

Deep Cascaded Bi-Network for Face Hallucination

Shizhan Zhu¹, Sifei Liu^{1,2}, Chen Change Loy^{1,3(✉)}, and Xiaoou Tang^{1,3}

¹ Department of Information Engineering,
The Chinese University of Hong Kong, Hong Kong, China
{zs014, ccloy, xtang}@ie.cuhk.edu.hk, sliu32@ucmerced.edu

² University of California, Merced, Merced, USA

³ Shenzhen Institutes of Advanced Technology,
Chinese Academy of Sciences, Shenzhen, China

Abstract. We present a novel framework for hallucinating faces of unconstrained poses and with very low resolution (face size as small as 5pxIOD). In contrast to existing studies that mostly ignore or assume pre-aligned face spatial configuration (e.g. facial landmarks localization or dense correspondence field), we alternately optimize two complementary tasks, namely face hallucination and dense correspondence field estimation, in a unified framework. In addition, we propose a new gated deep bi-network that contains two functionality-specialized branches to recover different levels of texture details. Extensive experiments demonstrate that such formulation allows exceptional hallucination quality on in-the-wild low-res faces with significant pose and illumination variations.

1 Introduction

Increasing attention is devoted to detection of small faces with an image resolution as low as 10 pixels of height [1]. Meanwhile, facial analysis techniques, such as face alignment [2, 3] and verification [4, 5], have seen rapid progress. However, the performance of most existing techniques would degrade when given a low resolution facial image, because the input naturally carries less information, and images corrupted with down-sampling and blur would interfere the facial analysis procedure. Face hallucination [6–13], a task that super-resolves facial images, provides a viable means for improving low-res face processing and analysis, e.g. person identification in surveillance videos and facial image enhancement.

Prior on face structure, or face spatial configuration, is pivotal for face hallucination [6, 7, 12]. The availability of such prior distinguishes the face hallucination task from the general image super-resolution problem [14–21], where the latter lacks of such global prior to facilitate the inference. In this study, we extend the notion of prior to pixel-wise dense face correspondence field. We observe that an informative prior provides a strong semantic guidance that enables face hallucination even from a very low resolution. Here the dense correspondence field is

Throughout this paper, we use the inter-ocular distance measured in pixels (denoted as pxIOD), to concisely and unambiguously represent the face size.

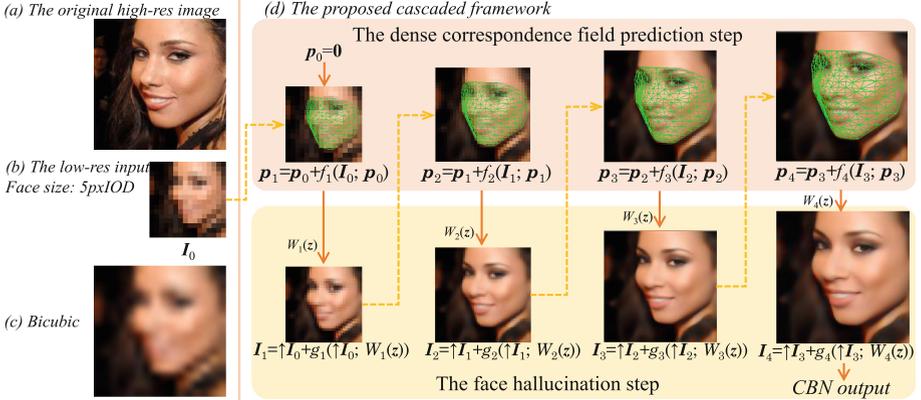


Fig. 1. (a) The original high-res image. (b) The low-res input with a size of $5pxIOD$. (c) The result of bicubic interpolation. (d) An overview of the proposed face hallucination framework. The solid arrows indicate the hallucination step that hallucinates the face with spatial cues, i.e. the dense correspondence field. The dashed arrows indicate the spatial prediction step that estimates the dense correspondence field.

necessary for describing the spatial configuration for its pixel-wise (not by facial landmarks) and correspondence (not by face parsing) properties. The importance of dense field will be reflected in Sect. 3.2. An example is shown in Fig. 1 – even an eye is only visible from a few pixels in a low-res image, one can still recover its qualitative details through inferring from the global face structure.

Nevertheless, obtaining an accurate high-res pixel-wise correspondence field is non-trivial given only the low-res input. First, the definition of the high-res dense field is by itself ill-posed because the gray-scale of each pixel is distributed to adjacent pixels on the interpolated image (Fig. 1(c)). Second, the blur causes difficulties for many existing face alignment or parsing algorithms [3, 22–24] because most of them rely on sharp edge information. Consequently, we face a chicken-and-egg problem - face hallucination is better guided by face spatial configuration, while the latter requires a high resolution face. This issue, however, has been mostly ignored or bypassed in previous works (Sect. 2).

In this study, we propose to address the aforementioned problem with a novel *task-alternating cascaded framework*, as shown as Fig. 1(d). The two tasks at hand - the high-level face correspondence estimation and low-level face hallucination are complementary and can be alternately refined with the guidance from each other. Specifically, motivated by the fact that both tasks are performed in a cascaded manner [15, 23, 25], they can be naturally and seamlessly integrated into an alternating refinement process. During the cascade iteration, the dense correspondence field is progressively refined with the increasing face resolution, while the image resolution is adaptively upscaled guided by the finer dense correspondence field.



Fig. 2. Examples for visualizing the effects of the proposed gated deep bi-network. (a) The bicubic interpolation of the input. (b) Results where only common branches are enabled. (c) Results where only high-frequency branches are enabled. (d) Results of the proposed CBN when both branches are enabled. (e) The original high-res image. Best viewed by zooming in the electronic version.

To better recover different levels of texture details on faces, we propose a new *gated deep bi-network* architecture in the face hallucination step in each cascade. Deep convolutional neural networks have demonstrated state-of-the-art results for image super resolution [14, 15, 17, 18]. In contrast to aforementioned studies, the proposed network consists two functionality-specialized branches, which are trained end-to-end. The first branch, referred as common branch, conservatively recovers texture details that are only detectable from the low-res input, similar to general super resolution. The other branch, referred as high-frequency branch, super-resolves faces with the additional high-frequency prior warped by the estimated face correspondence field in the current cascade. Thanks to the guidance of prior, this branch is capable of recovering and synthesizing un-revealed texture details in the overly low-res input image. A pixel-wise gate network is learned to fuse the results from the two branches. Figure 2 demonstrates the properties of the gated deep bi-network. As can be observed, the two branches are complementary. Although the high-frequency branch synthesizes the facial parts that are occluded (the eyes with sun-glasses), the gate network automatically favours the results from the common branch during fusion.

We refer the proposed framework as *Cascaded Bi-Networks* (CBN) hereafter. We summarize our contribution as follows:

1. While conducting face hallucination or dense face correspondence field is hard on low-res images, we circumvent this problem through a novel task-alternating cascade framework. In comparison to existing approaches, this framework has an appealing property of not assuming pre-aligned inputs or any availability of spatial information (e.g. landmark, parsing map).
2. We propose a gated deep bi-network that can effectively exploit face spatial prior to recover and synthesize texture details that even are not explicitly presented in the low-resolution input.
3. We provide extensive results and discussions to demonstrate and analyze the effectiveness of the proposed approach.

We perform extensive experiments against general super-resolution and face hallucination approaches on various benchmarks. Our method not only achieves high Peak Signal to Noise Ratio (PSNR), but also superior quality perceptually.

Demo codes will be available in our project page <http://mmlab.ie.cuhk.edu.hk/projects/CBN.html>.

2 Related Work

Face hallucination and spatial cues. There is a rich literature in face hallucination [6–13]. Spatial cues are proven essential in most of previous works, and are utilized in various forms. For example, Liu et al. [7, 12] and Jin et al. [6] devised a warping function to connect the face local reconstruction with the high-res faces in the training set. However, a low-res correspondence field¹ may not be sufficient for aiding the high-res face reconstruction process, while obtaining the high-res correspondence field is ill-posed with only a low-res face given. Yang et al. [8] assumed that facial landmarks can be accurately estimated from the low-res face image. This is not correct if the low-res face is rather small (e.g. 5pxIOD), since the gray-scale is severely distributed to the adjacent pixels (Fig. 1(c)). Wang et al. [10] and Kolouri et al. [9] only aligned the input low-res faces with an identical similarity transform (e.g. the same scaling and rotation). Hence these approaches can only handle canonical-view low-res faces. Zhou et al. [26] pointed out the difficulties of predicting the spatial configuration over a low-res input, and did not take any spatial cues into account for hallucination. In contrast to all aforementioned approaches, we adaptively and alternately estimate the dense correspondence field as well as hallucinate the faces in a cascaded framework. The two mutual tasks aid each other and hence our estimation of the spatial cues and hallucination can be better refined with each other.

Cascaded prediction. The cascaded framework is privileged both for image super-resolution (SR) [15, 25] and facial landmark detection [2, 3, 22, 23, 27–30]. For image SR, Wang et al. [15] showed that two rounds of 2× upscaling is better than a single round of 4× upscaling in their framework. For facial landmark detection, the cascaded regression framework has revolutionized the accuracy and has been extended to other areas [31]. The key success of the cascaded regression comes from its coarse-to-fine nature of the residual prediction. As pointed out by Zhang et al. [28], the coarse-to-fine nature can be better achieved by the increasing facial resolution among the cascades. To our knowledge, no existing work has integrated these two related tasks into a unified framework.

The bi-network architecture. The bi-network architecture [32–34] has been explored in various form, such as bilinear networks [35, 36] and two-stream convolutional network [37]. In [35], the two factors, namely object identification and localization, are modeled by the two branches respectively. This is different from our model, where the two factors, the low-res face and the prior, are jointly modeled in one branch (the high-frequency branch), while the other branch (the common branch) only models the low-res face. In addition, the two branches are joined via the gate network in our model, different from the outer-production in [35]. In [37], both spatial and temporal information are modeled by the network,

¹ We assume that we only correspond from pixel to pixel.

which is different from our model, where no temporal information is incorporated. Our architecture also differs from [26]. In [26], the output is the average weighted by a scalar between the result of one branch and the low-res input. Moreover, neither of the two branches utilizes any spatial cues or prior in [26].

3 Cascaded Bi-Network (CBN)

3.1 Overview

Problem and notation. Given a low-resolution input facial image, our goal is to predict its high-resolution image. We introduce the two main entities involved in our framework:

The *facial image* is denoted as a matrix \mathbf{I} . We use $\mathbf{x} \in \mathbb{R}^2$ to denote the (x, y) coordinates of a pixel on \mathbf{I} .

The *dense face correspondence field* defines a pixel-wise correspondence mapping from $M \subset \mathbb{R}^2$ (the 2D face region in the mean face template) to the face region in image \mathbf{I} . We represent the dense field with a warping function [38], $\mathbf{x} = W(\mathbf{z}) : M \rightarrow \mathbb{R}^2$, which maps the coordinates $\mathbf{z} \in M$ from the mean shape template domain to the target coordinates $\mathbf{x} \in \mathbb{R}^2$. See Fig. 3(a,b) for a clear illustration. Following [39], we model the warping residual $W(\mathbf{z}) - \mathbf{z}$ as a linear combination of the dense facial deformation bases, i.e.

$$W(\mathbf{z}) = \mathbf{z} + \mathbf{B}(\mathbf{z})\mathbf{p} \quad (1)$$

where $\mathbf{p} = [p_1 \dots p_N]^\top \in \mathbb{R}^{N \times 1}$ denotes the deformation coefficients and $\mathbf{B}(\mathbf{z}) = [\mathbf{b}_1(\mathbf{z}) \dots \mathbf{b}_N(\mathbf{z})] \in \mathbb{R}^{2 \times N}$ denotes the deformation bases. The N bases are chosen in the AAMs manner [40], that 4 out of N correspond to the similarity transform and the remaining for non-rigid deformations. Note that the bases are pre-defined and shared by all samples. Hence the dense field is actually controlled by the deformation coefficients \mathbf{p} for each sample. When $\mathbf{p} = \mathbf{0}$, the dense field equals to the mean face template.

We use the hat notation ($\hat{\cdot}$) to represent ground-truth in the learning step. For example, we denote the high-resolution training image as $\hat{\mathbf{I}}$.

Framework overview. We propose a principled framework to alternatively refine the face resolution and the dense correspondence field. Our framework consists of K iterations (Fig. 1(d)). Each iteration updates the prediction via

$$\mathbf{p}_k = \mathbf{p}_{k-1} + f_k(\mathbf{I}_{k-1}; \mathbf{p}_{k-1}); W_k(\mathbf{z}) = \mathbf{z} + \mathbf{B}_k(\mathbf{z})\mathbf{p}_k; \quad (2)$$

$$\mathbf{I}_k = \uparrow \mathbf{I}_{k-1} + g_k(\uparrow \mathbf{I}_{k-1}; W_k(\mathbf{z})); \quad (\forall \mathbf{z} \in M_k), \quad (3)$$

where k iterates from 1 to K . Here, Eq. 2 represents the dense field updating step while Eq. 3 stands for the spatially guided face hallucination step in each cascade. ‘ \uparrow ’ denotes the upscaling process ($2 \times$ upscaling with bicubic interpolation in our implementation). All the notations are now appended with the index

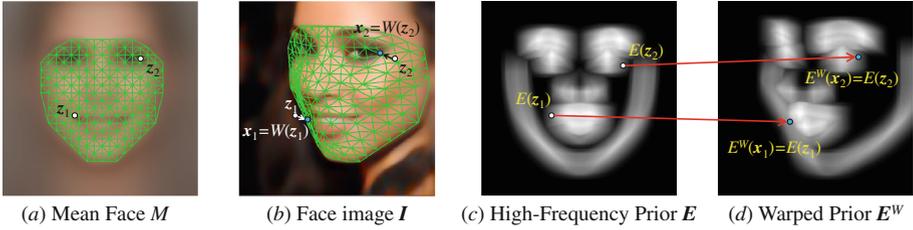


Fig. 3. (a,b) Illustration of the mean face template M and the facial image I . The grid denotes the dense correspondence field $W(\mathbf{z})$. The warping from \mathbf{z} to \mathbf{x} is determined by this warping function $W(\mathbf{z})$. (c,d) Illustration of the high-frequency prior \mathbf{E} and the prior after warping \mathbf{E}^W for the sample image in (b). Note that both \mathbf{E} and \mathbf{E}^W have C channels. Each channel only contains one ‘contour line’. For the purpose of visualization, in this figure, we reduce their channel dimension to one channel with max operation. We leave out all indices k for clarity. Best viewed in the electronic version.

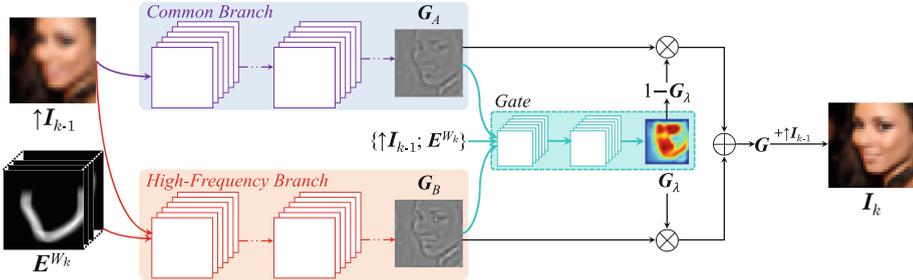


Fig. 4. Architecture of the proposed deep bi-network (for the k -th cascade). It consists of a common branch (blue), a high-frequency branch (red) and the gate network (cyan). (Color figure online)

k to indicate the iteration. A larger k in the notation of \mathbf{I}_k , W_k , \mathbf{B}_k and M_k ² indicates the larger resolution and the same k indicates the same resolution. The framework starts from \mathbf{I}_0 and \mathbf{p}_0 . \mathbf{I}_0 denotes the input low-res facial image. \mathbf{p}_0 is a zero vector representing the deformation coefficients of the mean face template. The final hallucinated facial image output is \mathbf{I}_K .

Model, inference and learning. Our model is composed of functions f_k (dense field estimation) and g_k (face hallucination with spatial cues). The deformation bases \mathbf{B}_k are pre-defined for each cascade and fixed during the whole training and testing procedures. During testing, we repeatedly update the image \mathbf{I}_k and the dense correspondence field $W_k(\mathbf{z})$ (basically the coefficients \mathbf{p}_k) with Eqs. 2, 3. The learning procedure works similarly to the inference but incorporating the learning process of the two functions - g_k for hallucination and f_k for predicting

² We also append the subscript k for M because the mean face template domain M_k do not have the same size in different iteration k .

the dense field coefficients. We present their learning procedures in Sects. 3.2 and 3.3 respectively.

3.2 g_k - Gated Deep Bi-Network: Face Hallucination with Spatial Cues

We propose a gated deep bi-network architecture for face hallucination with the guidance from spatial cues. We train one gated bi-network for each cascade. For the k -th iteration, we take in the input image $\uparrow \mathbf{I}_{k-1}$ and the current estimated dense correspondence field $W_k(\mathbf{z})$, to predict the image residual $\mathbf{G} = \mathbf{I}_k - \uparrow \mathbf{I}_{k-1}$.

As the name indicates, our gated bi-network contains two branches. In contrast to [35] where two branches are joined with outer production, we combine the two branches with a gate network. More precisely, if we denote the output from the common branch (A) and the high-frequency branch (B) as \mathbf{G}_A and \mathbf{G}_B respectively, we combine them with

$$g_k(\uparrow \mathbf{I}_{k-1}; W_k(\mathbf{z})) = \mathbf{G} = (\mathbf{1} - \mathbf{G}_\lambda) \otimes \mathbf{G}_A + \mathbf{G}_\lambda \otimes \mathbf{G}_B, \quad (4)$$

where \mathbf{G} denotes our predicted image residual $\mathbf{I}_k - \uparrow \mathbf{I}_{k-1}$ (i.e. the result of g_k), and \mathbf{G}_λ denotes the pixel-wise soft gate map that controls the combination of the two outputs \mathbf{G}_A and \mathbf{G}_B . We use \otimes to denote element-wise multiplication.

Figure 4 provides an overview of the gated bi-network architecture. Three convolutional sub-networks are designed to predict \mathbf{G}_A , \mathbf{G}_B and \mathbf{G}_λ respectively. The common branch sub-network (blue in Fig. 4) takes in only the interpolated low-res image $\uparrow \mathbf{I}_{k-1}$ to predict \mathbf{G}_A while the high-frequency branch sub-network (red in Fig. 4) takes in both $\uparrow \mathbf{I}_{k-1}$ and the warped high-frequency prior \mathbf{E}^{W_k} (warped according to the estimated dense correspondence field). All the inputs ($\uparrow \mathbf{I}_{k-1}$ and \mathbf{E}^{W_k}) as well as \mathbf{G}_A and \mathbf{G}_B are fed into the gate sub-network (cyan in Fig. 4) for predicting \mathbf{G}_λ and the final high-res output \mathbf{G} .

We now introduce the high-frequency prior and the training procedure of the proposed gated bi-network.

High-frequency prior. We define high-frequency prior as the indication for location with high-frequency details. In this work, we generate high-frequency prior maps to enforce spatial guidance for hallucination. The prior maps are obtained from the mean face template domain. More precisely, for each training image, we compute the residual image between the original image $\hat{\mathbf{I}}$ and the bicubic interpolation of \mathbf{I}_0 , and then warp the residual map into the mean face template domain. We average the magnitude of the warped residual maps over all training images and form the preliminary high-frequency map. To suppress the noise and provide a semantically meaningful prior, we cluster the preliminary high-frequency map into C continuous contours (10 in our implementation). We form a C -channel maps, with each channel carrying one contour. We refer this C -channel maps as our high-frequency prior, and denote it as $E_k(\mathbf{z}) : M_k \rightarrow \mathbb{R}^C$. We use \mathbf{E}_k to represent $E_k(\mathbf{z})$ for all $\mathbf{z} \in M_k$. An illustration of the prior is shown in Fig. 3(c).

Learning the gated bi-network. We train the three parts of convolutional neural networks to predict \mathbf{G}_A , \mathbf{G}_B and \mathbf{G}_λ in our unified bi-network architecture. Each part of the network has a distinct training loss. For training the common branch, we use the following loss over all training samples

$$L_A = \|\hat{\mathbf{I}}_k - \uparrow\mathbf{I}_{k-1} - \mathbf{G}_A\|_F^2. \quad (5)$$

The high-frequency branch has two inputs: $\uparrow\mathbf{I}_{k-1}$ and the *warped* high-frequency prior \mathbf{E}^{W_k} (see Fig. 3(d) for illustration) to predict the output \mathbf{G}_B . The two inputs are fused in the channel dimension to form a $(1 + C)$ -channel input. We use the following loss over all training samples

$$L_B = \sum_{c=1}^C \|(\mathbf{E}^{W_k})_c \otimes (\hat{\mathbf{I}}_k - \uparrow\mathbf{I}_{k-1} - \mathbf{G}_B)\|_F^2, \quad (6)$$

where $(\mathbf{E}^{W_k})_c$ denotes the c -th channel of the warped high-frequency prior maps. Compared to the common branch, we additionally utilize the prior knowledge as input and only penalize over the high-frequency area. Learning to predict the gate map \mathbf{G}_λ is supervised by the final loss

$$L = \|\hat{\mathbf{I}}_k - \uparrow\mathbf{I}_{k-1} - \mathbf{G}\|_F^2. \quad (7)$$

We train the proposed gated bi-network with three steps. Step *i*: We only enable the supervision from L_A (Eq. 5) to pre-train the common branch; Step *ii*: We only enable L_B (Eq. 6) to pre-train the high-frequency branch; Step *iii*: We finally fine-tune the whole gated bi-network with the supervision from L (Eq. 7). In the last step, we set the learning rate of the parameters related to the gate map to be 10 times as the parameters in the two branches. Note that we can still use back-propagation to learn the whole bi-network in our last step.

3.3 f_k - Dense Field Deformation Coefficients Prediction

We apply a simple yet effective strategy to update the correspondence field coefficients estimation (f_k). Observing that predicting a sparse set of facial landmarks is more robust and accurate under low resolution, we transfer the facial landmarks deformation coefficients to the dense correspondence field. More precisely, we simultaneously obtain two sets of N deformation bases: $\mathbf{B}_k(\mathbf{z}) \in \mathbb{R}^{2 \times N}$ for the dense field, and $\mathbf{S}_k(l) \in \mathbb{R}^{2 \times N}$ for the landmarks, where l is the landmark index. The bases for the dense field and landmarks are one-to-one related, i.e. both $\mathbf{B}_k(\mathbf{z})$ and $\mathbf{S}_k(l)$ share the same deformation coefficients $\mathbf{p}_k \in \mathbb{R}^N$:

$$W_k(\mathbf{z}) = \mathbf{z} + \mathbf{B}_k(\mathbf{z})\mathbf{p}_k; \quad \mathbf{x}_k(l) = \bar{\mathbf{x}}_k(l) + \mathbf{S}_k(l)\mathbf{p}_k, \quad (8)$$

where $\mathbf{x}_k(l) \in \mathbb{R}^2$ denotes the coordinates of the l -th landmark, and $\bar{\mathbf{x}}_k(l)$ denotes its mean location.

To predict the deformation coefficients \mathbf{p}_k in each cascade k , we utilize the powerful cascaded regression approach [23] for estimation. A Gauss-Newton

steepest descent regression matrix \mathbf{R}_k is learned in each iteration k to map the observed appearance to the deformation coefficients update:

$$\mathbf{p}_k = \mathbf{p}_{k-1} + f_k(\mathbf{I}_{k-1}; \mathbf{p}_{k-1}) = \mathbf{p}_{k-1} + \mathbf{R}_k(\phi(\mathbf{I}_{k-1}; \mathbf{x}_{k-1}(l)|_{l=1,\dots,L}) - \bar{\phi}), \quad (9)$$

where ϕ is the shape-indexed feature [2, 27] that concatenates the local appearance from all L landmarks, and $\bar{\phi}$ is its average over all the training samples.

To learn the Gauss-Newton steepest descent regression matrix \mathbf{R}_k , we follow [23] to learn the Jacobian \mathbf{J}_k and then obtain \mathbf{R}_k via constructing the project-out Hessian: $\mathbf{R}_k = (\mathbf{J}_k^\top \mathbf{J}_k)^{-1} \mathbf{J}_k^\top$. We refer readers to [23] for more details.

It is worth mentioning that the face flow method [39] that applies a landmark-regularized Lucas-Kanade variational minimization [38] is also a good alternative to our problem. Since we have obtained satisfying results with our previously introduced deformation coefficients transfer strategy, which is purely discriminative and much faster than face flow (8 ms per cascade in our approach v.s. 1.4 s for face flow), we use the coefficients transfer approach in our experiments.

4 Experiments

Datasets. Following [6, 8], we choose the following datasets that contain both in-the-wild and lab-constrained faces with various poses and illuminations.

1. *MultiPIE* [41] was originally proposed for face recognition. A total of more than 750,000 faces from 337 identities are collected under lab-constrained environment. We use the same 351 images as used in [8] for evaluation.
2. *BioID* [42] contains 1521 faces also collected in the constrained settings. We use the same 100 faces as used in [6] for evaluation.
3. *PubFig* [43] contains 42461 faces (the evaluation subset) from 140 identities originally for evaluating face verification and later used for evaluating face hallucination [8]. The faces are collected from the web and hence in-the-wild. Due to the existence of invalid URLs, we use a total of 20991 faces for evaluation. Further, following [6], we use *PubFig83* [44], a subset of PubFig with 13838 images, to experiment with input blurred by unknown Gaussian kernel. Similar to [6], we test with the same 100-image-subset of PubFig83.
4. *Helen* [45] contains 2330 in-the-wild faces with high resolution. The mean face size is as large as 275pxIOD. We evaluate with the 330-image test set.

Metric. We follow existing studies [6, 8, 12, 14, 15] to adopt PSNR (dB) and only evaluate on the luminance channel of the facial region. The definition of the facial region is the same as used in [6]. Similar to [6], SSIM is not reported for in-the-wild faces due to irregular facial shape.

Implementation details. Our framework consists of $K = 4$ cascades, and each cascade has its specific learned network parameters and Gauss-Newton steepest descent regression matrix. During training, our model requires two parts of training data, one for training the cascaded dense face correspondence field,

and the other for training the cascaded gated bi-networks for hallucination. The model is trained by iterating between these two parts of the training data. For the former part, we use the training set from 300W [46] (the same 2811 images used in [23]) for estimating deformation coefficient and BU4D [47, 48] dataset for obtaining dense face correspondence basis (following [39]). For the latter part, as no manual labeling is required, we leverage the existing large face database CelebA [49] for training the gated bi-network.

4.1 Comparison with State-of-the-Art Methods

We compare our approach with two types of methods: (I) general super resolution (SR) approaches and (II) face hallucination approaches. For SR methods, we compare with the recent state-of-the-art approaches [14, 15, 19, 50] based on the original released codes. For face hallucination methods, we report the result of [6, 12, 51] by directly referring to the literature [6]. We compare with [8, 52] by following the implementation of [8]. We re-transform the input face to canonical-view if the method assumes the input must be aligned. Hence, such method would enjoy extra advantages in the comparison. If the method requires exemplars, we feed in the same in-the-wild samples in our training set. We observe that such in-the-wild exemplars improve the exemplar-based baseline methods compared to their original implementation. Codes for [7] is not publicly available. Similar to [6], we provide the qualitative comparison with [7].

We conduct the comparison in two folds: 1. The input is the down-sampled version of the original high-res image as many of the previous SR methods are evaluated on [7, 14, 15, 19, 50] (referred as the conventional SR setting, Sect. 4.1); 2. The input is additionally blurred with unknown Gaussian kernel before down-sampling as in [6, 8, 12] (referred as the Gaussian-blurred setting, Sect. 4.1).

The Conventional SR Evaluation Setting We experiment with two scenarios based on two different types of input face size configuration:

1. *Fixed up-scaling factors* – The input image is generated by resizing the original image with a fixed factor. For MultiPIE, following [8] we choose the fixed factor to be 4. For the in-the-wild datasets (PubFig and Helen), we evaluate for scaling factors of 2, 3 and 4 as in [14, 15, 19, 50] (denoted as $2\times$, $3\times$, $4\times$ respectively in Table 1). In this case, different inputs might have different face sizes. The proposed CBN is flexible to handle such scenario. Other existing face hallucination approaches [8, 12, 51, 52] cannot handle different input face sizes and their results in this scenario are omitted.
2. *Fixed input face sizes* – Similar to the face hallucination setting, the input image is generated by resizing the original image to ensure the input face size to be fixed (e.g. 5 or 8 pxIOD, denoted as 5/8px in Table 1). Hence, the required up-scaling factor is different for each input. For baseline approaches, [15] can naturally handle any up-scaling requirement. For other approaches, we train a set of models for different up-scaling factors. During testing, we pick up the most suitable model based on the specified up-scaling factor.

Table 1. Results under the conventional SR setting (for Sect. 4.1). Numbers in the parentheses indicate SSIM and the remaining represent PSNR (dB). The first part of the results are from Scenario 1 where each method super-resolves for a fixed factor ($2\times$, $3\times$ or $4\times$), while the latter part are from Scenario 2 that each method begins from the same face size (5 or 8 pxIOD, i.e. the inter-ocular distance is 5 or 8 pixels). The omitted results (-) are due to their incapability of handling varying input face size.

Dataset	Input size	Bicubic	(I) General super-resolution				(II) Face hallucination				CBN
			A+ [50]	SRCNN [14]	CSCN [15]	NBF [19]	PCA [12,51]	[52]	[8]		
MultiPIE	4 \times	33.66	34.53	34.75	35.10	34.73	33.98	34.07	34.31	35.65	
		(.900)	(.910)	(.913)	(.920)	(.912)	(.904)	(.907)	(.903)	(.926)	
PubFig	2 \times	34.78	35.89	36.12	36.47	35.98	-	-	-	36.66	
	3 \times	31.52	32.02	32.13	32.88	32.09	-	-	-	33.17	
	4 \times	29.61	30.02	30.15	30.79	30.16	-	-	-	31.28	
HELEN	2 \times	41.96	42.77	42.95	43.37	43.01	-	-	-	43.51	
	3 \times	38.52	38.89	39.10	39.57	39.15	-	-	-	39.78	
	4 \times	36.59	36.81	36.87	37.61	36.89	-	-	-	37.94	
MultiPIE	5px	25.39	25.63	25.72	25.93	25.75	25.62	25.83	25.72	27.14	
		(.752)	(.767)	(.771)	(.773)	(.769)	(.767)	(.774)	(.769)	(.808)	
PubFig	8px	22.32	22.79	22.98	23.25	23.08	23.37	23.57	23.10	26.83	
	5px	20.63	20.96	21.07	21.33	21.04	21.42	21.58	21.19	25.31	
HELEN	8px	21.86	22.24	22.47	22.69	22.53	22.95	23.01	22.62	26.36	
	5px	20.28	20.50	20.59	20.84	20.57	21.09	21.13	20.64	25.09	

Table 2. Results under the Gaussian-blur setting (for Sect. 4.1). Numbers in parentheses indicate SSIM and the remaining represent PSNR (dB). Settings adhere to [6]. For a fair comparison, we feed in the same number of *in-the-wild* exemplars from CelebA when evaluating [8], instead of the originally used MultiPIE in the released codes.

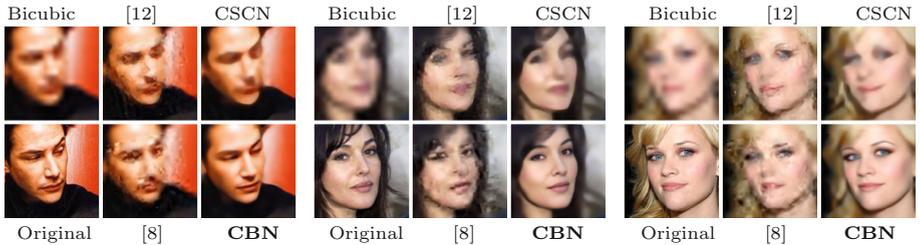
Dataset	Bicubic	(I) General super-resolution				(II) Face hallucination				CBN
		A+ [50]	SRCNN [14]	CSCN [15]	NBF [19]	PCA [12,51]	[52]	[8]	[6]	
BioID	19.67	20.47	20.59	20.86	20.60	21.51	21.77	20.01	22.32	24.55
	(.670)	(.684)	(.685)	(.695)	(.688)	(.770)	(.776)	(.689)	(.810)	(.852)
PubFig83	24.78	25.20	25.22	25.65	25.47	25.72	25.83	25.02	26.17	29.83

We need to point out that the latter scenario is more challenging and appropriate for evaluating a face hallucination algorithm, because recovering the details of the face with the size of 5/8pxIOD is more applicable for low-res face processing applications. In the former scenario, the input face is not small enough (as revealed in the bicubic PSNR in Table 1), such that it is more like a facial image enhancement problem rather than the challenging face hallucination task.

We report the results in Table 1, and provide qualitative results in Fig. 5. As can be seen from the results, our proposed CBN outperforms all general SR and face hallucination methods in both scenarios. The improvement is especially significant in the latter scenario because our incorporated face prior is more critical when hallucinating face from very low resolution. We observe that the general SR algorithms did not obtain satisfying results because they take full efforts to recover only the detectable high-frequency details, which obviously contain noise.

Table 3. PSNR results (dB) of in-house comparison of the proposed CBN (for Sect. 4.2).

Dataset	1a. Only common branch i.e. Vanilla cascaded CNN	1b. Only high-freq. branch	2. Fixed correspondence	3. Single cascade	Full Model
PubFig	23.76	24.66	23.85	22.09	25.31
HELEN	23.57	24.53	23.77	21.83	25.09
PubFig83	28.06	29.31	28.34	26.70	29.83

**Fig. 5.** Qualitative results from PubFig/HELEN with input size 5pxIOD (for Sect. 4.1, detailed results refer Table 1). Best viewed by zooming in the electronic version.

In contrast, our approach recovers the details according to the high-frequency prior as well as the estimated dense correspondence field, thus achieving better performance. The existing face hallucination approaches did not perform well either. In comparison to the evaluation under the constrained or canonical-view condition (e.g. [8]), we found that these algorithms are more likely to fail under in-the-wild setting with substantial shape deformation and appearance variation.

The Gaussian-Blur Evaluation Setting It is also important to explore the capability of handling blurred input images [53]. Our method demonstrates certain degrees of robustness toward unknown Gaussian blur. Specifically, in this section, we still adopt the same model as in Sect. 4.1, with no extra efforts spent in the training to specifically cope with blurring. To compare with [6], we add Gaussian blur to the input facial image in the same way as [6]. The experimental settings are precisely the same as in [6] - the input faces have the same size (around 8pxIOD); the up-scaling factor is set to be 4; and σ for Gaussian blur kernel is set to be 1.6 for *PubFig83* and 2.4 for *BioID*. Additional Gaussian noise with $\eta = 2$ is added in *BioID*. We note that our approach only uses single frame for inference, unlike multiple frames in [6].

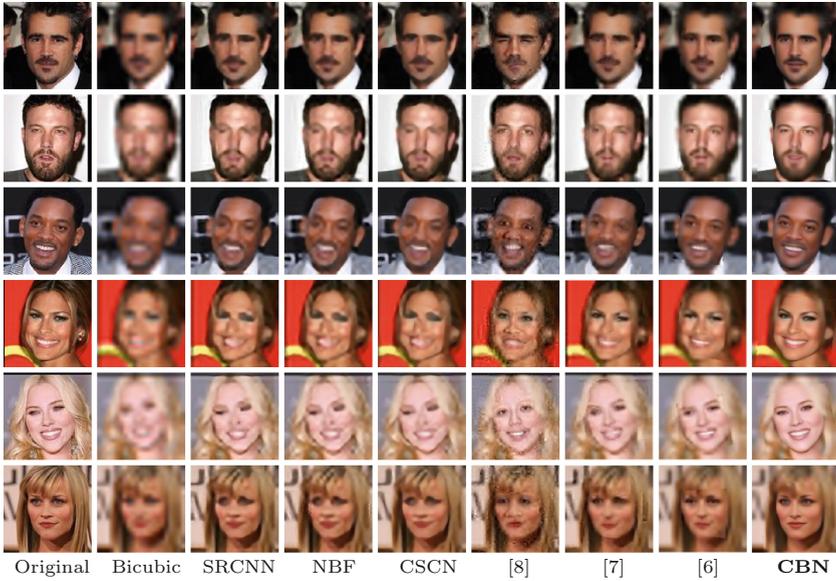


Fig. 6. Qualitative results from the PubFig83 dataset (for Sect. 4.1, detailed results refer Table 2). The six test samples presented are chosen by strictly following [6].

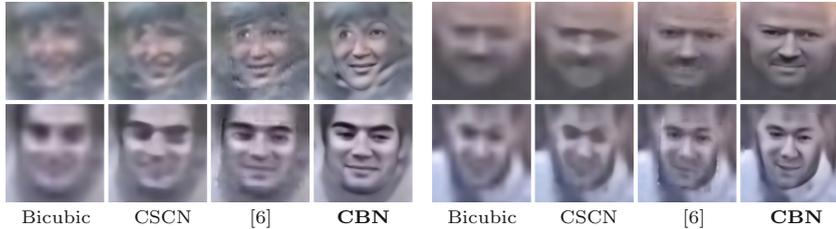


Fig. 7. Qualitative results for real surveillance videos (for Sect. 4.1). The test samples are directly imported from [6]. Best viewed by zooming in the electronic version.

We summarize the results in Table 2. Qualitative results are shown in Fig. 6. From the results it is observed that again CBN significantly outperforms all the compared approaches. We attribute the robustness toward the unknown Gaussian blur on the spatial guidance provided by the face high-frequency prior.

Taking advantages of such robustness of our approach, we further test the proposed algorithm over the faces from real surveillance videos. In Fig. 7, we compare our result with [6, 15]. Note that the presented test cases are directly imported from [6]. Again, our result demonstrates the most appealing visual quality compared to existing state-of-the-art approaches, suggesting the potential of our proposed framework in real-world applications.

Run Time The major time cost of our approach is consumed on the forwarding process of the gated deep bi-networks. On a single core i7-4790 CPU, the face hallucination steps for the four cascades (from 5pxIOD to 80pxIOD) require 0.13 s, 0.17 s, 0.70 s, 2.76 s, respectively. The time cost of the dense field prediction steps is negligible compared to the hallucination step. Our framework totally consumes 3.84 s, which is significantly faster than existing face hallucination approaches, for examples, 15–20 min for [6], 1 min for [8], 8 min for [12], thanks to CBN’s purely discriminative inference procedure and the non-exemplar and parametric model structure.

4.2 An Ablation Study

We investigate the effects of three important components in our framework:

1. **Effects of the gated bi-network.** (a) We explore the results if we replace the cascaded gated bi-network with the vanilla cascaded CNN, in which only the common branch (the blue branch in Fig. 4) is remained. In this case, the spatial information, i.e. the dense face correspondence field is not considered or optimized at all. (b) We also explore the case where only the high-frequency branch (the red branch in Fig. 4) is remained.
2. **Effects of the progressively updated dense correspondence field.** In our framework, the pixel-level correspondence field is refined progressively to better facilitate the subsequent hallucination process. We explore the results if we only use the correspondence estimated from the input low-res image³. In this case, the spatial configuration estimation is not updated with the growth of the resolution.
3. **Effects of the cascade.** The cascaded alternating framework is the core for our framework. We explore the results if we train one network and directly super resolve the input to the required size. High-frequency prior is still used in this baseline. We observe an even worse result without this prior.

We present the results in Table 3. The experimental setting follows the same setting in Sect. 4.1 - The PubFig and HELEN datasets super-resolve from 5pxIOD while the PubFig83 dataset up-scales 4 times with unknown Gaussian blur. The results suggest that all components are important to our proposed approach.

4.3 Discussion

Despite the effectiveness of our method, we still observe a small set of failure cases. Figure 8 illustrates three typical types of failure: (1) Over-synthesis of occluded facial parts, e.g., the eyes in Fig. 8(a). In this case, the gate network might have been misled by the light-colored sun-glasses and therefore favours the results from the high-frequency branch. (2) Ghosting effect, which is caused by inaccurate spatial prediction under low-res. It is rather challenging to localize

³ As the correspondence estimation is by itself a cascaded process, in this case, we re-order the face corresponding cascades before the super resolution cascades.



Fig. 8. Three types of representative failure cases of our approach (for Sect. 4.3).

facial parts with very large head pose in the low-res image. (3) Incorrect details such as gaze direction. We found that there is almost no reliable gaze direction information presented in the input. Our method only synthesizes the eyes with the most probable gaze direction. We leave it as future works to address the aforementioned drawbacks.

5 Conclusion

We have presented a novel framework for hallucinating faces under substantial shape deformation and appearance variation. Owing to the specific capability to adaptively refine the dense correspondence field and hallucinate faces in an alternating manner, we obtain state-of-the-art performance and visually appealing qualitative results. Guided by the high-frequency prior, our framework can leverage spatial cues in the hallucination process.

Acknowledgment. This work is partially supported by SenseTime Group Limited and the Hong Kong Innovation and Technology Support Programme.

References

1. Yang, S., Luo, P., Loy, C.C., Tang, X.: Wider face: A face detection benchmark. arXiv preprint [arXiv:1511.06523](https://arxiv.org/abs/1511.06523) (2015)
2. Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. *IJCV* **107**(2), 177–190 (2014)
3. Xiong, X., De la Torre, F.: Supervised descent method and its applications to face alignment. In: *CVPR*, pp. 532–539 (2013)
4. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: a unified embedding for face recognition and clustering. In: *CVPR* (2015)
5. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: DeepFace: Closing the gap to human-level performance in face verification. In: *CVPR* (2014)
6. Jin, Y., Bouganis, C.S.: Robust multi-image based blind face hallucination. In: *CVPR* (2015)
7. Tappen, M.F., Liu, C.: A bayesian approach to alignment-based image hallucination. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012, Part VII. LNCS*, vol. 7578, pp. 236–249. Springer, Heidelberg (2012)
8. Yang, C.Y., Liu, S., Yang, M.H.: Structured face hallucination. In: *CVPR* (2013)
9. Kolouri, S., Rohde, G.K.: Transport-based single frame super resolution of very low resolution face images. In: *CVPR* (2015)

10. Wang, X., Tang, X.: Hallucinating face by eigentransformation. *IEEE Trans. Syst. Man Cybern. Part C: Appl. Rev.* **35**, 425–434 (2005)
11. Chakrabarti, A., Rajagopalan, A., Chellappa, R.: Super-resolution of face images using kernel PCA-based prior. *IEEE Trans. Multimedia* **9**(4), 888–892 (2007)
12. Liu, C., Shum, H.Y., Freeman, W.T.: Face hallucination: theory and practice. *IJCV* **75**, 115–134 (2007)
13. Baker, S., Kanade, T.: Hallucinating faces. In: *AFGR* (2000)
14. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. In: *PAMI* (2015)
15. Wang, Z., Liu, D., Yang, J., Han, W., Huang, T.: Deep networks for image super-resolution with sparse prior. In: *ICCV* (2015)
16. Huang, J.B., Singh, A., Ahuja, N.: Single image super-resolution from transformed self-exemplars. In: *CVPR* (2015)
17. Gu, S., Zuo, W., Xie, Q., Meng, D., Feng, X., Zhang, L.: Convolutional sparse coding for image super-resolution. In: *ICCV* (2015)
18. Bruna, J., Sprechmann, P., LeCun, Y.: Super-resolution with deep convolutional sufficient statistics. In: *ICLR* (2016)
19. Salvador, J., Perez-Pellitero, E.: Naive bayes super-resolution forest. In: *ICCV* (2015)
20. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016, PART II. LNCS*, vol. 9906, pp. 391–407. Springer, Heidelberg (2016)
21. Hui, T.W., Loy, C.C., Tang, X.: Depth map super resolution by deep multi-scale guidance. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016, PART III. LNCS*, vol. 9907, pp. 353–369. Springer, Heidelberg (2016)
22. Ren, S., Cao, X., Wei, Y., Sun, J.: Face alignment at 3000 fps via regressing local binary features. In: *CVPR* (2014)
23. Tzimiropoulos, G.: Project-out cascaded regression with an application to face alignment. In: *CVPR* (2015)
24. Smith, B.M., Zhang, L., Brandt, J., Lin, Z., Yang, J.: Exemplar-based face parsing. In: *CVPR*, pp. 3484–3491 (2013)
25. Cui, Z., Chang, H., Shan, S., Zhong, B., Chen, X.: Deep network cascade for image super-resolution. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part V. LNCS*, vol. 8693, pp. 49–64. Springer, Heidelberg (2014)
26. Zhou, E., Fan, H., Cao, Z., Jiang, Y., Yin, Q.: Learning face hallucination in the wild. In: *AAAI* (2015)
27. Dollár, P., Welinder, P., Perona, P.: Cascaded pose regression. In: *CVPR*, pp. 1078–1085 (2010)
28. Zhang, J., Shan, S., Kan, M., Chen, X.: Coarse-to-fine auto-encoder networks (CFAN) for real-time face alignment. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part II. LNCS*, vol. 8690, pp. 1–16. Springer, Heidelberg (2014)
29. Zhu, S., Li, C., Loy, C.C., Tang, X.: Face alignment by coarse-to-fine shape searching. In: *CVPR* (2015)
30. Zhu, S., Li, C., Loy, C.C., Tang, X.: Unconstrained face alignment via cascaded compositional learning. In: *CVPR* (2016)
31. Wang, X., Valstar, M., Martinez, B., Haris Khan, M., Pridmore, T.: Tric-track: tracking by regression with incrementally learned cascades. In: *ICCV* (2015)
32. Tenenbaum, J.B., Freeman, W.T.: Separating style and content with bilinear models. *Neural Comput.* **12**, 1247–1283 (2000)

33. Pirsiavash, H., Ramanan, D., Fowlkes, C.C.: Bilinear classifiers for visual recognition. In: NIPS (2009)
34. Xiong, Y., Zhu, K., Lin, D., Tang, X.: Recognize complex events from static images by fusing deep channels. In: CVPR (2015)
35. Lin, T.Y., RoyChowdhury, A., Maji, S.: Bilinear CNN models for fine-grained visual recognition. In: ICCV (2015)
36. Gao, Y., Bejbom, O., Zhang, N., Darrell, T.: Compact bilinear pooling. In: CVPR (2016)
37. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
38. Alabort-i Medina, J., Zafeiriou, S.: Unifying holistic and parts-based deformable model fitting. In: CVPR (2015)
39. Snape, P., Roussos, A., Panagakis, Y., Zafeiriou, S.: Face flow. In: ICCV (2015)
40. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *TPAMI* **23**(6), 681–685 (2001)
41. Gross, R., Matthews, I., Cohn, J., Kanade, T., Baker, S.: Multi-pie. *Image Vis. Comput.* **28**(5), 807–813 (2010)
42. Jesorsky, O., Kirchberg, K.J., Frischholz, R.W.: Robust face detection using the Hausdorff distance. In: Bigun, J., Smeraldi, F. (eds.) AVBPA 2001. LNCS, vol. 2091, pp. 90–95. Springer, Heidelberg (2001). doi:[10.1007/3-540-45344-X_14](https://doi.org/10.1007/3-540-45344-X_14)
43. Kumar, N., Berg, A.C., Belhumeur, P.N., Nayar, S.K.: Attribute and simile classifiers for face verification. In: CVPR (2009)
44. Pinto, N., Stone, Z., Zickler, T., Cox, D.: Scaling up biologically-inspired computer vision: a case study in unconstrained face recognition on facebook. In: CVPRW (2011)
45. Le, V., Brandt, J., Lin, Z., Bourdev, L., Huang, T.S.: Interactive facial feature localization. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part III. LNCS, vol. 7574, pp. 679–692. Springer, Heidelberg (2012)
46. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: 300 faces in-the-wild challenge: the first facial landmark localization challenge. In: ICCVW (2013)
47. Zhang, X., Yin, L., Cohn, J.F., Canavan, S., Reale, M., Horowitz, A., Liu, P., Girard, J.M.: BP4D-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image Vis. Comput.* **32**(10), 692–706 (2014)
48. Zhang, X., Yin, L., Cohn, J.F., Canavan, S., Reale, M., Horowitz, A., Liu, P.: A high-resolution spontaneous 3d dynamic facial expression database. In: 2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG) (2013)
49. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: ICCV (2015)
50. Timofte, R., De Smet, V., Van Gool, L.: A+: adjusted anchored neighborhood regression for fast super-resolution. In: Cremers, D., Reid, I., Saito, H., Yang, M.-H. (eds.) ACCV 2014. LNCS, vol. 9006, pp. 111–126. Springer, Heidelberg (2015)
51. Capel, D., Zisserman, A.: Super-resolution from multiple views using learnt image models. In: CVPR (2001)
52. Ma, X., Zhang, J., Qi, C.: Hallucinating face by position-patch. *Pattern Recogn.* **43**(6), 2224–2236 (2010)
53. Efrat, N., Glasner, D., Apartsin, A., Nadler, B., Levin, A.: Accurate blur models vs. image priors in single image super-resolution. In: ICCV (2013)