A Unified Tensor-based Active Appearance Model

ZHEN-HUA FENG, University of Surrey, UK

JOSEF KITTLER, University of Surrey, UK

BILL CHRISTMAS, University of Surrey, UK

XIAO-JUN WU, Jiangnan University, China

Appearance variations result in many difficulties in face image analysis. To deal with this challenge, we present a Unified Tensor-based Active Appearance Model (UT-AAM) for jointly modelling the geometry and texture information of 2D faces. For each type of face information, namely shape and texture, we construct a unified tensor model capturing all relevant appearance variations. This contrasts with the variation-specific models of the classical tensor AAM. To achieve the unification across pose variations, a strategy for dealing with self-occluded faces is proposed to obtain consistent shape and texture representations of pose-varied faces. In addition, our UT-AAM is capable of constructing the model from an incomplete training dataset, using tensor completion methods. Last, we use an effective cascaded-regression-based method for UT-AAM fitting. With these advancements, the utility of UT-AAM in practice is considerably enhanced. As an example, we demonstrate the improvements in training facial landmark detectors through the use of UT-AAM to synthesise a large number of virtual samples. Experimental results obtained on a number of well-known face datasets demonstrate the merits of the proposed approach.

CCS Concepts: • **Computing methodologies** → **Computer vision tasks**;

Additional Key Words and Phrases: Face image analysis, Active appearance model, Tensor algebra, Missing training samples, Cascaded regression

ACM Reference Format:

Zhen-Hua Feng, Josef Kittler, Bill Christmas, and Xiao-Jun Wu. 0. A Unified Tensor-based Active Appearance Model. *ACM Trans. Multimedia Comput. Commun. Appl.* 0, 0, Article 0 (0), 23 pages. <https://doi.org/0000000.0000000>

1 INTRODUCTION

Geometry and texture information of an object plays important roles in a variety of image- and video-based multimedia applications. For automatic face analysis systems, e.g. face recognition [7, 41, 62], emotion analysis [51] as well as 3D face reconstruction [37, 39], face geometry and texture provide crucial clues for processing and interpreting face instances. In general, the face geometry (also known as face shape) is in the form of a vector consisting of the 2D coordinates of a set of pre-defined key points with semantic meaning, e.g. nose tip, eye corners and face outline; and the texture refers to pixel intensities of the face. In practice, face shape and texture are usually used jointly.

Authors' addresses: Zhen-Hua Feng, University of Surrey, 388 Stag Hill, Guildford, Surrey, GU2 7XH, UK, z.feng@surrey.ac.uk; Josef Kittler, University of Surrey, 388 Stag Hill, Guildford, Surrey, GU2 7XH, UK, j.kittler@surrey.ac.uk; Bill Christmas, University of Surrey, 388 Stag Hill, Guildford, Surrey, GU2 7XH, UK, w.christmas@surrey.ac.uk; Xiao-Jun Wu, Jiangnan University, 1800 Lihu Avenue, Wuxi, Jiangsu, 214122, China, wu_xiaojun@jiangnan.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 0 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1551-6857/0/0-ART0 \$15.00

<https://doi.org/0000000.0000000>

Fig. 1. Example images of a subject in an incomplete training dataset with pose, expression and illumination variations.

Face shape is used for aligning/registering face texture as a prerequisite to extracting meaningful textural features for the following analysis steps in the processing pipeline.

A well-known approach that is capable of recovering and representing the shape and texture information of faces is the Active Appearance Model (AAM) [9], which has also been widely used for modelling other deformable objects such as human hands [53]. However, to build such a face model is non-trivial, due to a wide range of appearance variations in pose, expression, illumination and occlusion [27]. One challenge is how to represent the geometry and texture information of a face compactly. To this end, tensor-based AAM (T-AAM) [41] has been proposed. T-AAM decomposes the original shape and texture space into a set of factor-related subspaces using Higher-Order Single Value Decomposition (HOSVD). Its success derives from the capability of multilinear subspace analysis to decouple multi-factor signals. However, the use of T-AAM has some issues in practice: 1) To build a tensor-based model, we have to collect a large number of training samples with different variations. For example, to build a T-AAM consisting of 10 pose, 10 expression and 10 illumination variations, the required number of training images is 1000 per subject. It is a tedious task to capture and manually label such a large number of face images. As shown in Fig. 1, in practice, we may have missing training samples in a dataset. 2) To fit a new face image, T-AAM generates a number of variation-specific models. Then, a tensor-based state estimation method is used to select a suitable variation-specific model for T-AAM fitting. This task is appearance-based and relies on a face/eye detector; hence the accuracy of state estimation cannot be guaranteed. 3) T-AAM does not consider the problem of the key points located in the self-occluded facial parts that are invisible in a 2D face image, which is usually caused by extreme pose variations. T-AAM has avoided this problem by focusing on relatively slight pose deviations (up to 22.5°) from the frontal one. 4) The classical T-AAM fitting algorithm is gradient-descent-based, and depends on the estimation of the state of each variation mode for a new face. It can easily be trapped in local minima, especially when the state estimation of a face is inaccurate.

To address these issues, we present a Unified Tensor-based AAM (UT-AAM). In contrast to the variation-specific models used in the classical T-AAM, our UT-AAM creates a single shape or texture model based on a multi-linear analysis approach. In such a case, the fitting process of UT-AAM does not require variation or state estimation for the selection of a variation-specific model. Additionally, to further improve the fitting accuracy, UT-AAM exploits the advanced cascaded regression method for model parameter estimation. Last, we can create the proposed UT-AAM model from an incomplete dataset, which is important for practical applications. In practice, it is usually hard to obtain a complete training dataset that consisting of all the images of a subject under all different variation types. In summary, the proposed unified tensor AAM method has four main innovations:

- The unification is achieved by adopting two measures. First, we create a single model across different variation modes, rather than using a number of variation-specific models. Second, we tackle the problem of self-occlusion of faces from large-scale pose variations by proposing a uniform landmarking strategy.
- Enabling the proposed UT-AAM to be constructed from an incomplete dataset with missing training samples. To this end, we investigate tensor completion algorithms to construct our UT-AAM from incomplete training datasets. Furthermore, a new initialisation method is developed to obtain a better reconstruction of missing training samples.
- We develop a learning-based model fitting algorithm. The proposed algorithm does not require prior knowledge of the state of each variation mode for fitting a new face, hence we do not have to perform state estimation before UT-AAM fitting. Moreover, according to our experimental results obtained from the Multi-PIE dataset, the proposed algorithm offers more accurate fitting results.
- We demonstrate the merit of the proposed UT-AAM in synthesising realistic virtual faces with arbitrary pose variations. We perform data augmentation for the training of a facial landmark detector, using a large number of synthetic faces. Experimental results validate the advancements of the use of virtual faces synthesised by our UT-AAM in facial landmark detection.

The rest of this paper is organised as follows: Section 2 gives a brief introduction to related work. Section 3 overviews the classical AAM and T-AAM algorithms. The proposed UT-AAM algorithm is presented in Section 4, and the experimental results are reported in Section 5. Last, the conclusion is drawn in Section 6.

2 RELATED WORK

To obtain the geometry and texture information of 2D faces, a variety of methods have been developed during the past decades, *e.g.* Active Shape Model (ASM) [10], Active Appearance Model (AAM) [9], Constrained Local Model (CLM) [13] and a wide spectrum of facial landmark detection methods [5, 15, 22, 24, 43, 54, 73, 74, 76–78]. Among these algorithms, ASM, CLM and other landmark detection approaches are mainly used to obtain the shape information conveyed by facial landmarks. In contrast, AAM is capable of jointly modelling the shape and texture information of faces.

However, fitting AAM to 2D faces is non-trivial, especially for faces exhibiting a wide range of appearance variations. The developments of AAM addressing this issue can be divided into two categories. In the *first* category, the aim is to improve the structure of the underlying AAM models for better representation of shape or texture information, which is also expected to benefit the subsequent model fitting phase. The *second* category focuses on developing fitting algorithms that generalise well for unseen faces with higher accuracy and less computational cost.

2.1 Model structure

A generative AAM has two PCA-based models for shape and texture, respectively. AAM is capable of generating arbitrary face instances by adjusting the shape and texture models' parameters. The texture model of AAM is usually constructed from raw pixel intensities. Hence an AAM fitting algorithm based on the optimisation of a cost function using texture information is sensitive to appearance variations. To eliminate this dependence, the use of variation-invariant features has been suggested, such as multi-band Value-Hue-Edge [63], image gradients [36], Haar-Like features [44] and image gradient orientations [3, 67]. In addition, local image features have been successfully used to fit a shape model to faces. For example, ASM uses local profiles for a statistical

shape model fitting [10]. Local image patches and edge information are also widely used to optimise the shape fitting process for a given face image [13, 76, 77]. More recently, variation-invariant local descriptors have been very popular in Cascaded Regression (CR-) based facial landmark detection. To extract local features, we usually apply a local descriptor, such as HOG [20, 34, 79], SIFT [74, 78], local pixel difference [5, 54] or Sparse Auto-Encoder [21, 82], in the neighbourhood of facial landmarks.

Another way to modify the model structure is to rely on different underlying methods. The representation capacity of a PCA-based AAM, constructed from a small number of training samples, is limited. For unseen faces, the PCA-based model may miss some details and in consequence it is not able to represent complex faces faithfully. To cope with this issue, more advanced techniques, *e.g.* kernel methods [30, 55, 56] and Deep Boltzmann Machines [19, 49], have been suggested for model construction. Fan *et al.* proposed a new projective invariant method that provides strong shape priors for shape fitting [76, 77]. Note that the aforementioned limitation can also be addressed by using more PCA components trained from more samples with a wide range of appearance variations, as demonstrated by [4, 68]. Besides the representation capacity of AAM, a more important issue is how to construct a compact and structural model. A common way to do this is to use multi-view models, *e.g.* the View-based AAM (V-AAM) [11, 56]. However, this strategy is resource- and time-consuming because we have to construct, store and fit multiple models to a face. As an alternative, Bilinear AAM (B-AAM) constructs a unified model across pose variations [27]. Nevertheless, both V-AAM and B-AAM can only deal with a single variation type among pose, expression and illumination modes.

To analyse multi-factor signals, in recent years, tensor algebra has produced impressive results for many image and video analysis applications [29, 33, 50, 52, 69, 70]. As far as human faces are concerned, an instantaneous observation of the appearance of a subject depends on many factors (as shown in Fig. 1), hence its representation is naturally amended to a tensor-based modelling. Tensors have been shown to be a powerful tool for overcoming difficulties posed by appearance variations in AAM modelling and fitting. For instance, T-AAM builds a set of variation-specific AAMs and then fit the corresponding model to an input image by estimating the states of pose, illumination and expression of the face [41]. It should be noted that the classical AAM, V-AAM and B-AAM can be viewed as degenerations of T-AAM. However, as discussed in the last section, the use of T-AAM has some issues in practice. To tackle these issues, we propose a new framework that introduces a unified tensor-based model across different variation modes.

2.2 Model fitting

The target of AAM fitting is to find the parameters of a face model which best reconstruct the shape and texture information of a new face image. This is a non-linear multi-variable optimisation problem. Depending on the underlying model, the fitting algorithms can either be gradient-descent-based or learning-based. In a gradient-descent-based method, a cost function, designed to recover the model parameters, is optimised by calculating its partial derivatives with respect to the model parameters. For example, in classical AAM, the cost is the pixel intensity difference between the model and an input image. Given their initial values, the model parameters are iteratively updated using gradient descent, such as steepest descent and Gauss-Newton methods [47, 60, 64]. Although some techniques can be used to speed up the fitting phase and prevent the algorithms falling into local minima (*e.g.* a pyramid scheme [9]), the iterative computation of partial derivatives is time-consuming. A distinctive milestone in the history of development is the inverse compositional AAM fitting algorithm that neatly avoids the iterative computation of partial derivatives by inversely compositing model parameters in a projecting-out image space [47, 48]. This seminal work has greatly improved the speed of AAM fitting and consequently broadened its applications.

The extended approaches based on the inverse compositional fitting algorithm have achieved competitive results even for faces in the wild [2, 3, 68].

An alternative way of AAM fitting is to use machine learning techniques, either classification-based or regression-based. For a given observation, the goal of classification-based methods is to maximise the probability of model parameters. For example, Liu considered AAM fitting as a binary classification problem and achieved favourable fitting results with Haar-Like features [44]. As another example, the Support Vector Machine (SVM) has been used as a local expert to identify the best candidate point in the vicinity of a face key point in a CLM-based framework [45, 72]. Unlike classification-based methods, regression-based approaches estimate the partial derivatives by learning from a set of examples. For instance, Cootes *et al.* reported great success with linear regression in their early studies of AAM [8], in which they assumed that there is a constant linear relationship between fitting residuals and parameters updates. However, the simple linear regressor is incapable of solving such a complicated non-linear multi-variable optimisation problem [47]. To deal with this issue, more powerful regression methods have been used, such as Canonical Correlation Analysis (CCA) [18], decision stumps [58] and random forests [12, 59]. More recently, regression-based approaches, in particular the cascaded regression, have been widely used for shape model fitting (also known as facial landmark detection), delivering promising results in both constrained and unconstrained scenarios [15, 24, 43, 78]. The key idea of cascaded regression is to form a strong regressor by cascading a set of weak regressors in series. A weak regressor in cascaded regression could be any regression method, such as linear regression [20, 43, 78, 79], random ferns [5, 54, 74, 77, 83] and even deep neural networks [23, 25, 35, 65, 80, 81, 84]. However, cascaded regression is usually based on a non-parametric Point Distribution Model (PDM) and merely recovers face shapes. In this paper, we extend cascaded regression to our proposed UT-AAM fitting that recovers the shape and texture information jointly. Moreover, we demonstrate the capacity of the proposed model to synthesise realistic 2D face instances for learning-based facial landmark detector training.

3 OVERVIEW OF TENSOR-BASED AAM

In this paper, scalars, vectors, matrices and higher-order tensors are denoted by lower-case letters (a , b , ...), bold lower-case letters ($\mathbf{a}$, $\mathbf{b}$, ...), bold upper-case letters ($\mathbf{A}$, $\mathbf{B}$, ...) and calligraphic upper-case letters ($\mathcal{A}$, $\mathcal{B}$, ...) respectively.

3.1 Active Appearance Model (AAM)

The classical AAM has two PCA-based parametric models, *i.e.* shape and texture models. For a 2D face image, the face shape $\mathbf{s} = [x_1, y_1, \dots, x_L, y_L]^T$ is a vector formed by concatenating the 2D coordinates of L pre-defined key points. Given a set of annotated face images, a PCA-based shape model can be obtained:

$$\mathbf{s} = \bar{\mathbf{s}} + \sum_{k=1}^{N_s} \alpha_k \mathbf{s}_k, \quad (1)$$

where $\bar{\mathbf{s}}$ is the mean shape, $\mathbf{s}_k$ is the k th shape eigenvector obtained by applying PCA to all the training shapes aligned with Procrustes analysis, α_k is the corresponding model parameter and N_s is the number of shape eigenvectors used for the shape model.

To obtain a texture model, pixels inside the face shape of a training image are first warped to a reference shape, *e.g.* the mean shape, using a piece-wise affine transformation [47]. Then raster-scanning is applied to convert the 2D texture of a warped face to a texture vector $\mathbf{t} \in \mathbb{R}^{I_t}$.

Last, AAM applies PCA to construct a texture model:

$$\mathbf{t} = \bar{\mathbf{t}} + \sum_{k=1}^{N_t} \beta_k \mathbf{t}_k, \quad (2)$$

where $\bar{\mathbf{t}}$ is the mean texture, $\mathbf{t}_k$ is the k th eigenvector obtained by PCA, β_k is the corresponding model parameter and N_t is the number of eigenvectors used for the texture model.

The classical PCA-based AAM is capable of representing the majority of shape and texture variations observed in a training dataset. These variations are parametrised by the coefficients of shape and texture models, *i.e.* α and β . Given a new face image $\mathbf{I}$, AAM can reconstruct and model the shape and texture information of the face using a fitting algorithm. The goal of AAM fitting is to adjust model parameters to minimise the pixel intensity difference between a generated face instance and an input image:

$$\|\bar{\mathbf{t}} + \sum_{k=1}^{N_t} \beta_k \mathbf{t}_k - W(\mathbf{I}, \alpha)\|_2^2, \quad (3)$$

where $W(\mathbf{I}, \alpha)$ is a function that warps the face texture inside the shape generated by the shape model with parameter α to the reference shape. This non-linear optimisation problem can be solved using either gradient-descent-based or learning-based approaches [47].

3.2 Tensor-based AAM

3.2.1 Higher-order Singular Value Decomposition. Tensors are higher-order extensions of vectors and matrices. An N th-order tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is an N -dimensional array with multiple indices. Given a face dataset with I_i identity, I_p pose, I_e expression and I_l illumination variations, the shape or texture information of the dataset can naturally be expressed as a tensor. In the shape tensor $\mathcal{S} \in \mathbb{R}^{I_i \times I_p \times I_l \times I_e \times I_s}$, the element $s_{(i_i, i_p, i_l, i_e, i_s)}$ denotes the i_s th entry of the face shape vector for the i_i th identity with the i_p th pose, i_l th illumination and i_e th expression states, where $I_s = 2L$ is the dimensionality of a face shape vector. Similarly, the texture tensor $\mathcal{T} \in \mathbb{R}^{I_i \times I_p \times I_l \times I_e \times I_t}$ reorganises the texture vectors of a training dataset in a tensor fashion.

In contrast to the PCA method used in AAM, T-AAM uses a multilinear subspace analysis to construct tensor-based shape and texture models. To this end, tensor decomposition algorithms are used [6, 31, 38, 66]. The two most important and popular tensor decomposition methods are CANDECOMP/PARAFAC (CP) [6, 31] and Tucker [66] tensors. In one work, T-AAM uses Tucker tensor decomposition to obtain shape and texture models.

Given an N th-order tensor $\mathcal{X}$, Tucker tensor decomposition results in:

$$\mathcal{X} = \mathcal{C} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \dots \times_N \mathbf{U}_N, \quad (4)$$

where $\mathcal{C} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is the core tensor with the same dimensionality of the input tensor $\mathcal{X}$, which models the interaction between the orthonormal mode matrices $\mathbf{U}_n \in \mathbb{R}^{I_n \times I_n} (n = 1, \dots, N)$. The mode- n product ' $\times_n$ ' between a tensor $\mathcal{X}$ and a matrix $\mathbf{Y} \in \mathbb{R}^{J \times I_n}$ results in a new tensor $\mathcal{Z} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N}$, in which each element is calculated by:

$$z_{(i_1, \dots, i_{n-1}, j, i_{n+1}, \dots, i_N)} = \sum_{i_n=1}^{I_n} x_{(i_1, \dots, i_n)} y_{(j, i_n)}. \quad (5)$$

To perform Tucker tensor decomposition, Higher Order Singular Value Decomposition (HOSVD) is usually used, which is also known as the Tucker-1 tensor decomposition method [14, 66]. HOSVD calculates the mode- n matrix $\mathbf{U}_n$ using the left singular matrix of SVD decomposition to the mode- n unfolded matrix $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times I_1 \dots I_{n-1} I_{n+1} \dots I_N}$ of $\mathcal{X}$. To unfold a tensor along the n th mode, we reorder

all the entries in the tensor by stacking the vectors along the n th mode as column vectors in a matrix. Last, the core tensor C is obtained by:

$$C = X \times_1 U_1^T \times_2 U_2^T \dots \times_N U_N^T. \quad (6)$$

3.2.2 Constructing T-AAM. Given a shape tensor S , the use of HOSVD results in:

$$S = C_s \times_1 S_i \times_2 S_p \times_3 S_l \times_4 S_e \times_5 S_s, \quad (7)$$

where $C_s \in \mathbb{R}^{I_l \times I_p \times I_l \times I_e \times I_s}$ is the shape core tensor, $S_i \in \mathbb{R}^{I_l \times I_l}$, $S_p \in \mathbb{R}^{I_p \times I_p}$, $S_l \in \mathbb{R}^{I_l \times I_l}$, $S_e \in \mathbb{R}^{I_e \times I_e}$ and $S_s \in \mathbb{R}^{I_s \times I_s}$ are mode matrices representing the decomposed identity, pose, illumination, expression and shape subspaces. In the same manner, a texture tensor $\mathcal{T}$ is decomposed as:

$$\mathcal{T} = C_t \times_1 T_i \times_2 T_p \times_3 T_l \times_4 T_e \times_5 T_s, \quad (8)$$

in which the elements have similar meanings as those in the shape tensor decomposition.

As a 2D face shape is mainly influenced by pose and expression variations, T-AAM constructs a shape basis sub-tensor:

$$\mathcal{B}_s = C_s \times_2 \mathbf{c}_p^T S_p \times_4 \mathbf{c}_e^T S_e \times_5 S_s, \quad (9)$$

where $\mathbf{c}_p \in \mathbb{R}^{I_p}$ and $\mathbf{c}_e \in \mathbb{R}^{I_e}$ are pose and expression mixture coefficient vectors that indicate a linear combination of different pose or expression variation states, satisfying $\sum_{k=1}^{I_p} c_p(k) = 1$, $\sum_{k=1}^{I_e} c_e(k) = 1$, $0 \leq c_p(k) \leq 1$ and $0 \leq c_e(k) \leq 1$.

Because the pose and expression variations have already been considered in the shape basis sub-tensor, a texture basis sub-tensor is obtained by:

$$\mathcal{B}_t = C_t \times_3 \mathbf{c}_l^T T_l \times_5 T_s, \quad (10)$$

where $\mathbf{c}_l \in \mathbb{R}^{I_l}$ is an illumination mixture coefficient vector that indicates a linear combination of different illumination variation states, satisfying $\sum_{k=1}^{I_l} c_l(k) = 1$ and $0 \leq c_l(k) \leq 1$.

Last, T-AAM constructs a variation-specific shape model:

$$\mathbf{s} = \bar{\mathbf{s}}(\mathbf{c}_p, \mathbf{c}_e) + \sum_{k=1}^{N_s} \alpha_k \mathbf{s}_k(\mathbf{c}_p, \mathbf{c}_e), \quad (11)$$

where $\bar{\mathbf{s}}(\mathbf{c}_p, \mathbf{c}_e)$ is the weighted mean shape computed over all training shapes using the pose and expression mixture coefficient vectors, $\mathbf{s}_k(\mathbf{c}_p, \mathbf{c}_e)$ is the k th column vector of the unfolded matrix $\mathbf{B}_s \in \mathbb{R}^{I_s \times I_l I_l}$ of the shape basis tensor along the 5th mode and α_k is the corresponding model parameter. Similarly, a variation-specific texture model is constructed by T-AAM:

$$\mathbf{t} = \bar{\mathbf{t}}(\mathbf{c}_l) + \sum_{k=1}^{N_t} \beta_k \mathbf{t}_k(\mathbf{c}_l), \quad (12)$$

where $\bar{\mathbf{t}}(\mathbf{c}_l)$ is the weighted mean texture computed over all training texture vectors using the illumination mixture coefficient vector, $\mathbf{t}_k(\mathbf{c}_l)$ is the k th column vector of the unfolded matrix $\mathbf{B}_t \in \mathbb{R}^{I_l \times I_l I_p I_e}$ of the texture basis tensor along the 5th mode and β_k is the corresponding model parameter.

3.2.3 T-AAM fitting. Given a new image, T-AAM first estimates the states of pose, expression and illumination of the face in the image. T-AAM applies face and eye detection algorithms to perform rigid face alignment. Then a tensor-based prediction method is used to estimate the mixture coefficient vectors $\mathbf{c}$ for pose, expression and illumination. According to the constraint used for the value of each element in a mixture coefficient vector, T-AAM can be divided into discrete or continuous T-AAM. The assumption of discrete T-AAM is that the state of a variation mode of a new

face belongs to one of the states of the corresponding variation mode in the training dataset. Hence the value of each element in a mixture coefficient vector can only be either 0 or 1, *i.e.* $c \in \{0, 1\}$. In fact, the state of a specific variation mode of a face could be a linear combination of many discrete variations states in the training dataset. In such a case, continuous T-AAM sets the value of each element in a mixture coefficient vector as a continuous variable, *i.e.* $c \in [0, 1]$.

For discrete T-AAM, we can pre-compute a set of variation-specific models offline and select the best matched one during online fitting. In contrast, continuous T-AAM generates variation-specific shape and texture models online during model fitting. Both of them rely on the estimation step in the fitting phase, *i.e.* obtaining the mixture coefficient vector of each variation mode by estimating the states of pose, expression and illumination modes. Given the estimated mixture coefficient vectors, the corresponding variation-specific shape and texture models are selected to fit the input image using a gradient-descent-based optimisation approach [41].

Compared with AAM, the multilinear subspace analysis used in T-AAM decouples shape and texture information into different variation-related subspaces, and provides a structured and compact representation of the shape and texture information of a dataset. For model fitting, T-AAM creates variation-specific models by estimating the variation states of a new face. This benefits the gradient-descent-based fitting algorithm in two ways. On one hand, the use of variation-specific models initialises model parameters closer to the global minima. On the other hand, by fixing variation types, the search space shrinks to a smaller subspace that is easier to handle. In fact, T-AAM generates a variation-specific model using only the shape and texture eigenvectors related to a specific variation type, hence it fits a new face exhibiting the same variation states much more readily.

4 THE PROPOSED UT-AAM FRAMEWORK

Although the tensor algebra has the capacity to support multiple-factor data analysis, and T-AAM has been reported to be a powerful approach for dealing with appearance variations in face analysis, the use of T-AAM is not without difficulties. The major issue is that the classical T-AAM is a collection of many variation-specific models. For T-AAM fitting, the estimation of variation states of a new face is in practice hard and restricted by the accuracy of face and eye detectors. In addition, T-AAM cannot deal with the problems of self-occlusion and missing training samples. To address these issues, we propose a unified framework for tensor-based AAM. To achieve unification, the proposed UT-AAM framework introduces four techniques. First, UT-AAM creates a unified tensor model across different variation modes. Second, to deal with the problem of self-occlusion of large-scale pose variations, a uniform representation strategy is advocated. Third, the proposed UT-AAM method can be created from an incomplete training dataset with missing training samples by introducing tensor completion approaches. Last, a new cascaded-regression-based model fitting algorithm is presented, which does not require estimating the variation states for fitting a new face.

4.1 Unified shape and texture models

Unlike the variation specific shape and texture models used in the classical T-AAM method, UT-AAM constructs a single tensor-based shape or texture model. Given a set of annotated face images, the corresponding shape tensor $\mathcal{S} \in \mathbb{R}^{I_i \times I_p \times I_l \times I_e \times I_s}$ and texture tensor $\mathcal{T} \in \mathbb{R}^{I_i \times I_p \times I_l \times I_e \times I_t}$, we apply HOSVD to the shape and texture tensors for Tucker tensor decomposition. The resulting unified tensor-based shape model is given as:

$$\mathbf{s} = \bar{\mathbf{s}} + \mathbf{C}_s \times_1 \mathbf{a}_i^T \mathbf{S}_i \times_2 \mathbf{a}_p^T \mathbf{S}_p \times_3 \mathbf{a}_l^T \mathbf{S}_l \times_4 \mathbf{a}_e^T \mathbf{S}_e \times_5 \mathbf{S}_s, \quad (13)$$

where $\bar{\mathbf{s}} \in \mathbb{R}^{2L}$ is the mean shape, $\mathbf{a}_i \in \mathbb{R}^{I_i}$, $\mathbf{a}_p \in \mathbb{R}^{I_p}$, $\mathbf{a}_l \in \mathbb{R}^{I_l}$ and $\mathbf{a}_e \in \mathbb{R}^{I_e}$ are the identity, pose, illumination and expression mixture coefficient vectors. Because the shape of a subject is

Fig. 2. Different definitions for self-occluded facial landmarks. *Left*: using 3D face models; *Middle*: the proposed strategy; *Right*: the corresponding frontal face.

independent of illumination variations, we can compress the shape model as:

$$\mathbf{s} = \bar{\mathbf{s}} + \tilde{\mathbf{C}}_s \times_1 \mathbf{a}_i^T \mathbf{S}_i \times_2 \mathbf{a}_p^T \mathbf{S}_p \times_4 \mathbf{a}_e^T \mathbf{S}_e, \quad (14)$$

where $\tilde{\mathbf{C}}_s = \mathbf{C}_s \times_3 \mathbf{S}_l \times_5 \mathbf{S}_s$. In addition, we have to apply a global affine transform $G(\mathbf{s}, \mathbf{p}_g)$ to the shape, where $\mathbf{p}_g = [s, \theta, t_x, t_y]^T$ is the global affine transform parameter controlling scale, rotation and translation. In this unified tensor-based shape model, a new shape can be represented by a long parameter vector $\mathbf{p} = [\mathbf{p}_g^T, \mathbf{a}_i^T, \mathbf{a}_p^T, \mathbf{a}_e^T]^T$.

Similarly, we can obtain a unified texture model:

$$\mathbf{t} = \bar{\mathbf{t}} + \tilde{\mathbf{C}}_t \times_1 \mathbf{b}_i^T \mathbf{T}_i \times_2 \mathbf{b}_p^T \mathbf{T}_p \times_3 \mathbf{b}_l^T \mathbf{T}_l \times_4 \mathbf{b}_e^T \mathbf{T}_e, \quad (15)$$

where $\mathbf{b}_i \in \mathbb{R}^{I_i}$, $\mathbf{b}_p \in \mathbb{R}^{I_p}$, $\mathbf{b}_l \in \mathbb{R}^{I_l}$ and $\mathbf{b}_e \in \mathbb{R}^{I_e}$ are the identity, pose, illumination and expression mixture coefficient vectors for texture, and $\tilde{\mathbf{C}}_t = \mathbf{C}_t \times_5 \mathbf{T}_t$. The texture of a face can be expressed by the texture model parameter vector $\mathbf{q} = [\mathbf{b}_i^T, \mathbf{b}_p^T, \mathbf{b}_l^T, \mathbf{b}_e^T]^T$. Last, a new face instance can be represented by a unified parameter vector concatenating the shape and texture model parameters $[\mathbf{p}^T, \mathbf{q}^T]^T$. For a new face image, the goal of UT-AAM fitting is to find the model parameter vector best representing the input face.

4.2 Uniform landmarking strategy for self-occluded faces

As discussed at the beginning of this paper, another practical issue for building a unified AAM is the problem of self-occlusion posed by large-scale head rotations. Extreme pose variation of a face often results in some facial parts being invisible. However, classical T-AAM does not consider the problem of self-occlusion for extreme pose variations. This is a crucial issue for constructing unified shape and texture models, as T-AAM would require consistent representation for a face shape or texture across large-scale pose variations, *i.e.* having the same number of face key points and the same dimensionality of extracted texture vectors.

To mitigate this problem, two different strategies have been developed in previous work. The first one is to use view-specific models, in which only visible facial parts are considered. For example, view-based AAM builds 5 different AAMs for the viewpoints at -90° , -45° , 0° , 45° and 90° , using different definitions of face key points [11]. However, this strategy is not feasible for our UT-AAM because it requires consistent shape/texture vectors across poses. Another solution is to use 3D face models. Feng *et al.* used the ground-truth of self-occluded face key points by projecting 3D face vertices to 2D images [20]. The same strategy has also been used in 3D-assisted 2D facial landmark detection [84]. However, it is very hard to manually annotate/guess such kind of face key points, as the red points shown in the left column of Fig. 2. We have to build a 3D face model and fit it to 2D face images to obtain self-occluded face key points. Both the construction and fitting procedures of a 3D face model are hard tasks to implement.

In contrast to these two approaches, we propose a uniform landmarking strategy that is simple to implement and easy for a human to annotate self-occluded facial landmarks. In our UT-AAM, a key point of invisible facial parts is re-defined as the cross point of the face outline and the horizontal line passing the occluded key point, as illustrated at the middle column of Fig. 2. This simple landmarking strategy guarantees the same dimensionality of all the face shapes across pose variations. For face texture, we use the Delaunay triangulation method to create the corresponding face mesh of a given face shape. In such a case, different face shapes with different viewpoints may result in different triangulation results. To address this issue, in the proposed method, the triangulation is created from the reference shape, *i.e.* the average value over all the frontal faces. Then all the face shapes with different viewpoints share the same triangulation result derived from the reference shape. This guarantees the same triangulation result of all the face shapes across different viewpoints. Last, shape-free face patches are obtained by piece-wise affine warp that maps the global texture inside an original face shape to the reference shape [47]. The proposed universal landmarking strategy ensures the same representation of face shapes regardless of pose angles, so that we can obtain a single reference shape and a consistent reference mesh for extracting uniform face texture vectors.

4.3 Coping with missing training samples

In practice, one obstacle impeding the use of a traditional tensor-based model is the problem of missing training samples. Sometimes it is not feasible to collect a complete dataset with all possible combinations of variations, which results in an incomplete training set. For example, some subjects may fail to participate in a particular data collection session. Thus it is clearly important to be able to build a tensor-based model from an incomplete training dataset and to investigate how this model performs in practice.

As a tensor can be unfolded to a matrix, any matrix completion method dealing with missing values can be used for tensor completion [71], in which missing entries are usually randomly distributed in a matrix. However, a missing training sample in our case leads to the whole column vectors of the shape and texture matrices being missing. In addition, matrix completion methods do not consider the interaction and structure of the signals in a dataset. In this section, to address the problem of missing training samples, we use two tensor completion methods, *i.e.* the CP-based CP-WOPT [1] and Tucker-based M²SA algorithms [26].

Given an incomplete tensor $X \in \mathbb{R}^{I_1 \times \dots \times I_N}$ with missing entries, the goal of a tensor completion algorithm is to find a complete tensor X' to minimise $\|X' - X^*\|$, where X^* is the ground truth tensor without missing entries. ‘ $\| \cdot \|$ ’ is the norm of a tensor, which is defined as the square root of the sum of the squares of all the elements in the tensor. However, this complete ground truth tensor is unknown in practice; hence M²SA modifies the cost to find the best low-rank approximation X' for available entries by minimising

$$\|O * (X - X')\|, \quad (16)$$

where ‘ $*$ ’ is the entry-wise product of two tensors, $O \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is an index tensor with the same size as X , in which the value of an entry is set to 1 (or 0) when it is available (or missing). Given an initial estimate of X' , M²SA applies a power method that iteratively updates X' to solve the above optimisation problem. In contrast, CP-WOPT treats the task as a weighted least squares problem and uses a first-order optimisation approach to solve it. For more details of these two tensor completion algorithms, the reader is referred to [26] and [1].

To use the M²SA and CP-WOPT methods, we have to first initialise the missing values for an incomplete tensor. A straightforward initialisation method is to use random values or the average value of some other available entries. However, these approaches do not consider the variations

of a specific missing entry. A better way is to predict the missing entry using only the available entries that share the same variability types. To simplify the discussion, we take the shape tensor as an example to introduce the proposed initialisation algorithm. Note that the algorithm discussed below is also applicable to texture tensors.

Given an incomplete shape tensor $\mathcal{S} \in \mathbb{R}^{I_i \times I_p \times I_l \times I_e \times I_s}$ with missing entries, we first initialise a missing entry using the average value of all available entries with the same pose, illumination and expression variations:

$$\mathbf{s}_{i_i, i_p, i_l, i_e, :}^m \leftarrow \frac{\sum_{(i'_p=i_p) \wedge (i'_l=i_l) \wedge (i'_e=i_e)} \mathbf{s}_{i_i, i'_p, i'_l, i'_e, :}^a}{N_a}, \quad (17)$$

where the superscripts ‘ m ’ and ‘ a ’ stand for missing entry and available entry respectively. The subscripts stand for the position of the corresponding entry in $\mathcal{S}$, and N_a is the number of all the available entries with the same variations as the missing one.

It is obvious that the constraint in equation (17) is very strong. The ‘AND’ operator used in the initialisation method has only one free factor ‘identity’, whereas all the other factors are fixed. Thus the initialised missing entries do not contain other types of variations. However, we may not have enough available entries with the same variation modes as the missing one, especially when the proportion of missing entries is very high. This problem can be solved by replacing the ‘AND’ operator with the ‘OR’ operator:

$$\mathbf{s}_{i_i, i_p, i_l, i_e, :}^m \leftarrow \frac{\sum_{(i'_p=i_p) \vee (i'_l=i_l) \vee (i'_e=i_e)} \mathbf{s}_{i_i, i'_p, i'_l, i'_e, :}^a}{N_a}, \quad (18)$$

in which the ‘OR’ operator makes it easier to find enough available entries sharing the variations of a missing entry.

In summary, for a missing entry, we first use the ‘AND’ operator for initialisation. If no available entries exists, we switch to using the ‘OR’ operator. However, in some extreme cases, even the use of ‘OR’ operator may not find an available entry to initialise the missing one. In such a case, we initialise the missing entry with a random value in $[0, 1]$.

4.4 Cascaded regression for UT-AAM fitting

Given a new face image, the classical T-AAM first applies a tensor-based estimation algorithm to predict the states of pose, expression and illumination variations of the face [41]. Then a gradient-descent-based model fitting algorithm is used to fit the generated variation-specific models to the image. This pipeline has two main drawbacks. The first and most important one is that the estimation step highly relies on the accuracy of a face or eye detector, which cannot be guaranteed, especially for faces exhibiting extreme appearance variations. With an inaccurate estimation result, the generated variation-specific shape and texture models cannot be well fitted to the image. The second drawback is that the fitting algorithm is gradient-descent-based, hence can easily be trapped in local minima.

To deal with the above issues, we propose a new fitting algorithm for our UT-AAM, using a learning-based, coarse-to-fine framework, *i.e.* cascaded regression [15, 78]. Given a face image $\mathbf{I}$ and the initial model parameter vector $\mathbf{p}$, the aim of a regression method is to construct a mapping function:

$$\begin{aligned} \phi : f(\mathbf{I}, \mathbf{p}) &\mapsto \delta \mathbf{p}, \\ \text{s.t. } \|\mathbf{p} + \delta \mathbf{p} - \mathbf{p}^*\|_2^2 &= 0, \end{aligned} \quad (19)$$

Algorithm 1 The proposed T-AAM fitting algorithm

-
- 1: **input** An image $\mathbf{I}$, the trained cascaded regressors $\Phi = \{\phi_1, \dots, \phi_M\}$, initialised model parameters $\mathbf{p}$ and $\mathbf{q}$;
 - 2: **for** $m = 1$ to M **do**
 - 3: Extract local features $f(\mathbf{I}, \mathbf{p})$;
 - 4: Apply the m th weak regressor ϕ_m to obtain $\delta\mathbf{p}$;
 - 5: Update the current model parameter $\mathbf{p} \leftarrow \mathbf{p} + \delta\mathbf{p}$;
 - 6: **end for**
 - 7: Warp the texture in the current shape estimate to the reference shape and estimate $\mathbf{q}$;
 - 8: **return** Predicted model parameters $\mathbf{p}$ and $\mathbf{q}$.
-

where $f(\mathbf{I}, \mathbf{p}) \in \mathbb{R}^{N_f}$ is a feature extraction function that is related to the current model parameter vector, N_f is the dimensionality of an extracted feature vector, $\delta\mathbf{p}$ is the update to the current model parameter vector and $\mathbf{p}^*$ is the ground truth parameter vector of the face. Given a set of training examples, we can learn this mapping function by any regression method, *e.g.* linear regression, random ferns or even deep neural networks. However, a single regressor will not handle the task very well. To address this issue, cascaded regression constructs a strong regressor by cascading M weak regressors, $\Phi = \{\phi_1, \dots, \phi_M\}$.

To construct these cascaded weak regressors, we first train the first weak regressor using the original training samples. Then we apply the first trained weak regressor to update all the initial model parameters, $\mathbf{p} \leftarrow \mathbf{p} + \delta\mathbf{p}$, for the second weak regressor training. The required number of weak regressors can be trained by iteratively repeating this procedure. In this paper, each weak regressor is a linear regressor, *i.e.* $\phi_m : \delta\mathbf{p} = \mathbf{A}_m f(\mathbf{I}, \mathbf{p}) + \mathbf{b}_m$, where $\mathbf{A}_m \in \mathbb{R}^{N_p \times N_f}$ is the projection matrix, $\mathbf{b}_m \in \mathbb{R}^{N_p}$ is the offset and N_p is the dimensionality of a parameter vector $\mathbf{p}$. For the training of the m th weak regressor, the cost function is:

$$\underset{\mathbf{A}_m, \mathbf{b}_m}{\operatorname{argmin}} \sum_{n=1}^N \|\mathbf{A}_m f(\mathbf{I}_n, \mathbf{p}_n) + \mathbf{b}_m - \delta\mathbf{p}_n\|_2^2 + \lambda \|\mathbf{A}_m\|_F^2, \quad (20)$$

where $f(\mathbf{I}_n, \mathbf{p}_n)$ is the extracted feature vector of the n th training example, $\delta\mathbf{p}_n = \mathbf{p}_n^* - \mathbf{p}_n$ is the difference between the current model parameter and the ground truth model parameter, λ is the weight of the regularisation term and $\|\cdot\|_F$ is the Frobenius norm for a matrix. It should be noted that $\mathbf{p}_n$ and $\delta\mathbf{p}_n$ are updated after each weak regressor training.

Given a new face image, the initial parameter estimation $\mathbf{p}$ and a trained cascaded regressor Φ , the model parameter is iteratively updated using the weak regressors in Φ . In this paper, we extract HOG features around each face key point and concatenate them to a long vector as our extracted features, *i.e.* $f(\mathbf{I}, \mathbf{p})$. Once we obtain the final estimate of the shape model parameter $\mathbf{p}$, the shape of a face can be calculated using our tensor-based shape model. Then the global texture of the face image is obtained by wrapping the pixels in the estimated face shape to the reference shape, and the texture model parameter $\mathbf{q}$ can be estimated as introduced in [41, 42]. The proposed fitting algorithm is summarised in Algorithm 1.

5 EXPERIMENTAL RESULTS

In this section, we first compare the proposed UT-AAM with the classical T-AAM on the Multi-PIE face dataset [28]. Then we demonstrate the capacity of the proposed UT-AAM to synthesise a large number of virtual faces and examine how these synthesised faces can improve the training of a facial landmark detector, using the 300-W [57] and AFLW [40] face datasets.

Fig. 3. Variations of the selected subset from Multi-PIE.

5.1 Datasets and experimental settings

The Multi-PIE face dataset has more than 750000 images of 377 subjects, captured from 4 different sessions over the span of five months. The images of a subject in the Multi-PIE dataset were captured across 15 poses, 20 lighting conditions and a range of expression variations. In our experiments, each face in Multi-PIE was manually annotated using 68 facial key points for model training and providing ground truth. The locations of these 68 landmarks were defined as same as the 300-W face dataset. For a self-occluded point, it was annotated using the landmarking strategy in Section 4.2. However, it is laborious to manually annotate all the Multi-PIE images; hence we only annotated a subset containing 60 subjects with 7 poses (12_0, 08_1, 13_0, 05_1, 04_1, 19_1 and 01_0), 3 expressions (neutral from session-1, smile from session-3 and scream from session-4) and all the 20 illumination variations. In total, 25200 images were manually annotated. Fig. 3 shows the variations of the subset.

The 300-W [57] and AFLW [40] datasets have been widely used for benchmarking a facial landmark detection algorithm. For 300-W, the XM2VTS, FRGC, AFW, HELEN and LFPW datasets are semi-automatically annotated to provide training samples, each with 68 facial landmarks. The test set of 300-W contains 600 images, 300 indoor and 300 outdoor. For the AFLW dataset, we follow the AFLW-FULL protocol [83], in which 20000 facial images are used to create the training set and 4386 images are used for testing. In the AFLW-FULL protocol, each AFLW face image has 19 facial landmarks.

The accuracy of different algorithms on Multi-PIE was measured in terms of the widely used point-to-point (pt-pt) error, *i.e.* the average across all the landmarks of the Euclidean distance between the ground truth and fitted face shapes. The Multi-PIE dataset was captured under a controlled scenario and the imaging parameters were fixed, hence the resolution and scale of all the images are unified. However, for the 300-W and AFLW datasets, the images were downloaded from the Internet so there is not consistency among face scales. To address this issue, the normalised pt-pt error is used to measure the accuracy of a facial landmark detector on 300-W and AFLW. Specifically, the inter-ocular distance and face size are used to normalise pt-pt errors for 300-W and AFLW, respectively.

Table 1. A comparison of different algorithms on the Multi-PIE face dataset in terms of fitting error and speed

Algorithm	Initial Error (pixel)	Fitting Error (pixel)	Speed (fps)
AAM	10.4143±0.1399	10.5387±0.2084	3
V-AAM	5.0248±0.0657	3.1367±0.0745	3
T-AAM	4.6182±0.1728	3.3438±0.1378	3
Fast-SIC	10.4143±0.1399	6.7840±0.2027	2
UT-AAM	9.4655±0.0374	2.6964±0.0816	13

5.2 UT-AAM versus T-AAM

This part first compares our UT-AAM with the classical T-AAM, as well as the classical AAM, the view-based AAM (V-AAM) and the Fast Simultaneous Inverse Compositional (Fast-SIC) algorithm [68]. Then we investigate the performance of our UT-AAM trained from an incomplete dataset with missing training samples.

5.2.1 Model fitting. In this experiment, we randomly selected 30 subjects with 12600 ($30 \times 20 \times 3 \times 7$) images from our annotated subset of Multi-PIE as the training set and the remaining 30 subjects with 12600 images were used as the test set. We repeated this procedure 10 times and used the average pt-pt error to assess accuracy. To initialise V-AAM and T-AAM, we assumed that the states of pose, expression and illumination of a test image were already known. A benefit of our UT-AAM in practical applications is that we do not have to estimate the state of each variation type for a test face image. For AAM, V-AAM and T-AAM, the gradient-descent-based Gauss-Newton fitting algorithm was used. Fast-SIC is an advanced AAM fitting algorithm developed for unconstrained scenarios, which is based on the inverse compositional algorithm [68]. For the proposed UT-AAM fitting algorithm, five linear regressors were cascaded.

The initial and fitting errors of different algorithms are shown in Table 1. It should be noted that both V-AAM and T-AAM require knowledge of the states of the pose, expression and illumination variation modes of a test image. Then the corresponding variation-specific mean face is used to initialise the model fitting process. This is the main reason why the initial errors of V-AAM and T-AAM are much lower than those of the classical AAM, Fast-SIC and our UT-AAM. In contrast, UT-AAM does not require any prior information for model fitting. In addition, gradient-descent-based fitting algorithms can be trapped by local minima when the initialisation is very far away from the global optimum, so that the fitting errors of AAM and Fast-SIC are very high. In contrast, both V-AAM and T-AAM obtain much lower fitting errors even using the Gauss-Newton solver. The success of V-AAM and T-AAM has two main origins. The first is the use of variation-specific shape and texture models, which can be fitted to a new face image with the same variation states better than a generic model. The second one, which may be more important, owes to the variation state estimation step that provides better initialisation for model fitting. Note, the proposed UT-AAM with our learning-based model fitting algorithm beats all the other algorithms, including the state-of-the-art Fast-SIC, in terms of accuracy without requiring the estimation of the variation state of each test image. Last, our UT-AAM fitting is much faster than the other algorithms with the speed of 13 fps (frames per second). The speed was obtained on a Intel Xeon E5-2643 v3 CPU.

To further investigate the performance of different algorithms, we present their fitting errors parametrised by pose variations in Fig. 4. As can be seen in the figure, the classical AAM and Fast-SIC are incapable of fitting a new face with combined appearance variations in pose, expression and illumination, especially for faces with extreme pose variations (up to $\pm 90^\circ$ in yaw). For V-AAM

Fig. 4. A comparison of different algorithms on the Multi-PIE face dataset parametrised by pose variations.

and T-AAM, as conjectured above, the use of variation-specific models provides a good mechanism for dealing with appearance variations. However, the main drawback of both methods is their need for a separate step to estimate the state of variations for model initialisation, which is both time-consuming and in practice difficult for faces with a wide range of appearance variations. In contrast, the proposed UT-AAM does not require the variation estimation step and outperforms all the other approaches.

5.2.2 Dealing with missing training samples. In this section, we evaluate the robustness of the proposed UT-AAM to missing training samples. To this end, we first test the reconstruction accuracy of two tensor completion algorithms, *i.e.* M^2SA and CP-WOPT, using both the random initialisation method and the proposed initialisation method. Second, we evaluate the accuracy of the proposed UT-AAM in terms of model fitting error.

To evaluate the capacity of the M^2SA and CP-WOPT algorithms to reconstruct missing training samples of an incomplete tensor, we used the following root mean square (RMS) error as our performance criterion. The RMS reconstruction error for shape or texture was calculated between the ground truth shape/texture and the reconstructed shape/texture of a missing training sample. In this experiment, we randomly selected 30 subjects including 12600 face images to compare the reconstruction accuracy of M^2SA and CP-WOPT. The incomplete shape and global texture tensors were obtained by randomly removing 10%, 20%, ..., 90% and 95% samples from the original complete shape and global texture tensors. Then the proposed UT-AAM was created based on the completed shape and texture tensors using M^2SA and CP-WOPT and tested on the remaining 30 subjects with 12600 images. We repeated this experiment 10 times and reported the average results.

Fig. 5a and Fig. 5b show the reconstruction RMS errors of the M^2SA and CP-WOPT methods using two different initialisation methods for completion of the incomplete shape and texture tensors, parametrised by the proportion of missing training samples. It is evident that the proposed initialisation method performs much better than the random initialisation method when using the M^2SA algorithm, in terms of both the average error and standard deviation of the mean. In contrast, the CP-WOPT algorithm appears to be insensitive to different initialisation methods,

Fig. 5. A comparison of the accuracy of M²SA and CP-WOPT tensor completion algorithms, parametrised by different proportions of missing training samples: (a) the shape reconstruction RMS errors; (b) the texture reconstruction RMS errors. Both the proposed and random initialisation methods are used.

until the proportion of missing samples is larger than 70%. Note that the reconstruction error of the proposed initialisation method increases rapidly when the proportion of missing entries is higher than 80%. The reason is that the restrictive nature of the ‘AND’ operator results in entries being unavailable for initialising missing items and we switch to the ‘OR’ operator. However, the proposed initialisation method still performs better than that of the random initialisation method. The CP-WOPT and M²SA methods have similar performance when the proportion of missing entries is lower than 80%. However, M²SA performs much better than CP-WOPT when more than 80% training samples are missing.

We also evaluate the fitting error of UT-AAM parametrised by different proportions of missing training samples. The results are shown in Fig. 6. Note that, in this experiment, we also compared our UT-AAM with the classical T-AAM using the gradient-descent-based Gauss-Newton fitting algorithm. First, it is evident that the proposed UT-AAM performs well even when a large proportion of training samples are missing. Compared to the model constructed from a complete training dataset (the point ‘0’ on the X-axis), the fitting errors grow slightly as the proportion of missing training samples increases. Second, the M²SA algorithm provides better fitting accuracy than CP-WOPT for the classical T-AAM. In contrast, for the proposed UT-AAM using cascaded-regression-based fitting method, the difference in using M²SA and CP-WOPT is minor. The main reason is that the tensor completion algorithms are only used to build the tensor-based shape and texture model. For cascaded regression based model fitting, the regressor was trained only from available training samples. Last, an important finding here is that the proposed cascaded regression based fitting algorithm is more robust to the variation in the proportion of missing training samples. As the proportion increases, the fitting error of the proposed UT-AAM using cascaded regression grows slowly. In contrast, the classical T-AAM using the gradient-descent-based fitting algorithm is more sensitive to the proportion of missing training samples and has higher fitting error.

Fig. 6. A comparison of the M^2SA and CP-WOPT tensor completion algorithms in terms of the UT-AAM fitting error, parametrised by different proportions of missing training samples. We use our proposed initialisation method for tensor completion, and use the gradient-descent-based (GD) and proposed cascaded-regression-based (CR) algorithms for UT-AAM fitting.

Fig. 7. 2D face instances synthesised using the proposed UT-AAM method: (a) the original 7 poses; (b) synthesised 6 additional poses.

5.3 The use of UT-AAM in facial landmark detection

A potential use of our UT-AAM is through its capacity to generate 2D face instances. In this section, we demonstrate the utility of UT-AAM for the training of 2D facial landmark detectors. Recently, most cutting-edge facial landmark detection algorithms are data-driven and require a large number of training samples. However, the laborious work of manually annotating facial landmarks for face images is tedious. One alternative is to synthesise virtual training samples using a generative model, such as the 3D morphable face model [20, 37, 84]. However, the collection of 3D face scans and the construction of a 3D face model are very involved compared with the data collection and model construction of a 2D face model. The most important advantage of UT-AAM is its capacity to generate realistic 2D face images by changing its mode-related model parameters. Fig. 7 shows some examples synthesised by the proposed UT-AAM, by interpolating the tensor-based pose-mode coefficient vectors between two original faces with different poses. We can see from the synthesised faces with new pose variations that UT-AAM is capable of performing realistic image synthesis.

Fig. 8. A comparison of the cumulative error distribution curves of SDM and CCR, as well as a set of state-of-the-art methods from Baltrusaitis, Hasan, Jaiswal, Miborrow, Yan and Zhou [57], on the 300-W face dataset: (a) results on the 300 outdoor face images; (b) results on the 300 indoor faces.

Despite the capacity of UT-AAM to synthesise realistic face instances, the use of synthesised faces for facial landmark detector training presents some challenges. As discussed in [20], synthesised faces are from different domains than real faces. Synthesised face instances often lack complicated appearance variations in background and occlusion compared with real faces. In that work, to gain a maximum benefit from synthesised faces, we proposed a Cascaded Collaborative Regression (CCR) that was trained on a mixture of real faces and synthesised faces by dynamically reducing the weights of synthesised training samples in the cascade [20]. Motivated by this, we compared the classical supervised descent method (SDM) [78] with CCR for facial landmark detection. For SDM training, we first used the training samples provided by the 300-W dataset, marked by ‘SDM (Real)’. Then we used both the training images provided by 300-W and additional face instances synthesised by UT-AAM for SDM training, marked by ‘SDM (Real+Syn.)’. Last, we trained the CCR model using both the real faces provided by 300-W and our synthesised faces (‘CCR (Real+Syn.)’). In total, 46800 synthesised face instances were used, comprising 60 subjects with 3 expression, 20 illumination and 7 + 6 pose variations (Fig. 7).

The results obtained on the 300-W dataset are shown in Fig. 8. It should be noted that, SDM is one of the most popular algorithms for detecting facial landmarks in unconstrained face images, and is usually used as a baseline. CCR is an improved version of SDM, developed in particular for the purpose of using a mixture of real and synthesised faces. Both CCR and SDM are based on a set of linear regressors in cascade. Fig. 8 shows that SDM performs well on the 300-W benchmark and beats most of the other algorithms. In addition, the use of synthesised 2D face instances improves the performance of SDM significantly. Last, the joint use of CCR and synthesised faces further improves the performance than SDM and beats all the other methods.

More recently, powerful algorithms such as Deep Neural Networks (DNN) have been successfully used for facial landmark detection, delivering impressive results [23, 25, 65]. One underlying assumption for successfully training a DNN is a large number of training samples. However, the size of existing facial landmark detection datasets cannot meet this requirement, limiting the performance of a trained DNN model. To explore the merit of incorporating DNN and synthesised 2D faces in facial landmark detection, we evaluate the performance of three deep neural network architectures, including CNN-6 [23], VGG-16 [61] and ResNet-50 [32], equipped with Wing loss [23]

Table 2. A comparison of different algorithms on the AFLW dataset, in terms of the normalised pt-pt error by face size.

Method	CCL [83]	DAC-CSR [24]	TR-DRN [46]
Error	2.72%	2.27%	2.17 %
Method	CPM+SBR [17]	SAN [16]	GoDP [75]
Error	2.14%	1.91%	1.84%
Method	CNN-6 (Real)	VGG-16 (Real)	ResNet-50 (Real)
Error	1.83%	1.62%	1.51%
Method	CNN-6 (Real+Syn.)	VGG-16 (Real+Syn.)	ResNet-50 (Real+Syn.)
Error	1.77%	1.59%	1.47%

on the AFLW dataset. It should be noted that both VGG-16 and ResNet-50 were trained from scratch, which is different from the use of pre-trained models in [23]. Additionally, a similar annealing strategy as used in CCR is adopted to address the issue that the synthesised 2D faces are from a different domain of real ones.

The final results obtained on AFLW are shown in Table 2. Due to the use of the advanced Wing loss framework, even the simple CNN-6 architecture trained on real AFLW images outperforms the other state-of-the-art facial landmark detectors. In addition, VGG-16 and ResNet-50 perform much better than the simple CNN-6 model in accuracy. Last, the use of synthesised 2D facial images further improves the performance of the three used network architectures in accuracy. The results obtained on AFLW further validate the merits of the use of synthesised 2D faces for the training of a learning-based face landmarking system.

6 CONCLUSION

In this paper, we proposed a unified tensor-based AAM. Compared with the classical tensor-based AAM, the proposed UT-AAM can be created from an incomplete training dataset and results in a unified single tensor model across different variation modes. To deal with the problem of self-occlusion, a unified landmarking strategy was advocated for obtaining universal shape and texture representations of faces across large pose variations. A more efficient and accurate cascaded-regression-based model fitting algorithm was also proposed for UT-AAM fitting. Experiments conducted on the Multi-PIE face dataset demonstrate the merits of the proposed UT-AAM algorithm. Last, we showed that the use of our UT-AAM to augment the volume of training data for learning-based facial landmark detectors improved their performance on the 300-W and AFLW datasets.

ACKNOWLEDGMENT

This work was supported in part by the EPSRC Programme Grant (FACER2VM) EP/N007743/1, EP-SRC/dstl/MURI project EP/R018456/1, the National Natural Science Foundation of China (61373055, 61672265, 61876072, 61602390) and the NVIDIA GPU Grant Program.

REFERENCES

- [1] Evrim Acar, Daniel M Dunlavy, Tamara G Kolda, and Morten Mørup. 2011. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems* 106, 1 (2011), 41–56.
- [2] Joan Alabort-i Medina and Stefanos Zafeiriou. 2017. A Unified Framework for Compositional Fitting of Active Appearance Models. *International Journal of Computer Vision* 121, 1 (2017), 26–64.
- [3] Epameinondas Antonakos, Joan Alabort-i Medina, Georgios Tzimiropoulos, and Stefanos P Zafeiriou. 2015. Feature-based lucas-kanade and active appearance models. *IEEE Trans. on Image Processing* 24, 9 (2015), 2617–2632.

- [4] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. 2016. A 3D morphable model learnt from 10,000 faces. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 5543–5552.
- [5] Xudong Cao, Yichen Wei, Fang Wen, and Jian Sun. 2014. Face alignment by explicit shape regression. *International Journal of Computer Vision* 107, 2 (2014), 177–190.
- [6] J.D. Carroll and J.J. Chang. 1970. Analysis of individual differences in multidimensional scaling via an N-way generalization of Eckart-Young decomposition. *Psychometrika* 35, 3 (1970), 283–319.
- [7] Sang-Il Choi, Sung-Sin Lee, Sang Tae Choia, and Won-Yong Shin. 2018. Face Recognition Using Composite Features Based on Discriminant Analysis. *IEEE Access* 6 (2018), 13663–13670.
- [8] T. Cootes, G. Edwards, and C. Taylor. 1998. Active appearance models. In *European Conference on Computer Vision*. Lecture Notes in Computer Science, Vol. 1407. Springer Berlin / Heidelberg, 484–498.
- [9] T. Cootes, G. Edwards, and C. Taylor. 2001. Active appearance models. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 23, 6 (2001), 681–685.
- [10] T.F. Cootes, C.J. Taylor, D.H. Cooper, J. Graham, et al. 1995. Active shape models-their training and application. *Computer Vision and Image Understanding* 61, 1 (1995), 38–59.
- [11] T.F. Cootes, K. Walker, and C.J. Taylor. 2000. View-based active appearance models. In *IEEE International Conference on Automatic Face and Gesture Recognition*. 227–232.
- [12] Tim F Cootes, Mircea C Ionita, Claudia Lindner, and Patrick Sauer. 2012. Robust and accurate shape model fitting using random forest regression voting. In *European Conference on Computer Vision*. Springer, 278–291.
- [13] David Cristinacce and Tim Cootes. 2006. Feature detection and tracking with constrained local models. In *British Machine Vision Conference*, Vol. 3. 929–938.
- [14] L. De Lathauwer, B. De Moor, and J. Vandewalle. 2000. A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.* 21, 4 (2000), 1253–1278.
- [15] Piotr Dollár, Peter Welinder, and Pietro Perona. 2010. Cascaded pose regression. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1078–1085.
- [16] Xuanyi Dong, Yan Yan, Wang Ouyang, and Yi Yang. 2018. Style Aggregated Network for Facial Landmark Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 379–388.
- [17] Xuanyi Dong, Shou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, and Yaser Sheikh. 2018. Supervision-by-Registration: An Unsupervised Approach to Improve the Precision of Facial Landmark Detectors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 360–368.
- [18] R. Donner, M. Reiter, G. Langs, P. Peloschek, and H. Bischof. 2006. Fast active appearance model search using canonical correlation analysis. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 28, 10 (2006), 1690–1694.
- [19] Chi Nhan Duong, Khoa Luu, Kha Gia Quach, and Tien D Bui. 2019. Deep appearance models: A deep boltzmann machine approach for face modeling. *International Journal of Computer Vision* 127, 5 (2019), 437–455.
- [20] Zhen-Hua Feng, Guosheng Hu, Josef Kittler, William Christmas, and Xiao-Jun Wu. 2015. Cascaded collaborative regression for robust facial landmark detection trained using a mixture of synthetic and real images with dynamic weighting. *IEEE Trans. on Image Processing* 24, 11 (2015), 3425–3440.
- [21] Zhen-Hua Feng, Patrik Huber, Josef Kittler, William Christmas, and Xiao-Jun Wu. 2015. Random cascaded-regression copse for robust facial landmark detection. *IEEE Signal Processing Letters* 1, 22 (2015), 76–80.
- [22] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. 2017. Face Detection, Bounding Box Aggregation and Pose Estimation for Robust Facial Landmark Localisation in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [23] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. 2018. Wing Loss for Robust Facial Landmark Localisation with Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2235–2245.
- [24] Zhen-Hua Feng, Josef Kittler, William Christmas, Patrik Huber, and Xiao-Jun Wu. 2017. Dynamic Attention-controlled Cascaded Shape Regression Exploiting Training Data Augmentation and Fuzzy-set Sample Weighting. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2481–2490.
- [25] Zhen-Hua Feng, Josef Kittler, and Xiao-Jun Wu. 2019. Mining Hard Augmented Samples for Robust Facial Landmark Localization With CNNs. *IEEE Signal Processing Letters* 26, 3 (2019), 450–454.
- [26] X. Geng, K. Smith-Miles, Z.H. Zhou, and L. Wang. 2011. Face image modeling by multilinear subspace analysis with missing values. *IEEE Trans. on Systems, Man, and Cybernetics, Part B: Cybernetics* 41, 3 (2011), 881–892.
- [27] J. Gonzalez-Mora, F. De la Torre, R. Murthi, N. Guil, and E. L. Zapata. 2007. Bilinear Active Appearance Models. In *IEEE International Conference on Computer Vision*. 1–8.
- [28] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. 2010. Multi-PIE. *Image and Vision Computing* 28, 5 (2010), 807–813.
- [29] Yanan Guo, Dapeng Tao, Jun Cheng, Alan Dougherty, Yaotang Li, Kun Yue, and Bob Zhang. 2016. Tensor Manifold Discriminant Projections for Acceleration-Based Human Activity Recognition. *IEEE Trans. on Multimedia* 18, 10 (2016), 1907–1918.

- 1977–1987.
- [30] Onur C Hamsici and Aleix M Martinez. 2009. Active appearance models with rotation invariant kernels. In *IEEE International Conference on Computer Vision*. 1003–1009.
 - [31] Richard A Harshman. 1970. Foundations of the PARAFAC procedure: models and conditions for an "explanatory" multimodal factor analysis. *UCLA Working Papers in Phonetics* (1970), 1–84.
 - [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.
 - [33] Xiaofei He, Deng Cai, and Partha Niyogi. 2006. Tensor subspace analysis. *Advances in Neural Information Processing Systems* 18 (2006), 499.
 - [34] Patrik Huber, Zhen-Hua Feng, William Christmas, Josef Kittler, and Matthias Rätzsch. 2015. Fitting 3D Morphable Face Models using local features. In *IEEE International Conference on Image Processing*. IEEE, 1195–1199.
 - [35] Amin Jourabloo and Xiaoming Liu. 2016. Large-Pose Face Alignment via CNN-Based Dense 3D Model Fitting. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4188–4196.
 - [36] P Kittipanya-ngam and TF Cootes. 2006. The effect of texture representations on AAM performance. In *International Conference on Pattern Recognition*, Vol. 2. IEEE, 328–331.
 - [37] Josef Kittler, Patrik Huber, Zhen-Hua Feng, Guosheng Hu, and William Christmas. 2016. 3D Morphable Face Models and Their Applications. In *International Conference on Articulated Motion and Deformable Objects*. Springer, 185–206.
 - [38] T.G. Kolda and B.W. Bader. 2009. Tensor decompositions and applications. *SIAM Rev.* 51, 3 (2009), 455–500.
 - [39] Paul Koppen, Zhen-Hua Feng, Josef Kittler, Muhammad Awais, William Christmas, Xiao-Jun Wu, and He-Feng Yin. 2018. Gaussian mixture 3D morphable face model. *Pattern Recognition* 74 (2018), 617–628.
 - [40] Martin KÄustinger, Paul Wohlhart, Peter M. Roth, and Horst Bischof. 2011. Annotated Facial Landmarks in the Wild: A large-scale, real-world database for facial landmark localization. In *IEEE International Conference on Computer Vision Workshops*. IEEE, 2144–2151.
 - [41] Hyung-Soo Lee and Daijin Kim. 2009. Tensor-based AAM with Continuous Variation Estimation: Application to Variation-Robust Face Recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 31, 6 (2009), 1102–1116.
 - [42] Dahua Lin, Yingqing Xu, Xiaoou Tang, and Shuicheng Yan. 2005. Tensor-based factor decomposition for relighting. In *IEEE International Conference on Image Processing*, Vol. 2. IEEE, II–386.
 - [43] Q. Liu, J. Deng, and D. Tao. 2015. Dual Sparse Constrained Cascade Regression for Robust Face Alignment. *IEEE Transactions on Image Processing* 25, 2 (2015), 700–712.
 - [44] Xiaoming Liu. 2009. Discriminative face alignment. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 31, 11 (2009), 1941–1954.
 - [45] Simon Lucey, Yang Wang, Mark Cox, Sridha Sridharan, and Jeffery F Cohn. 2009. Efficient constrained local model fitting for non-rigid face alignment. *Image and vision computing* 27, 12 (2009), 1804–1813.
 - [46] Jiangjing Lv, Xiaohu Shao, Junliang Xing, Cheng Cheng, and Xi Zhou. 2017. A Deep Regression Architecture With Two-Stage Re-Initialization for High Performance Facial Landmark Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3317–3326.
 - [47] Iain Matthews and Simon Baker. 2004. Active Appearance Models Revisited. *International Journal of Computer Vision* 60, 2 (2004), 135–164.
 - [48] L Matthews, Takahiro Ishikawa, and Simon Baker. 2004. The template update problem. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 26, 6 (2004), 810–815.
 - [49] Chi Nhan Duong, Khoa Luu, Kha Gia Quach, and Tien D Bui. 2015. Beyond principal components: Deep boltzmann machines for face modeling. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4786–4794.
 - [50] Xiushan Nie, Yilong Yin, Jiande Sun, Ju Liu, and Chaoran Cui. 2017. Comprehensive feature-based robust video fingerprinting using tensor model. *IEEE Trans. on Multimedia* 19, 4 (2017), 785–796.
 - [51] Chao Qi, Min Li, Qiushi Wang, Huiquan Zhang, Jinling Xing, Zhifan Gao, and Huailing Zhang. 2018. Facial Expressions Recognition Based on Cognition and Mapped Binary Patterns. *IEEE Access* 6 (2018), 18795–18803.
 - [52] Na Qi, Yunhui Shi, Xiaoyan Sun, and Baocai Yin. 2016. TenSR: Multi-Dimensional Tensor Sparse Representation. In *The IEEE Conference on Computer Vision and Pattern Recognition*. 5916–5925.
 - [53] Michael J Reale, Shaun Canavan, Lijun Yin, Kaoning Hu, and Terry Hung. 2011. A multi-gesture interaction system using a 3-D iris disk model for gaze estimation and an active appearance model for 3-D hand pointing. *IEEE Trans. on Multimedia* 13, 3 (2011), 474–486.
 - [54] Shaoqing Ren, Xudong Cao, Yichen Wei, and Jian Sun. 2014. Face alignment at 3000 fps via regressing local binary features. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1685–1692.
 - [55] S. Romdhani, S. Gong, A. Psarrou, et al. 1999. A multi-view nonlinear active shape model using kernel PCA. In *British Machine Vision Conference*, Vol. 99. 483–492.
 - [56] Sami Romdhani, Alexandra Psarrou, and Shaogang Gong. 2000. On utilising template and feature-based correspondence in multi-view appearance models. In *European Conference on Computer Vision*. Springer, 799–813.

- [57] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 2016. 300 faces in-the-wild challenge: Database and results. *Image and Vision Computing* 47 (2016), 3–18.
- [58] Jason Saragih and Roland Goecke. 2007. A nonlinear discriminative approach to AAM fitting. In *International Conference on Computer Vision*. 1–8.
- [59] P. Sauer, T. Cootes, and C. Taylor. 2011. Accurate Regression Procedures for Active Appearance Models. In *British Machine Vision Conference*. 1–11.
- [60] S. Sclaroff and J. Isidoro. 1998. Active blobs. In *International Conference on Computer Vision*. 1146–1153.
- [61] K. Simonyan and A. Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv abs/1409.1556* (2014).
- [62] Xiaoning Song, Zhen-Hua Feng, Guosheng Hu, Josef Kittler, and Xiao-Jun Wu. 2018. Dictionary integration using 3D morphable face models for pose-invariant collaborative-representation-based classification. *IEEE Transactions on Information Forensics and Security* 13, 11 (2018), 2734–2745.
- [63] Mikkel B Stegmann and Rasmus Larsen. 2003. Multi-band modelling of appearance. *Image and Vision Computing* 21, 1 (2003), 61–67.
- [64] J. Sung, T. Kanade, and D. Kim. 2007. A unified gradient-based approach for combining ASM into AAM. *International Journal of Computer Vision* 75, 2 (2007), 297–309.
- [65] George Trigeorgis, Patrick Snape, Mihalisis A. Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. 2016. Mnemonic Descent Method: A Recurrent Process Applied for End-To-End Face Alignment. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4177–4187.
- [66] L.R. Tucker. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika* 31, 3 (1966), 279–311.
- [67] Georgios Tzimiropoulos, Joan Alabort-i Medina, Stefanos P Zafeiriou, and Maja Pantic. 2014. Active orientation models for face alignment in-the-wild. *IEEE Trans. on Information Forensics and Security* 9, 12 (2014), 2024–2034.
- [68] Georgios Tzimiropoulos and Maja Pantic. 2013. Optimization problems for fast AAM fitting in-the-wild. In *IEEE International Conference on Computer Vision*. 593–600.
- [69] M. Alex O. Vasilescu and Demetri Terzopoulos. 2002. Multilinear Analysis of Image Ensembles: TensorFaces. In *European Conference on Computer Vision*. 447–460.
- [70] M Alex O Vasilescu and Demetri Terzopoulos. 2007. Multilinear projection for appearance-based recognition in the tensor framework. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1–8.
- [71] Daniel Vlasic, Matthew Brand, Hanspeter Pfister, and Jovan Popović. 2005. Face transfer with multilinear models. 24, 3 (2005), 426–433.
- [72] Yang Wang, Simon Lucey, and Jeffrey F Cohn. 2008. Enforcing convexity for improved alignment with constrained local models. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1–8.
- [73] Renliang Weng, Jiwen Lu, Yap-Peng Tan, and Jie Zhou. 2016. Learning Cascaded Deep Auto-Encoder Networks for Face Alignment. *IEEE Trans. Multimedia* 18, 10 (2016), 2066–2078.
- [74] Xinrong Wu, Junwei Zhou, and Yiyun Pan. 2017. Initial Shape Pool Construction for Facial Landmark Localization Under Occlusion. *IEEE Access* 5 (2017), 16649–16655.
- [75] Yuhang Wu, Shishir K Shah, and Ioannis A Kakadiaris. 2018. GoDP: Globally Optimized Dual Pathway deep network architecture for facial landmark localization in-the-wild. *Image and Vision Computing* 73 (2018), 1–16.
- [76] Fan Xin, Wang Hao, Luo Zhongxuan, Li Yuntao, Hu Wenyu, and Luo Daiyun. 2015. Fiducial facial point extraction using a novel projective invariant. *IEEE Transactions on Image Processing* 24, 3 (2015), 1164–1177.
- [77] Fan Xin, Risheng Liu, Zhongxuan Luo, Yuntao Li, and Yuyao Feng. 2018. Explicit Shape Regression with Characteristic Number for Facial Landmark Localization. *IEEE Transactions on Multimedia* 20, 3 (2018), 567–579.
- [78] Xuehan Xiong and Fernando De la Torre. 2013. Supervised descent method and its applications to face alignment. In *IEEE Conference on Computer Vision and Pattern Recognition*. 532–539.
- [79] Junjie Yan, Zhen Lei, Dong Yi, and Stan Li. 2013. Learn to combine multiple hypotheses for accurate face alignment. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*. 392–396.
- [80] Junfeng Zhang and Haifeng Hu. 2018. Joint Head Attribute Classifier and Domain-Specific Refinement Networks for Face Alignment. *ACM Trans. Multimedia Comput. Commun. Appl.* 14, 4, Article 79 (Oct. 2018), 19 pages.
- [81] Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. 2016. Occlusion-Free Face Alignment: Deep Regression Networks Coupled With De-Corrupt AutoEncoders. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3428–3437.
- [82] Jie Zhang, Shiguang Shan, Meina Kan, and Xilin Chen. 2014. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In *European Conference on Computer Vision*. 1–16.
- [83] Shizhan Zhu, Cheng Li, Chen-Change Loy, and Xiaoou Tang. 2016. Unconstrained Face Alignment via Cascaded Compositional Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3409–3417.
- [84] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z. Li. 2016. Face Alignment Across Large Poses: A 3D Solution. In *IEEE Conference on Computer Vision and Pattern Recognition*. 146–155.

Received 0; revised 0; accepted 0