic identities.

a) CoMA identities: The CoMA dataset [25] is largely used for the analysis of dynamic facial expressions. An important characteristic of this dataset that contributed to its large use is the fixed topology, according to which all the scans have 5,023 vertices that are connected in a fixed way to form meshes with 9,976 triangular facets. The dataset includes 12 real identities (5 females and 7 males).

b) Synthetic identities: On the Web, a large number of 3D models of synthetic facial characters, either females or males, can be purchased or downloaded for free. Using these

online resources, we were able to add 63 synthetic identities (33 females and 30 males) to the data, selecting those that allow editing and redistribution for non-commercial purposes. Subjects are split in three ethnic groups, Afro (16%), Asian (13%), and Caucasian (71%). Because such identities are synthetic, the resulting meshes are defect free, and perfectly symmetric, which is different from real faces. To make models more realistic, morphing solutions were applied to include face asymmetries.

c) 3D real scans: We acquired 3D scans of 20 subjects (5 females and 15 males) with a 3DmD HR scanner. Subjects are mainly students and university personnel, 30 years old on average. Meshes have approximately 30k vertices. Written consents were collected for these subjects for using their 3D face scans.

B. Data pre-processing

Combining together the identities from the three sources indicated above, we obtained an overall number of 95 identities, 43 females and 52 males. Identities corresponding to synthetic 3D models and 3D scans of real subjects have different topology when compered with CoMA, and a variable number of facets and vertices. Instead, one objective of our dataset was that of providing identities with the same topology as the CoMA dataset (i.e., 5,023 vertices and 9,976 triangular facets). To this end, we used a workflow that involved the joint use of the DAZ Studio [8] and R3DS Wrap 3 [17] software to homogenize the correspondence of the identity meshes. All identities were converted into morphs of the DAZ Studio's Genesis 8 Female (G8-F) base mesh using the Wrap 3 software that allows one mesh to be wrapped over another by selecting corresponding points of the two meshes. The wrapped meshes were then associated with the G8-F mesh as morphs. At the end of the process, we got a G8-F mesh with 95 morphs of different identities. After animating the facial expressions and before exporting the sequence of meshes, we restored the animated G8-F to the original topology of the CoMA dataset. The overall process is illustrated in Figure 1.

C. Facial expressions

With the basic Genesis 8 mesh, we also got a set of facial expressions, in the form of morphs that we used for our dataset. The number of presets was expanded by downloading free and paid packages from the DAZ Studio online shop and from other sites. The base set included 40 different expressions. A paid package of 30 more expressions was added, obtaining a total of 70 different expressions. These expressions were classified according to the Plutchik's wheel of emotions [23], which is illustrated in Figure 2. Following this organization of expressions, we generated a set of secondary expressions from the eight primary ones (for each primary expression, the number of expression per class is indicated): *anger*, AR (6), *fear*, FR (6), *sadness*, SS (13), *disgust*, DT (9), *surprise*, SE (11), *anticipation*, AN (4), *trust*, TT (6), *joy*, JY (15). Details are given in Table II.



Fig. 2. Plutchik's wheel of emotions [23], illustrating expression relations.

гл і	рτ	Е	п	
IAI	ЫĽ	E.	п	

Primary expression	Expressions		
Anger, AR (6)	Angry1, Angry2, Fierce, Glare, Rage, Snarl		
Fear, FR (6)	Afraid, Ashamed, Fear, Scream, Terrified, Wor-		
	ried		
Sadness, SS (13)	Agony, Bereft, Ill, Mourning, Pain, Pouting,		
	Pouty, Sad1, Sad2, Serious, Tired1, Tired2, Upset		
Disgust, DT (9)	Arrogant, Bored, Contempt, Disgust, Displeased,		
	Ignore, Irritated1, Irritated2, Unimpressed		
Surprise, SE (11)	Awe, Confused, Ditzy, Drunk1, Frown, Hurt, In-		
-	credulous, Moody, Shock, Surprised, Suspicious		
Anticipation, AN (4)	Cheeky, Concentrate, Confident, Cool		
Trust, TT (6)	Desire, Drunk2, Flirting, Hot, Kissy, Wink		
Joy, JY (15)	Amused, Dreamy, Excitement, Happy, Inno-		
- · · ·	cent, Laughing, Pleased, Sarcastic, Silly, Smile1,		
	Smile2, Smile3, Smile4, Triumph, Zen		

In the dataset, we named the expressions with pairs of names representing the abbreviation of the primary emotion and the facial expression represented, *e.g.*, JY-smile or SE-incredulous. The Genesis 8 mesh also has 70 morphs of facial expressions available, in addition to 95 identity morphs.

D. Creation of expression sequences

Using the above expression classification, we generated the expression sequences of each identity by iterating through the activation of the expression morphs for each identity morph. The dataset includes two types of sequences for each identity: *single expression* and *multiple expressions*.

a) Single expression: For each identity, the animation of each morph expression is generated as follows:

- Frame 0 neutral expression (morph with weight 0);
- Random frame between 10 and 50¹ expression climax (morph with weight 1);
- Frame 60 neutral expression (morph with weight 0).

The meshes in a sequence are named with the name of the expression and the number of the corresponding frame as a

¹With the randomization of the climax frame, we generated a greater variability in the speed of the transition from the neutral to the climax expression and back to the neutral expression for each identity.



Fig. 3. Sample frames from a generated sequence: (top) the expression passes from neutral to apex and to neutral again; (bottom) the expression passes from neutral to apex for expr. 1, then to apex for expr. 2, and finally to neutral again.

suffix (*e.g.*, *Smile_*01. An example is shown in the top row of Figure 3.

b) Multiple expressions: For each identity, we created mesh sequences of transitions from a neutral expression to a first expression (expr. 1), then from this expression to a second one (expr. 2), then back from the latter to the neutral expression. Also in this case, the climax frames of the two expressions were randomized to obtain greater variability (*i.e.*, the apex frame for each expression can occur at different times of the sequence). Summarizing, these sequences were created following this criterion:

- Frame 0 neutral expression (morph expr. 1 weight 0);
- Random frame between 15 and 40 morph expr. 1 with weight 1, and morph expr. 2 with weight 0;
- Random frame between 50 and frame 75 morph expr. 1 with weight 0, and morph expr. 2 with weight 1;
- Frame 90 neutral expression (morph expr. 2 with weight 0).

Meshes in a sequence are named with the initials of the primary emotions to which the two expressions involved in the animation belong to, followed by the name of the first and second expression plus a numeric suffix for the frame (*e.g.*, *AN-AR_Confident_Glare_01*. An example is shown in the bottom row of Figure 3.

E. Released data

Table III reports a quick summary of the main characteristics of the Florence 4D released data. In particular, we reported the number of identities (male and female), the number of vertices per mesh (same topology for all models), the number of different expressions per identity, the number of sequences that show a neutral-apex expression-neutral transition (6,650 in total); the number of sequences with neutral-expr. 1-expr. 2-neutral transition. Note that, in this latter case, all the possible expression combinations have been generated for a total of 198,550 sequences.

We also note the neutral-expr-neutral sequences include 60 frames each, with the apex intensity for the expression



Fig. 4. Examples frames from generated sequences: (top) apex frames of nine expression sequences for subject DAZ_M_CH20 ; (middle) *angry* expression for a male (DAZ_M_CH020) and a female (DAZ_F_CH073) subject; (bottom) For subject DAZ_F_CH046 the transitions *happy-pain*, and *confident-frown* are shown.

occurring around frame 30; 90 frames are instead generated for the sequences with an expression-to-expression transition, with the expr. 1 apex and the expr. 2 apex occurring around frame 30 and 60, respectively.

TABLE III

FLORENCE 4D EXPRESSION DATASET: SUMMARY OF RELEASED DATA

#IDs (m/f)	#vert	#exprs.	# n-exp-n/# f	# n-exp1-exp2-n/# f
95 (52/43)	5,023	70	70*95 / 60	2090*95 / 90

Some examples of the generated sequences are illustrated in Figure 4. In the top row, the apex frames of nine expression sequences (*i.e.*, smile, wink, disgust, sad, angry, arrogant, fear, happy irritated) of a male synthetic subject are illustrated. The second and third row compare frames of an *angry* expression for a male and a female subject. The two bottom rows, instead, show the transitions *happy-pain*, and *confident-frown* for a given subject.

IV. EXPERIMENTATION

In the following, we report a baseline evaluation for the proposed dataset. We are interested in assessing to what extent our dataset, composed of re-parameterized real scans and totally synthetic sequences, compares to a reference dataset of real scans. We do this by evaluating the task of landmark-based 3D model fitting. As reference datasets to compare with, we chose CoMA and D3DFACS as they share the same mesh topology as Florence 4D, and are composed of 4D expression sequences. They are also common benchmarks employed in other recent studies [3], [25]. For a consistent comparison and fulfill our goal, given the way larger amount and variability of sequences from it, corresponding to the 7

standard expressions, to make it comparable in size and content to CoMA and D3DFACS. Following similar previous works [3], [20], we performed experiments by splitting the data into train and test. To make sure they do not overlap, in one case, we divided the data based on the identities (Identity Split), in the other, based on expressions (Expression Split). In both the cases, we performed a 4-fold cross validation.

A. 3D Expression fitting

Since the main focus of Florence 4D is on expressions, we decided to exclude the problem of identity reconstruction, to avoid ambiguities in the results. The goal is to fit a neutral (not average) 3D face of a subject $\mathbf{S}^n \in \mathbb{R}^{N \times 3}$ to a target expressive face S^e guided by a set of 3D landmarks $Z^e \in \mathbb{R}^{68 \times 3}$. For evaluation, we set up a baseline by first comparing against standard 3DMM-based fitting methods. Similar to previous works [11], [16], we fit S^n to the set of target landmarks Z^e using the 3DMM components. Since the deformation is guided by the landmarks, we first retrieved the landmark coordinates in the neutral face by indexing into the mesh, *i.e.*, $Z^n = \mathbf{S}^n(\mathbf{I}_z)$, where $\mathbf{I}_z \in \mathbb{N}^{68}$ are the indices of the vertices that correspond to the landmarks. We then found the optimal deformation coefficients that minimize the Euclidean error between the target landmarks Z^e and the neutral ones Z^n , and use the coefficients to deform S^n . We experimented the standard PCA-based 3DMM and the DL-3DMM [11]. We also evaluated against recent deep models, including the Neural3DMM [3] and the very recent S2D-Dec [20]. In order to use Neural3DMM as a fitting method, we used the modified architecture as defined in [20], where the model was trained to generate an expressive mesh given its neutral counterpart and the target landmarks Z^e as input. The mean per-vertex Euclidean error between the reconstructed meshes and their ground truth was used as measure, as in the majority of works [3], [10], [24], [25].

Table IV reports the results. It can be noted that for the expression split, results are similar for all the compared datasets. We argue this represents a piece of evidence that the synthetic expressions are as difficult to reconstruct as the real ones, making them valid to be used in practice. Results for the identity split are instead much lower for the proposed Florence 4D. Likely, the variability of synthetic identities is lower than that of real ones, being obtained as a result of a generative software process.

V. DISCUSSION AND CONCLUSIONS

In this paper, a new dataset named Florence 4D, was presented. Its design and generation was guided by the goal of advancing the research in 4D facial analysis, with a particular focus on dynamic expressions. Compared to current datasets, its unique characteristic is that of including sequences of complex, non-standard expressions. Differently from the existing ones, Florence 4D also includes dynamic transitions across expressions, extending the standard neutral-peak-neutral setting. All the sequences were generated with randomized velocity for improved realism. The dataset is a combination of real and synthetic identities,

RECONSTRUCTION ERROR (MM) ON EXPRESSION-INDEPENDENT (LEFT) AND IDENTITY-INDEPENDENT (RIGHT) SPLITS

-		Expression Spl	it		Identity Split	
Method	CoMA	D3DFACS	Florence 4D	CoMA	D3DFACS	Florence 4D
PCA	0.76 ± 0.73	0.42 ± 0.44	0.70 ± 0.81	0.80 ± 0.73	0.56 ± 0.56	0.16 ± 0.17
DL3DMM [11]	$0,86 \pm 0,80$	0.73 ± 1.15	0.83 ± 1.03	0.89 ± 0.79	1.15 ± 1.50	0.17 ± 0.18
Neural3DMM [3]	0.75 ± 0.85	0.59 ± 0.86	1.45 ± 1.43	3.74 ± 2.34	2.09 ± 1.37	1.41 ± 1.09
S2D-Dec	0.52 ± 0.59	0.28 ± 0.31	0.57 ± 1.24	0.55 ± 0.62	0.27 ± 0.30	0.10 ± 0.08

while the expressions are fully synthetic. An experimental validation highlights the little domain gap with respect to real expressive scans, making it a valuable resource for real applications.

VI. ACKNOWLEDGMENTS

This research was partially supported by European Union's Horizon 2020 research and innovation program under grant number 951911 - AI4Media. Most of this work was done when Naima Otberdout was at University of Lille. Her work was supported partially by the French National Agency for Research (ANR) under the Investments for the future program with reference ANR-16-IDEX-0004 ULNE and by the ANR project Human4D ANR-19-CE23-0020.

REFERENCES

- L. Ariano, C. Ferrari, S. Berretti, and A. Del Bimbo. Action unit detection by learning the deformation coefficients of a 3d morphable model. *Sensors*, 21(2):589, 2021.
- [2] A. D. Bagdanov, A. Del Bimbo, and I. Masi. The florence 2d/3d hybrid face dataset. In *Joint ACM Workshop on Human Gesture and Behavior Understanding*, page 79–80, 2011.
 [3] G. Bouritsas, S. Bokhnyak, S. Ploumpis, S. Zafeiriou, and M. Bron-
- [3] G. Bouritsas, S. Bokhnyak, S. Ploumpis, S. Zafeiriou, and M. Bronstein. Neural 3D morphable models: Spiral convolutional networks for 3D shape representation learning and generation. In *IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, pages 7212–7221, 2019.
- [4] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou. Facewarehouse: A 3D facial expression database for visual computing. *IEEE Trans. on Visualization and Computer Graphics*, 20(3):413–425, 2013.
- [5] D. Cosker, E. Krumhuber, and A. Hilton. A facs valid 3D dynamic action unit database with applications to 3D dynamic morphable facial modeling. In *IEEE Int. Conf. on Computer Vision*, pages 2296–2303. IEEE, 2011.
- [6] D. Cudeiro, T. Bolkart, C. Laidlaw, A. Ranjan, and M. Black. Capture, learning, and synthesis of 3D speaking styles. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10101– 10111, 2019.
- [7] V.-T. Dang, H.-Q. Do, V.-V. Vu, and B. Yoon. Facial expression recognition: A survey and its applications. In 23rd Int. Conf. on Advanced Communication Technology (ICACT), pages 359–367, 2021.
 [8] I. Daz Productions. Daz 3D, 2022.
- [9] A. Dhall, R. Goecke, S. Lucey, and T. Gedeon. Static facial expression analysis in tough conditions: Data, evaluation protocol and benchmark. In *IEEE Int. Conf. on Computer Vision Workshops (ICCV Workshops)*, pages 2106–2112, 2011.
- [10] C. Ferrari, S. Berretti, P. Pala, and A. Del Bimbo. A sparse and locally coherent morphable face model for dense semantic correspondence across heterogeneous 3D faces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2021.
- [11] C. Ferrari, G. Lisanti, S. Berretti, and A. Del Bimbo. Dictionary learning based 3d morphable model construction for face recognition with varying expression and pose. In *IEEE Int. Conf. on 3D Vision*, pages 509–517, 2015.
- [12] G. Goswami, S. Bharadwaj, M. Vatsa, and R. Singh. On rgb-d face recognition using kinect. In *IEEE Int. Conf. on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–6, Sep. 2013.
- [13] H. Guerdelli, C. Ferrari, W. Barhoumi, H. Ghazouani, and S. Berretti. Macro- and micro-expressions facial datasets: A survey. *Sensors*, 22(4), 2022.

- [14] J. Kossaifi, G. Tzimiropoulos, S. Todorovic, and M. Pantic. Afew-va database for valence and arousal estimation in-the-wild. *Image and Vision Computing*, 65:23–36, 2017. Multimodal Sentiment Analysis and Mining in the Wild Image and Vision Computing.
- [15] T. Li, T. Bolkart, M. J. Black, H. Li, and J. Romero. Learning a model of facial shape and expression from 4d scans. ACM Trans. Graph., 36(6):194–1, 2017.
- [16] T. Li, T. Bolkart, M. Julian, H. Li, and J. Romero. Learning a model of facial shape and expression from 4D scans. ACM Trans. on Graphics, (Proc. SIGGRAPH Asia), 36(6), 2017.
- [17] R. LLC. R3DS wrap3, 2022.
- [18] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *IEEE Conf.* on Computer Vision and Pattern Recognition (CVPR Workshops), pages 94–101, 2010.
- [19] R. Min, N. Kose, and J.-L. Dugelay. Kinectfacedb: A kinect database for face recognition. *IEEE Trans. on Systems, Man, and Cybernetics: Systems*, 44(11):1534–1548, Nov 2014.
- [20] N. Otberdout, C. Ferrari, M. Daoudi, S. Berretti, and A. D. Bimbo. Sparse to dense dynamic 3d facial expression generation. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 20385–20394. IEEE, 2022.
- [21] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *IEEE Work. on Face Recognition Grand Challenge Experiments*, pages 947–954, June 2005.
- [22] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, and W. Worek. Preliminary face recognition grand challenge results. In *Int. Conf. on Automatic Face and Gesture Recognition*, pages 15–24, 2006.
- [23] R. PLUTCHIK. Chapter 1 a general psychoevolutionary theory of emotion. In R. Plutchik and H. Kellerman, editors, *Theories of Emotion*, pages 3–33. Academic Press, 1980.
- [24] R. A. Potamias, J. Zheng, S. Ploumpis, G. Bouritsas, E. Ververas, and S. Zafeiriou. Learning to generate customized dynamic 3D facial expressions. In *European Conf. on Computer Vision (ECCV)*, pages 278–294, 2020.
- [25] A. Ranjan, T. Bolkart, S. Sanyal, and M. J. Black. Generating 3D faces using convolutional mesh autoencoders. In *European Conf. on Computer Vision (ECCV)*, pages 725–741, 2018.
- [26] G. Sandbach, S. Zafeiriou, M. Pantic, and L. Yin. Static and dynamic 3d facial expression recognition: A comprehensive survey. *Image and Vision Computing*, 30(10):683–697, 2012.
- [27] A. Savran, N. Alyüz, H. Dibeklioğlu, O. Çeliktutan, B. Gökberk, B. Sankur, and L. Akarun. Bosphorus database for 3D face analysis. In *European Workshop on Biometrics and Identity Management*, pages 47–56. Springer, 2008.
- [28] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake. Real-time human pose recognition in parts from single depth images. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1297–1304, 2011.
 [29] E. Wood, T. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and C. Baltruaitis, C
- [29] E. Wood, T. Baltruaitis, C. Hewitt, S. Dziadzio, T. J. Cashman, and J. Shotton. Fake it till you make it: face analysis in the wild using synthetic data alone. In *IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, pages 3661–3671, 2021.
 [30] L. Yin, X. Chen, Y. Sun, T. Worm, and M. Reale. A high-resolution 3d
- [30] L. Yin, X. Chen, Y. Sun, T. Worm, and M. Reale. A high-resolution 3d dynamic facial expression database. In *IEEE Int. Conf. on Automatic Face & Gesture Recognition (FG)*, pages 1–6, 2008.
- [31] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato. A 3D facial expression database for facial behavior research. In *IEEE Conf. on Automatic Face and Gesture Recognition (FG)*, pages 211–216, 2006.
 [32] K. Zhu, Z. Du, W. Li, D. Huang, Y. Wang, and L. Chen. Discriminative
- [32] K. Zhu, Z. Du, W. Li, D. Huang, Y. Wang, and L. Chen. Discriminative attention-based convolutional neural network for 3d facial expression recognition. In *IEEE Conf. on Automatic Face & Gesture Recognition* (FG), pages 1–8, 2019.