



Masked Linear Regression for Learning Local Receptive Fields for Facial Expression Synthesis

Nazar Khan¹ · Arbish Akram¹ · Arif Mahmood² · Sania Ashraf¹ · Kashif Murtaza¹

Received: 29 October 2018 / Accepted: 16 October 2019 / Published online: 1 November 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Compared to facial expression recognition, expression synthesis requires a very high-dimensional mapping. This problem exacerbates with increasing image sizes and limits existing expression synthesis approaches to relatively small images. We observe that facial expressions often constitute sparsely distributed and locally correlated changes from one expression to another. By exploiting this observation, the number of parameters in an expression synthesis model can be significantly reduced. Therefore, we propose a constrained version of ridge regression that exploits the local and sparse structure of facial expressions. We consider this model as masked regression for learning local receptive fields. In contrast to the existing approaches, our proposed model can be efficiently trained on larger image sizes. Experiments using three publicly available datasets demonstrate that our model is significantly better than ℓ_0 , ℓ_1 and ℓ_2 -regression, SVD based approaches, and kernelized regression in terms of mean-squared-error, visual quality as well as computational and spatial complexities. The reduction in the number of parameters allows our method to generalize better even after training on smaller datasets. The proposed algorithm is also compared with state-of-the-art GANs including Pix2Pix, CycleGAN, StarGAN and GANimation. These GANs produce photo-realistic results as long as the testing and the training distributions are similar. In contrast, our results demonstrate significant generalization of the proposed algorithm over out-of-dataset human photographs, pencil sketches and even animal faces.

Keywords Expression · Face · Mapping · Synthesis · Regression · Masked · Local receptive field · Machine learning · Optimization · Quadratic · Convex · GAN · Adversarial · Generative · Discriminative · Ridge · Linear · Image-to-image translation

Communicated by Xavier Alameda-Pineda, Elisa Ricci, Albert Ali Salah, Nicu Sebe, Shuicheng Yan.

✉ Nazar Khan
nazarkhan@pucit.edu.pk

Arbish Akram
phdcsf18m002@pucit.edu.pk

Arif Mahmood
arif.mahmood@itu.edu.pk

Sania Ashraf
sashraf@pucit.edu.pk

Kashif Murtaza
kashifmurtaza@pucit.edu.pk

¹ Punjab University College of Information Technology (PUCIT), Lahore, Pakistan

² Department of Computer Science, Information Technology University (ITU), Lahore, Pakistan

1 Introduction

Affective human computer interaction requires both recognition as well as synthesis of different facial expressions and emotional states. Facial Expression Synthesis (FES) refers to the process of automatically changing the expression of an input face image to another desired expression (Wang et al. 2003; Susskind et al. 2008). Facial expressions are non-verbal visual cues which supplement or reinforce the meaning of spoken words. Therefore, facial expressions are a central element of visual communication for human and non-human characters (Bermano et al. 2014). Realistic FES is important because of its applications in animation of characters in video games and movies (Pighin et al. 2006; Rizzo et al. 2004) and avatar-based human-computer interaction (Saragih et al. 2011). It is also important in security and surveillance applications for the purpose of identifying persons across varying facial expressions (Elaiwat et al. 2016)

and can be useful in longitudinal face modelling as well (Nhan Duong et al. 2016).

A simple approach for generating expressions is by linear combinations of basis shapes each controlled by a scalar weight (Barsoum et al. 1997; Blanz et al. 1999). These linear weights may be considered as facial model parameters. Another face model parameterization is to simply represent a face by its vertices, splines and polygons (Patel and Zaveri 2010). However, this representation has significantly more degrees of freedom than an actual facial expression. Some facial animation systems use the Facial Action Coding System (FACS) (Ekman and Friesen 2013) to estimate facial models from motion capture data (Havaldar 2006). However, such methods require motion capture data along with extensive calibration and data cleansing. This limits their applicability in most cases where only face images and their expression labels are available.

Most existing approaches to FES involve separating the problem into two parts, a geometry adaptation step based on a 3D mesh or facial landmarks and then an appearance adaptation step based on texture. In contrast to these techniques, we present a landmark-free FES method which only requires aligned face images. That is, landmarks are used for alignment but not for any subsequent expression synthesis, mapping or warping.

FES has recently experienced a resurgence due to the introduction of Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Mirza 2014). GANs have enabled a new level of photo-realism by encouraging the generated images to be close to the manifolds of the real images instead of being close to the conditional mean, which may not be photo-realistic. GANs have been shown to be effective in a wide variety of applications such as image editing (Zhu et al. 2016), deblurring (Kupyn et al. 2018) and super-resolution (Ledig et al. 2017). They have been used for facial expression synthesis under the framework of image-to-image translation (Isola et al. 2016; Zhu et al. 2017; Choi et al. 2018). While GANs can generate photo-realistic expressions if the distribution of test images remains similar to the training images, their performance may degrade if the distribution of test images varies as can be seen in Fig. 1.

There is an important distinction to be made between expression *recognition* which typically maps to $O(1)$ classes and *synthesis* which is a very high-dimensional mapping of $O(mn)$ for $m \times n$ image size. Therefore, synthesis models use lots of parameters (even for small image sizes such as 56×56) and require much larger facial expression datasets than those currently used for learning expression recognition models. In the absence of such large datasets, learning FES models that generalize well requires architectures with relatively fewer parameters as we propose in the current manuscript. A key assumption in our work is that facial expressions often constitute sparsely distributed and locally correlated changes from

a neutral expression. This enables us to limit the number of parameters in the model at appropriate locations and achieve good generalization.

In our model, every output pixel directly observes only a localized region in the input image. In other words, each output pixel has a Local Receptive Field (LRF). This is in contrast to models such as ridge regression and multilayer perceptrons in which each output unit observes all input units and therefore has Global Receptive Fields (GRF). The difference between LRFs and GRFs is illustrated in Fig. 2. The simplicity introduced by LRFs is beneficial for FES since expressions constitute multiple local phenomena—so-called action units. GRFs force a pixel to observe too much unrelated information thereby making the learning task harder than it really should be. Therefore, for some problems, LRFs are sufficient and more effective (LeCun et al. 1998; Coates and Ng 2011) as they lead to less convoluted local minima by inducing a regularization effect.

We enforce sparsity in the model by making all the non-local weights zero. This greatly helps the learning task and improves generalization performance. The concept of locality has helped us to develop a memory-efficient, closed-form solution that is applicable to larger problem sizes.

The proposed model is equivalent to a masked version of ridge regression and hence has a global minimum. Due to LRFs, this minimum can be computed quickly with very low computational complexity using our proposed non-iterative, closed-form solution. Also due to LRFs, the number of parameters in our model becomes extremely small. This is important because real world applications of any good algorithm may be offset by the large number of parameters to be learned and stored. This leads to high computational cost at test time. This is especially true for deep network based GANs that contain a huge number of parameters. This leads to higher spatial and computational complexity at test time. This becomes more challenging if the trained models are to be deployed in resource-constrained environments such as mobile devices and embedded systems with limited memory, computational power, and stored energy. A comparison of the proposed algorithm with four state-of-the-art GAN models including Pix2Pix (Isola et al. 2016), CycleGAN (Zhu et al. 2017), StarGAN (Choi et al. 2018) and GANimation (Pumarola et al. 2019) is shown in Table 1. The proposed algorithm has more than two orders of magnitude fewer number of parameters than each of these GANs. In addition, it is more than two orders of magnitude faster in synthesizing an expression.

In contrast with other approaches, the role of weights and biases in our model can be clearly distinguished. The weights are predominantly used to transform the visible parts of the input expression into the target. The biases are used to insert hidden information such as teeth for a happy expression. The model also adjusts weights according to whether a particular

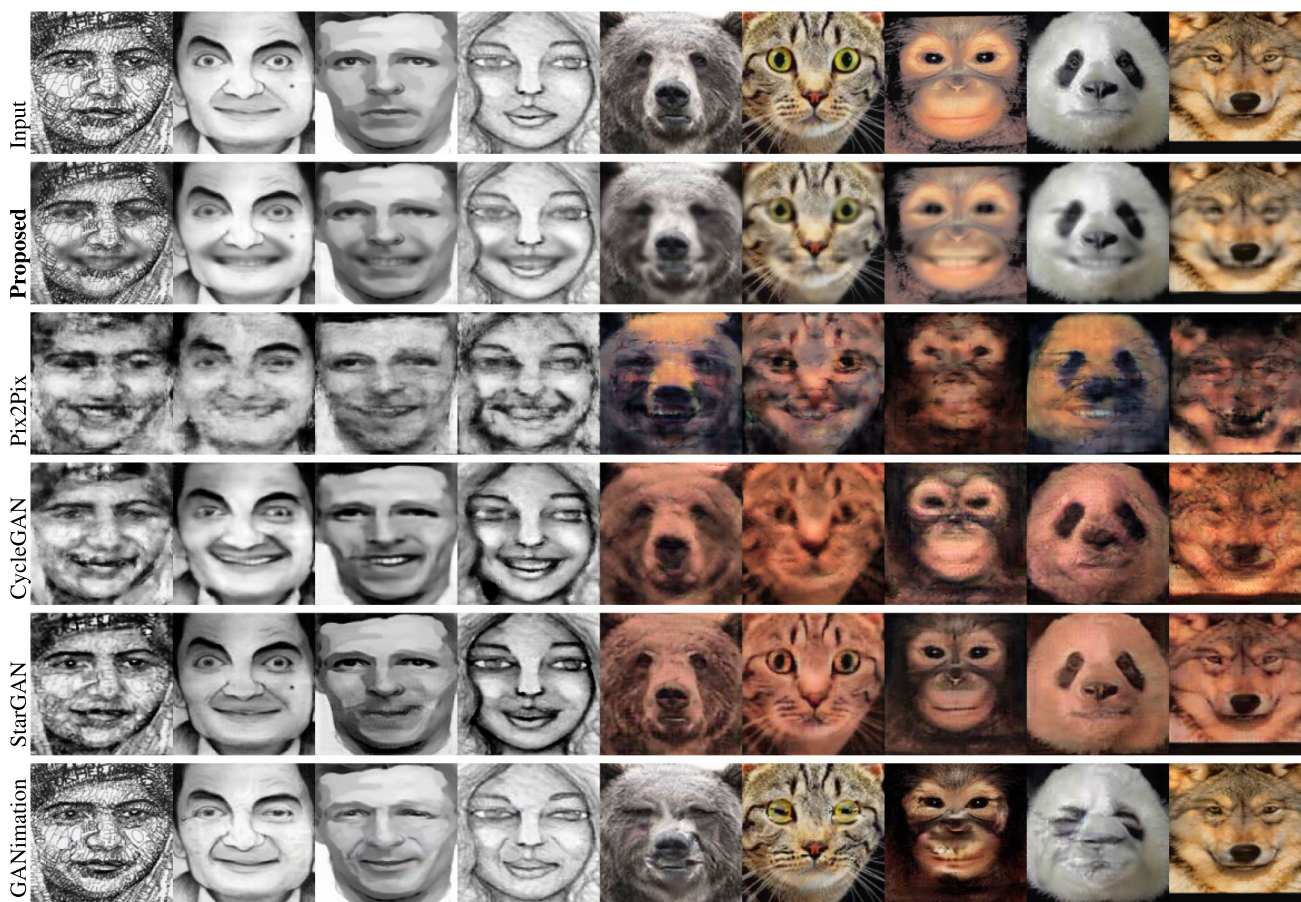


Fig. 1 Comparison of happy expressions synthesized by the proposed algorithm and different GANs. Training of the proposed method was performed on photographs of real human faces only. Columns 1–4: hand-drawn, gray-scale pencil sketches. Columns 5–9: colored animal faces. These results demonstrate the strength of the proposed algorithm in learning essential attributes of happy expressions from real

human face photographs and generalizing to images coming from significantly different distributions. Four state-of-the-art GANs found it very challenging to induce expressions in pencil sketches and for the case of animal faces, no satisfactory expression was induced (Color figure online)

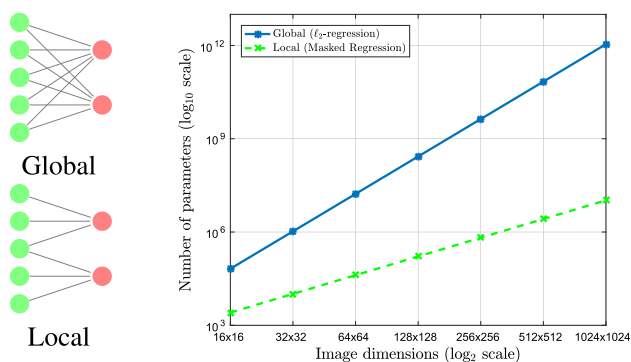


Fig. 2 Left: global versus local receptive fields. Local connections can convey only required information and reduce over-fitting. Right: growth in number of parameters as image size is increased. Local receptive fields remain practical for larger image sizes while regression with global receptive fields becomes impractical even for image sizes as small as 128 × 128 pixels

pixel is relevant for a particular expression. For example, an output pixel ‘looking at’ the mouth region might have a greater role in generating happy expressions than a pixel looking at the top of the forehead. We exploit these locally adaptive weights for identity preserving FES. Experiments performed on three publicly available datasets (Lundqvist et al. 1998; Savran et al. 2008; Lyons et al. 1998) demonstrate that our algorithm is significantly better than l_0 , l_1 and l_2 -regression, SVD based approaches (Tenenbaum and Freeman 2000), and bilinear kernel reduced rank regression (Huang and De la Torre 2010) in terms of mean-squared-error and visual quality.

The proposed approach also exhibits an advantage over GAN models in terms of generalization. Figure 1 shows a comparison of happy expressions synthesized for pencil sketches and several animal faces by the proposed algo-

Table 1 Comparison of different architectures in terms of model size (# of parameters) and average execution time (ms) for images of size 128×128

	Proposed	Pix2Pix	CycleGAN	StarGAN	GANimation
Size ($\times 10^4$)	1.68	4100	780	850	850
Time (ms)	2.70	320	710	580	507

rithm and by Pix2Pix (Isola et al. 2016), CycleGAN (Zhu et al. 2017), StarGAN (Choi et al. 2018) and GANimation (Pumarola et al. 2019). All methods were trained entirely on real human faces, therefore these test images may be considered as out-of-dataset. All four GANs found it very challenging to induce a happy expression in such out-of-dataset images. For the case of animal faces, none of the GANs was able to induce a happy expression. The proposed algorithm generalized well by learning essential attributes of happy expressions and it was able to induce the happy expression in non-human faces as well. Due to the small number of parameters, the proposed algorithm can be easily trained on quite small datasets and in very short time compared to the GANs. Despite using local receptive fields and a masked version of ridge regression, our objective function is still convex and we derive a non-iterative, closed-form solution for the global minimum. This is a fundamental algorithmic contribution of the current work. To the best of our knowledge, the proposed algorithm is novel and no such algorithm has been proposed before for the FES problem. In addition to FES, the proposed formulation can potentially be applied to the broader problem of image-to-image translation. The main contributions of the current work can be summarized as follows:

1. Convex optimization with closed-form solution of global minimum in a single iteration.
2. Extremely low spatial and computational complexity.
3. Trainable on very small datasets.
4. Intuitive interpretation of learned parameters can be exploited to improve results.
5. Good generalization over different types of images that state-of-the-art GANs find very challenging to synthesize.

The rest of the paper is organized as follows. Related work on traditional FES methods and GANs is given in Sect. 2. The proposed Masked Regression (MR) algorithm is given in Sect. 3 and its local receptive field learning formulation is compared with sparse receptive fields in Sect. 4. Experimental details and comparisons with traditional methods are given in Sect. 5. A blur refinement algorithm called Refined Masked Regression (RMR) is given in Sect. 6 and comparison with state-of-the-art GANs is given in Sect. 7. Conclusions and future directions are presented in Sect. 8.

2 Related Work

The Facial Expression Synthesis (FES) research can be divided into blending based techniques and learning based techniques. Blending based techniques primarily merge multiple images to synthesize new expressions (Zhang et al. 2006; Lin and Lin 2011; Pighin et al. 2006). However, such methods require multiple facial landmarks to be pre-identified and do not propose a unified framework for dealing with hidden information, such as teeth, that is usually added in a separate, post-processing step.

For the case of learning-based techniques, FES has received relatively less attention compared to expression recognition or face recognition across varying expressions (Zeng et al. 2009; Jain and Li 2011; Georgakis et al. 2016). Cootes et al. (2001) combined shape and texture information into an Active Appearance Model (AAM). Given facial landmarks, their model can be fit to an unseen face and subsequently used for synthesis and recognition. Liu et al. (2001) computed ratio between a neutral face and a face with an expression at each pixel to obtain an expression ratio image. A new neutral face can then be mapped to the corresponding expression via the ratio image. A bilinear model is employed by Tenenbaum and Freeman (2000) to learn the bases of person-space and expression-space in a single framework using SVD. Wang et al. (2003) learned a trilinear model using higher order SVD. Tensor-based AAM models have been employed for dynamic facial expression synthesis (Lee and Elgammal 2006) and transfer (Zhang and Wei 2012). Facial expression transfer differs from FES since it transfers the expression of a source face onto a different target face (Costigan and Prasad 2014; Zeiler et al. 2011; Wei et al. 2016; Thies et al. 2016). Expression transfer methods include (De La Hunty et al. 2010; Zeiler et al. 2011; Liu et al. 2014; Wei et al. 2016). Suwajanakorn et al. (2015) constructed a controllable 3D model of a person from a large number of photos. While they report impressive results, the large number of per-person training images required for model learning may not always be available. A bilinear model is employed by Tenenbaum and Freeman (2000) to learn the bases of person-space and expression-space in a single framework using SVD. Wang et al. (2003) learned a trilinear model for learning bases of person-space, expression-space and feature-space using higher order SVD. Lee and Elgammal (2006) incorporated the expression manifold with the Tensor-AAM model to synthesize dynamic expressions of the training face. Lee and

Kim (2009) aligned texture with the normalized shape of tensor based AAM. The expression coefficients of a test face were synthesized by linearly combining the expression coefficients of training faces. Zhang and Wei (2012) used Tensor Face combined with an expression manifold to synthesize the dynamic expressions of a training face, then extracted and transferred the dynamic expression details of the training face to the target face. Suwajanakorn et al. (2015) have made a system to construct a controllable 3D model of a person from a large number of photos. While they report impressive results, the large number of training images required for model learning may not always be available. More details and surveys on facial expression synthesis and transfer may be found in (Pantic and Rothkrantz 2000; Deng and Noh 2008; Zeng et al. 2009).

The kernelized regression-based FES method of Huang and De la Torre (2010) learns bases for neutral as well as expression faces. By using the neutral basis they can retain identity preserving details such as glasses and facial marks, using a post processing step. This method also improves generalization by limiting the effective number of free parameters. These properties are shared by our proposed method as well. The Bilinear Kernel Reduced Rank Regression method for static general FES was proposed by Huang and De la Torre (2010). It synthesizes general expressions on the face of a target subject. A relatively similar approach has been proposed by Jampour et al. (2015) for face recognition. Their approach employs local linear regression on localized sparse codes of non-frontal faces to obtain codes of frontal faces. Those codes are then used in a frontal-face classifier to indirectly classify non-frontal faces, though they do not synthesize expressions. In contrast to their approach for face recognition, we propose LRFs for facial expression synthesis.

A deep belief network for facial expression generation has been proposed by Susskind et al. (2008). However, unlike our approach, they cannot synthesize expressions for unseen faces. Their output is usually a semi-controllable mixture of different action units. In our proposed model, we have exact control over which expression is to be synthesized. Due to the use of Restricted Boltzmann Machines their expression generation phase has high computational cost.

The most recent advances in expression synthesis have been achieved via Generative Adversarial Networks (GANs). A typical GAN consists of two competing networks: a generator that takes a random noise vector (and conditioning input) and generates a fake image, and a discriminator network that predicts the probability of an input image being real or fake. These two networks compete against each other to update their weights via minimax learning. Conditional GANs (cGANs) condition their generator and discriminator with additional information such as images or labels. Recently, GAN based frameworks have shown impressive results in image-to-image translation tasks. Pix2pix (Isola

et al. 2016) is a paired image-to-image translation framework based on cGAN and ℓ_1 reconstruction loss. Unpaired image-to-image translation has also been successfully demonstrated by Zhu et al. (2017), Kim et al. (2017), Liu et al. (2017), Yi et al. (2017). CycleGAN (Zhu et al. 2017) learns a mapping between two different domains and incorporates a cycle-consistency loss with an adversarial loss to preserve key attributes between the two domains. Liu et al. (2017) have proposed UNIT framework that combines variational autoencoder with Coupled GAN (Liu and Tuzel 2016). UNIT consists of two generators that share the latent space between two different domains. All of the above-mentioned approaches are designed for translations between two domains at a time. More recently, multi-domain image-to-image translation frameworks have also been proposed. StarGAN (Choi et al. 2018) learns mappings among multiple domains using a single generator conditioned on the target domain labels. The GANimation model of Pumarola et al. (2019) introduced a framework that takes continuous target domain labels in the form of action units and can produce varying degrees of expressions containing multiple action units. Their method is more accurately described as an expression transfer method instead of synthesis. Their method translates a source face via automatically detected action units from a target face. Reliable automatic extraction of action units from face images is a prerequisite for their method to work properly.

Most of these GAN based frameworks share the same problem with other generative models, that is, partial control over the generated images. These methods synthesize the whole image, and therefore also influence attributes in addition to those that were targeted. Strict local control over generated faces is not guaranteed, though some recent GANs have attempted that as well (Shen and Liu 2017; Zhang et al. 2018). Image to image translation using GANs being a very recent research direction, has been quickly progressing.

In the current work we compare the performance of four GANs including Pix2Pix, CycleGAN, StarGAN and GANimation with the proposed Masked Regression (MR) algorithm. These GANs produce excellent results if the test image has similar distribution as the training dataset. As the distribution of test image diverges from the training dataset distribution, the performance of these GANs deteriorates. In contrast to these GANs, the proposed MR algorithm generalises well to very different type of images, can be trained using very small datasets, have a closed form solution with very small spatial as well as computational complexity. To the best of our knowledge, no such technique has been proposed before us for facial expression synthesis.

3 LRF Based Proposed Learning Formulation

We model the FES problem as a linear regression task whereby the output is compared with target faces. Denot-

ing every input face as a vector in \mathbb{R}^D and target face as a vector in \mathbb{R}^K , we can form the design and response matrices $X \in \mathbb{R}^{N \times D}$ and $T \in \mathbb{R}^{N \times K}$ respectively. Here N is the number of training pairs. The design/response matrices are formed by placing the input/target vectors in row-wise fashion. Standard linear regression can also be viewed as a single layer network with global receptive fields (GRF). Our goal is to learn a transformation matrix $W \in \mathbb{R}^{K \times D}$ that minimizes the ℓ_2 -regularized sum of squared errors

$$E^{RR}(W) = \frac{1}{2} \|WX^T - T^T\|_F^2 + \frac{\lambda_2}{2} \|W\|_F^2 \tag{1}$$

where regularization parameter $\lambda_2 > 0$ controls over-fitting and $\|\cdot\|_F^2$ is the squared Frobenious norm of a matrix. This is a quadratic optimization problem with a global minimizer obtained in closed-form as

$$W^{RR} = ((X^T X + \lambda_2 I)^{-1} X^T T)^T \tag{2}$$

As discussed earlier, we posit that transformations from one facial expression to another depend more on local information and less on global information. Therefore, we prune the global receptive fields to retain local weights only. This can be understood by considering faces as 2D images. An output unit at pixel (i, j) is then forced to ‘look at’ only a local window around pixel (i, j) in the input matrix. This can be a 3×3 window covering region $(i - 1, j - 1)$ to $(i + 1, j + 1)$ or an even larger window. Such localized windows are referred to as local receptive fields (LRF) and have been used in Convolutional Neural Networks (LeCun et al. 1998). In order to represent presence or absence of weights, we construct a mask matrix as large as the transformation matrix W where

$$M_{ij} = \begin{cases} 0 & \text{to fix } W_{ij} \text{ to } 0 \\ 1 & \text{to learn } W_{ij} \end{cases} \tag{3}$$

For every pixel in the output, there is a corresponding row in matrix M indexed according to row-major order. This row contains one entry for each pixel in the input which is also indexed according to row-major order. All entries are 0 except for those input pixels that are in the receptive field of the current output pixel. For example, let input and output images both be of size 5×5 . Then in vectorized form, input and output are vectors in \mathbb{R}^{25} . Matrix M will have 25 rows corresponding to output pixels and 25 columns corresponding to input pixels. Figure 3 shows the mask M constructed in this manner. Finally, to incorporate bias terms and treat them as learnable parameters, a column of ones is appended as the last column of M .

Since the local receptive fields obtained by masking the weights are subsets of global receptive fields, learning the

		Input pixel index j in row-major order																								
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
Output pixel index i in row-major order	1	1	1				1	1																		
	2	1	1	1			1	1	1																	
	3		1	1	1			1	1	1																
	4			1	1	1			1	1	1															
	5				1	1				1	1															
	6	1	1				1	1			1	1														
	7	1	1	1			1	1	1			1	1	1												
	8		1	1	1			1	1	1			1	1	1											
	9			1	1	1			1	1	1			1	1	1										
	⋮																									
	19											1	1	1				1	1	1				1	1	1
	20												1	1					1	1					1	1
	21																	1	1					1	1	
	22																		1	1				1	1	
	23																		1	1	1			1	1	1
	24																			1	1	1		1	1	1
	25																				1	1		1	1	1

Fig. 3 Mask M corresponding to input image of size 5×5 , output image of size 5×5 and receptive fields of size 3×3 . For clarity, entries equal to 0 are left blank. If the entry at row i and column j is 1, then output pixel i has input pixel j in its receptive field

optimal weights still involves a quadratic but masked objective function

$$E^{MR}(W) = \frac{1}{2} \|(W \circ M)X^T - T^T\|_F^2 + \frac{\lambda_M}{2} \|W \circ M\|_F^2$$

s.t. $W_{kd} = 0$ if $M_{kd} = 0, \quad 1 \leq k \leq K, \quad 1 \leq d \leq D, \tag{4}$

where \circ denotes the Hadamard product of two equal sized matrices and $\lambda_M > 0$ is a regularization parameter. This formulation fixes unwanted weights to 0 while encouraging the sum-squared-error and magnitudes of the wanted weights to be low. We term this as the *masked regression* (MR) problem. In contrast to ℓ_1 -penalized regression (Tibshirani 1996) that forces most weights to be zero without determining which ones exactly, our proposed masked regression makes specific, pre-determined weights equal to zero. That is, masked regression leads to localized sparsity. Our formulation (4) corresponds exactly to a single layer network with local receptive fields. The reduction in the number of parameters to be learned due to LRFs allows for very fast training of such systems.

Due to the presence of the Hadamard product, writing a closed-form solution for masked regression is not as straightforward as that for ridge regression (2). However, we handle this problem by writing out objective function (4) in terms of individual weights W_{kd} as

$$E^{MR}(W) = \frac{1}{2} \sum_{n=1}^N \sum_{k=1}^K \left\{ \left(\sum_{d=1}^D W_{kd} M_{kd} X_{nd} \right) - T_{nk} \right\}^2 + \frac{\lambda_M}{2} \sum_{k=1}^K \sum_{d=1}^D W_{kd}^2 M_{kd} \tag{5}$$

This allows us to compute entries of the gradient vector $\nabla E^{MR}(W) \in \mathbb{R}^{KD \times 1}$ as

$$\frac{\partial E^{MR}(W)}{\partial W_{ij}} = \sum_{n=1}^N \left\{ \left(\sum_{d=1}^D W_{id} M_{id} X_{nd} \right) - T_{ni} \right\} M_{ij} X_{nj} + \lambda_M W_{ij} M_{ij} \tag{6}$$

where $1 \leq i \leq K$ and $1 \leq j \leq D$. It must be noted that for LRFs looking at $r \times r$ pixels in the previous layer, the summation over d in (5) and (6) need not be performed more than $r^2 \ll D$ times since the corresponding row in mask matrix M contains not more than r^2 ones. Compared to ridge regression and its corresponding global receptive fields, this leads to a significant decrease in memory required for storing the transformation matrix W . We can also compute entries of the Hessian matrix $H \in \mathbb{R}^{KD \times KD}$ as

$$\frac{\partial^2 E^{MR}(W)}{\partial W_{ij} \partial W_{lm}} = \begin{cases} M_{ij} M_{lm} \sum_{n=1}^N X_{nj} X_{nm} & \text{if } i = l \text{ \& } j \neq m \\ M_{ij}^2 \sum_{n=1}^N X_{nj}^2 + \lambda_M M_{ij} & \text{if } i = l \text{ \& } j = m \\ 0 & \text{if } i \neq l \end{cases} \tag{7}$$

where $1 \leq \{i, l\} \leq K$ and $1 \leq \{j, m\} \leq D$. This allows us to compute the optimal solution via a single Newton-Raphson step as

$$\mathbf{w} = -H^{-1} \nabla E^{MR}(W) \tag{8}$$

where $\mathbf{w} \in \mathbb{R}^{KD \times 1}$ represents row-wise concatenated entries of W . That is, $\mathbf{w} = [W^1 \ W^2 \ \dots \ W^K]^T$ where W^k denotes the $1 \times D$ vector containing the values of the k -th row of W . The initial W_0 required for computing ∇E can be set as all zeros since the initial value does not affect the global solution. Therefore, we can find the transformation parameters vector \mathbf{w} by solving the linear system

$$H\mathbf{w} = -\nabla E^{MR}(W). \tag{9}$$

Since H is a block-diagonal matrix with K blocks of size $D \times D$, we can solve for each row separately instead of solving the complete linear system in KD variables involving a $KD \times KD$ system matrix. This means decomposing the larger linear system into K smaller linear systems in D variables involving a $D \times D$ system matrix. These systems can

also be solved in parallel. The k -th linear system can be written as

$$H_{\Omega_k, \Omega_k} \mathbf{w}_{\Omega_k} = -\nabla E_{\Omega_k}^{MR} \tag{10}$$

where Ω_k is the set of indices corresponding to the placement of the k -th row of W in vector \mathbf{w} . Because of the constraints in M , the solution vector \mathbf{w}_{Ω_k} can contain at most r^2 non-zero entries at pre-determined locations corresponding to receptive fields of size $r \times r$. We can solve for these non-zero entries only by removing those rows of $\nabla E_{\Omega_k}^{MR}$ and those rows and columns of H_{Ω_k} that correspond to zero elements of \mathbf{w}_{Ω_k} . This makes the linear system significantly smaller with at most r^2 variables. Denoting the indices of non-zero entries by $\hat{\Omega}_k$, the linear system becomes

$$H_{\hat{\Omega}_k, \hat{\Omega}_k} \mathbf{w}_{\hat{\Omega}_k} = -\nabla E_{\hat{\Omega}_k}^{MR} \tag{11}$$

This decomposition into K extremely small linear systems makes solving the masked regression problem extremely fast and with very low space complexity compared to traditional regression solutions. A comparison of model size of the proposed solution with traditional ridge regression based solutions for increasing problem sizes is shown in Fig. 2. It can be observed that memory required for storing ridge regression parameters quickly exceeds practical limits even for small images. In contrast, the use of LRFs in masked regression keeps the number of parameters and, consequently, memory requirement low even for large images.

4 Local Versus Sparse Receptive Fields

The local receptive fields that we propose can also be viewed as extremely sparse receptive fields with manually designed and fixed localizations. An interesting alternative is to learn sparse receptive fields. Will a sparsely learned topology also converge to our local receptive fields? To answer this question we learn a transformation matrix W that minimizes the ℓ_1 -regularized sum of squared errors

$$\frac{1}{2} \|WX^T - T^T\|_F^2 + \lambda_1 \|W\|_1 \tag{12}$$

where $\lambda_1 > 0$ controls the level of sparsity and therefore also controls over-fitting. The rows of the optimal transformation W will correspond to sparse receptive fields. Error function (12) can be decomposed into a sum of K indepen-

dent ℓ_1 -regression problems that can be solved in parallel. That is

$$\sum_{i=1}^K \frac{1}{2} \|XW^i - T_i\|_2^2 + \lambda_1 \|W^i\|_1 \quad (13)$$

where W^i is the $D \times 1$ vector containing the values in the i -th row of W and T_i is the $N \times 1$ vector containing the values in the i -th column of T . We solve the i -th sub-problem

$$\min_{W^i} \frac{1}{2} \|XW^i - T_i\|_2^2 + \lambda_1 \|W^i\|_1 \quad (14)$$

using the LASSO algorithm (Tibshirani 1996).

In order to provide a fair comparison with masked regression that limits the size of the receptive field, it is better to minimize the ℓ_0 -penalized regression error which can also be decomposed into K separate sub-problems

$$\min_{W^i} \frac{1}{2} \|XW^i - T_i\|_2^2 \text{ s.t. } \|W^i\|_0 \leq \lambda_0 \quad (15)$$

in which hyperparameter $\lambda_0 \in \mathbb{Z}^+$ acts as an upper-bound on the number of non-zero entries in the solution. Therefore, setting $\lambda_0 = r^2$ makes the sparse receptive fields obtained via (15) comparable to the local receptive fields of size $r \times r$ via masked regression. We approximated (15) using the Orthogonal Matching Pursuit algorithm (Pati et al. 1993; Tropp and Gilbert 2007). In the next section, we present a comparison of both ℓ_0 and ℓ_1 regression with the proposed masked regression method.

5 Experiments and Results

In order to provide enough data for learning useful mappings while avoiding over-fitting, we combine three datasets (Lundqvist et al. 1998; Savran et al. 2008; Lyons et al. 1998) containing the neutral and six basic expressions. The basic expressions include afraid, angry, disgusted, happy, sad and surprised. The KDEF dataset (Lundqvist et al. 1998) contains face images of 70 subjects (35 males and 35 females). The Bosphorous dataset (Savran et al. 2008) contains face images of 105 subjects, each subject having up to 35 expressions. The Japanese Female Facial Expression (JAFFE) dataset (Lyons et al. 1998) contains face images of 10 Japanese actresses in neutral and the six basic expressions. By combining these three datasets, we obtain a total of 1116 facial expression images. For each experiment we performed an 80%, 10%, 10% split of the image pairs from the input and target expressions as training, validation and testing sets. We performed alignment of all images with respect to a reference face image. All images were normalized to contain pixel values between 0 and 1.

5.1 Experiments on Grayscale Images

To evaluate the proposed masked regression (MR) method for synthesizing expressions on gray scale images, we compare it with existing regression based techniques including ℓ_0 , ℓ_1 and ℓ_2 -regression as well as Kernelized Reduced Rank Regression (KRRR) and its bilinear extension (BKRRR) (Huang and De la Torre 2010). In KRRR and BKRRR, a rank constraint is used to limit the number of free parameters in a kernel regression model for learning expression bases.

Table 2 Quantitative comparison of output and target images using mean-squared-error (MSE) scaled by 10^2

In	Out	PCA	SSC	ℓ_2	KRRR	BKRRR	ℓ_0	ℓ_1	MR ₃	MR ₅	MR ₇	MR ₉
Neutral	Afraid	2.366	2.365	2.402	2.37	2.30	2.040	1.970	1.813	1.812	1.870	1.940
Neutral	Angry	2.109	2.111	2.111	2.05	2.00	1.810	1.715	1.655	1.598	1.619	1.663
Neutral	Disgusted	2.028	2.028	2.152	2.13	2.12	1.862	1.764	1.625	1.588	1.639	1.694
Neutral	Happy	1.756	1.755	1.901	1.86	1.84	1.563	1.481	1.410	1.412	1.457	1.487
Neutral	Sad	1.623	1.621	1.816	1.77	1.80	1.518	1.429	1.301	1.309	1.350	1.379
Neutral	Surprised	2.499	2.500	2.112	2.07	2.04	1.983	1.789	1.820	1.770	1.791	1.838
Afraid	Neutral	2.537	2.530	1.994	1.85	1.85	1.733	1.611	1.401	1.411	1.500	1.589
Angry	Neutral	2.174	2.175	1.757	1.62	1.62	1.583	1.444	1.429	1.372	1.414	1.465
Disgusted	Neutral	2.218	2.216	1.765	1.61	1.61	1.465	1.371	1.395	1.348	1.393	1.433
Happy	Neutral	1.954	1.954	1.567	1.50	1.50	1.346	1.245	1.251	1.234	1.274	1.311
Sad	Neutral	1.714	1.712	1.505	1.42	1.42	1.375	1.233	1.188	1.171	1.210	1.243
Suprised	Neutral	2.682	2.680	1.776	1.66	1.66	1.697	1.522	1.562	1.496	1.492	1.532
	Mean MSE	2.138	2.137	1.904	1.83	1.82	1.665	1.548	1.487	1.460	1.501	1.548

The algorithms used are PCA, SSC (Tenenbaum and Freeman 2000), ℓ_2 -regression, Kernelized Regression (KRRR and BKRRR) (Huang and De la Torre 2010), ℓ_0 -regression, ℓ_1 -regression and the proposed masked regression (MR). MR_r refers to MR with receptive fields of size $r \times r$ pixels. Bold represents the minimum mean-squared-error among the different methods compared and which method is better.

We also compare with basis learning approaches including PCA and SVD-based bilinear model for separation of style and content (SSC) (Tenenbaum and Freeman 2000). In PCA, a basis is learned for each expression. A test face is mapped to a target expression by projection onto the target expression basis and then reconstructed from the projected coefficients. In SSC, bases are learned for expressions as well as persons.

For ℓ_2 -regression and masked regression, we cross validated the corresponding regularization parameters, λ_2 and λ_M respectively, over 10 equally spaced values between 0.1 and 10. For ℓ_1 -regression, λ_1 was cross-validated from 10^{-3} till 10^2 using 100 equally spaced values in log space. For ℓ_0 -regression, λ_0 was cross-validated for all integers from 1 till the number of training examples. For each method, the best value of the corresponding regularization parameter was used to finally train on the combined training and validation set. Weights learned from this final training were then used to compute mean-squared-errors (MSE) on the test data. We performed 12 experiments corresponding to the mapping of neutral to the six other expressions and vice versa. It can be seen from Table 2 that MR obtains the lowest MSE averaged over the 12 combinations. Visual comparison of different algorithms is presented in Fig. 4. It can be observed that only local receptive fields learned via MR were able to transform the expression while preserving identity and retaining facial details. Figure 5 contains visual results of transforming neutral expressions to the six basic expressions using MR. It

demonstrates that MR is a generic algorithm that can efficiently transform any expression into any other expression while preserving identities and individual facial details.

Role of receptive field size The proposed method can be easily modified to have not-so-local receptive fields. For example, a 3×3 field that looks at every other pixel in a 5×5 window or every third pixel in a 7×7 window. These modifications only involve setting the mask M in Fig. 3 appropriately. This way, an output pixel can ‘observe’ a larger region of the input while using the same number of weights. For example, 9 weights for any $r \times r$ receptive field. This helps to avoid over-fitting by limiting the complexity of the model. Table 2 compares performance of different receptive field sizes. For the dataset used, we observed minimum MSE for 5×5 receptive fields. Employing too large a receptive field increased the MSE since long-range receptive fields fail to capture the local nature of facial expressions.

Role of weights and biases In order to observe the role of only weights, we set the bias values to zero, and observe the resulting mappings. Figure 6 demonstrates that a major role of the weights is to wipe out the original expression while also sometimes inserting subtle intensity changes to affect the new expression. However, the weights cannot efficiently generate unseen content such as teeth that are hidden in the neutral and visible in the happy expressions. This inability to affect hidden expression units is overcome by the biases which

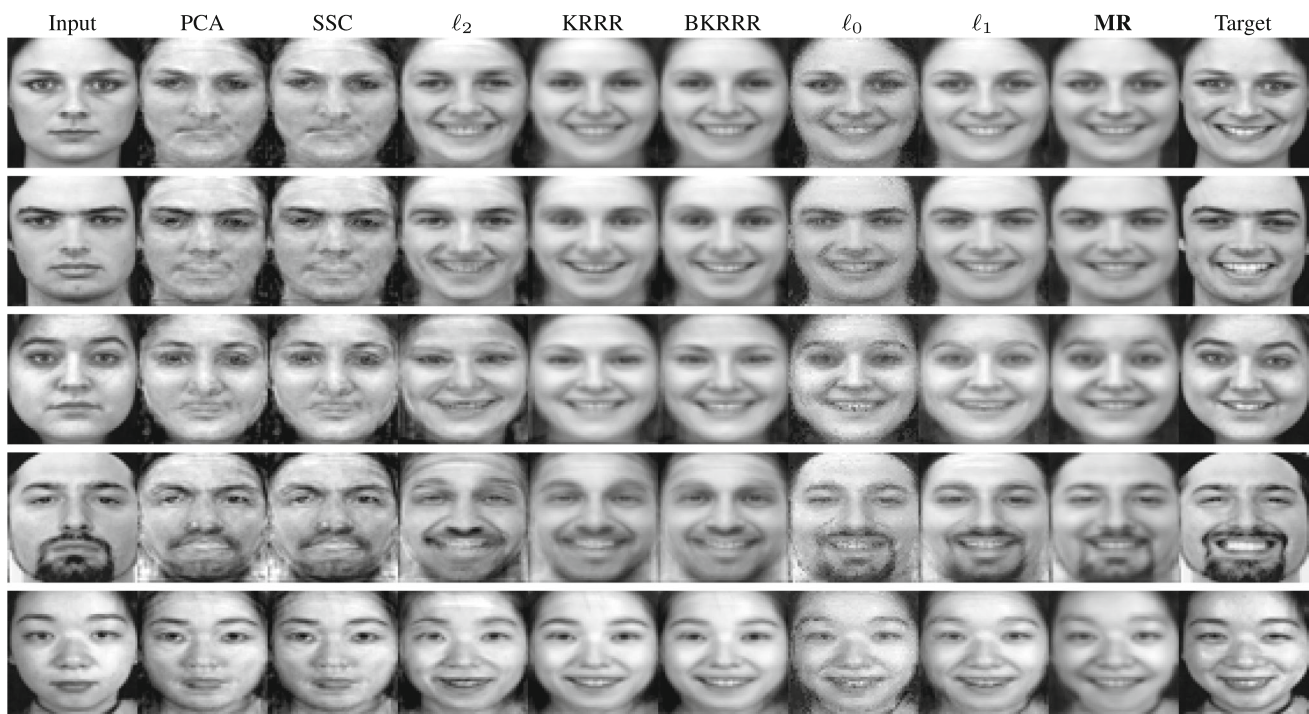


Fig. 4 Comparison of different techniques with the proposed MR method for the neutral to happy mapping. The proposed method was able to transform the expression while preserving identity and retaining facial details the most



Fig. 5 For each neutral input, rows 1, 3 and 5 show expressions generated via proposed MR and rows 2, 4 and 6 show ground-truth. MR effectively transformed expressions while preserving identities and facial details

adjust so that the major role is to produce the remaining, hidden expression units.

Once learned appropriately, the bias remains the same for all input test images. Therefore, it is not surprising to see in Fig. 7 that the learned model exploits the bias only to affect target expression units that the weights could not map. The biases have no role in identity preservation. Figure 8 compares the average absolute intensity of the transformation Wx produced by the weights only with the additive transformation b produced by the biases only. In this figure, for 12 transformations between expressions, weights learned via ℓ_2 -regression have less intensity than learned biases. This is a major cause of loss of identity in the transformed expres-

sion learned via ℓ_2 -regression. In contrast, for the proposed masked regression the transformation via the weights was roughly 5 times more important than the transformation produced by adding the biases only. This is why the proposed MR method has remained the best in preserving identity among all the considered methods.

5.2 Experiments on RGB Images

A straight-forward extension of the proposed method to color images is to learn a separate mapping for each channel. A visual comparison of learning per-channel mappings for MR and other methods in Fig. 9 for RGB images. It can be

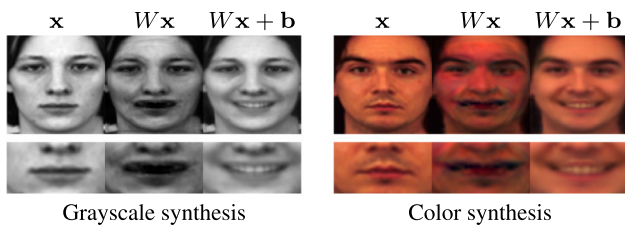


Fig. 6 Left to right x is the neutral input, Wx is the happy expression synthesized with bias $b = 0$ and $Wx + b$ is the complete synthesized happy expression. Second row shows the mouth regions zoomed in. The weights and biases learned via masked regression assumed distinct, complimentary roles. While the weights wiped out the mouth and surrounding regions, the biases (Fig. 7, column 4) then inserted missing information such as teeth. Regions not playing a significant role in the mapping (e.g. hair, forehead) were left unaffected which helps in preserving the identity of the input face (Color figure online)

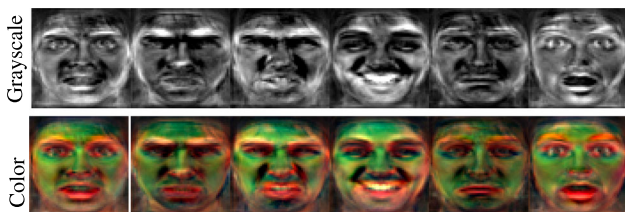
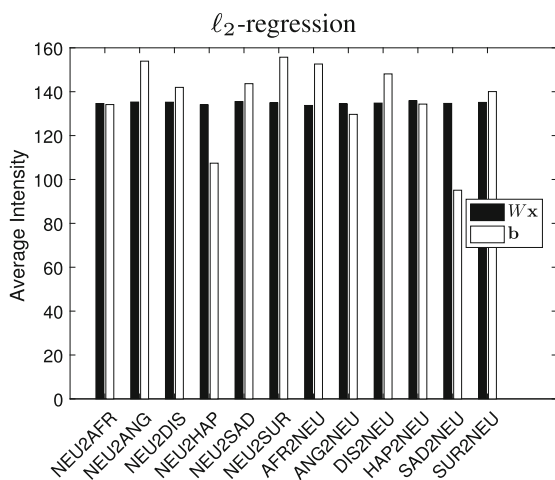


Fig. 7 Biases of masked regression corresponding to six basic expressions. Masked regression exploits the bias for learning expression specific action units such as eyebrow, lip or cheek movements. It is also exploited for adding content that cannot be captured by the weights. For example, appearance of teeth in happy expressions is not represented by any action unit but is still captured by the bias. The biases also represent some arbitrary face but compared to the weights, its importance is low (darker intensities). All images have been post-processed to increase visibility (Color figure online)

observed that MR is most successful in retaining background and other non-facial details that have no role in expression generation. The role of weights and biases for RGB images can be visualized in Figs. 6, 7. Table 3 shows that MR com-



pires favorably against all competing methods in terms of MSE on RGB images. The average training time for the closest competitor (ℓ_1 -regression) was much larger than MR as shown in Table 4.

A cheaper alternative is to replicate the mapping learned from gray-scale images for all color channels. Figure 10 demonstrates the effectiveness of this approach in preventing color leakage. In addition to retaining original color ratios, this solution causes no increase in the number of learnable parameters when scaling from gray-scale to color images. However, this approach can cause the resulting image to lose some of its colorfulness.

A third option is to learn a single mapping between color vectors. The error function for masked regression for multi-channel color images can be written as

$$E^{CMR}(W) = \frac{1}{2} \sum_{c=1}^C \|(W \circ M)X_c^T - T_c^T\|_F^2 + \frac{\lambda_M}{2} \|W \circ M\|_F^2 \tag{16}$$

where C is the number of channels and X_c and T_c are design matrices corresponding to channel c . In this way, the number of learnable parameters remains the same as for a gray-scale mapping but these parameters are now learned from color vectors instead of gray-scale pixels. Results of this approach can be seen in Figs. 1, 18 and 21.

Experiments are performed on other color spaces as well including YCbCr, Lab and HSV. However, best results were observed in the RGB color space. This may be due to the fact that the sparse, distributed, and local nature of facial expressions that is exploited by MR is better represented in the RGB color space.

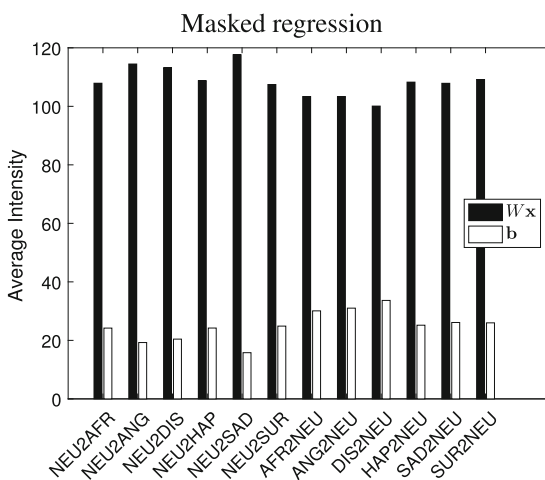


Fig. 8 Relative importance of weights and biases. Over 12 transformations, we compare the average absolute intensity of the transformation produced by the weights with the additive transformation learned as biases. For the case of ℓ_2 -regression, the bias often dominated the

weights, leading to loss of identity. For MR, the transformation via weights was roughly 5 times as important as the transformation produced by adding the bias only. This leads to better identity preservation

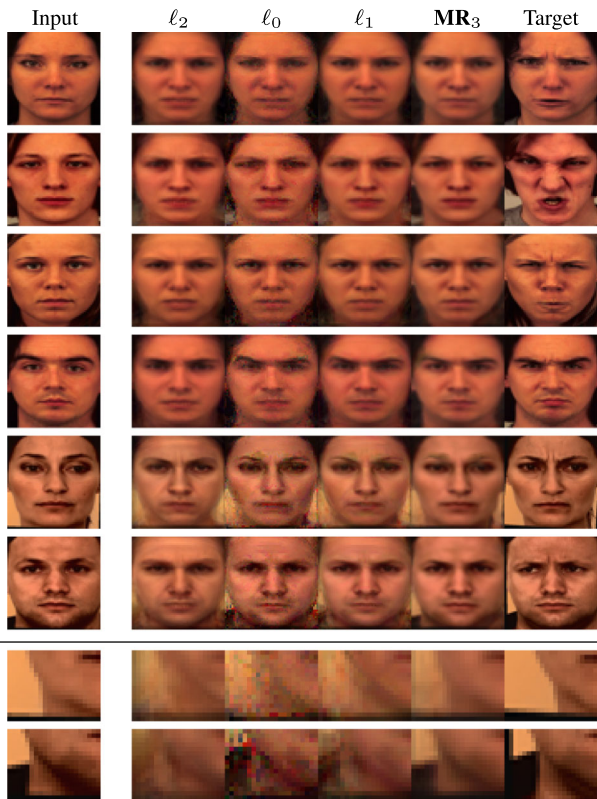


Fig. 9 Comparison of ℓ_2 -regression, ℓ_0 -regression, ℓ_1 -regression and Masked Regression (MR) results for the neutral to angry mapping on RGB images. Only local receptive fields were able to transform the expression while preserving identity the most and also retaining facial details the most. Last two rows are zoomed-in views of the bottom-left corners corresponding to rows 5 and 6 respectively. MR preserves background most successfully (Color figure online)

Table 3 Comparison of MR with ℓ_2 , ℓ_1 and ℓ_0 -regression on RGB images of size 56×56 in terms of mean-squared-error ($\times 10^2$)

In	Out	ℓ_2	ℓ_0	ℓ_1	MR ₃	MR ₅	MR ₇	MR ₉
Neu	Afr	1.183	1.209	1.081	1.027	1.026	1.046	1.073
Neu	Ang	1.088	1.114	0.954	0.909	0.887	0.898	0.919
Neu	Dis	1.067	1.136	0.996	0.914	0.898	0.916	0.939
Neu	Hap	0.962	0.940	0.836	0.792	0.789	0.803	0.818
Neu	Sad	0.977	0.985	0.835	0.760	0.768	0.783	0.794
Neu	Sur	1.069	1.124	0.997	1.034	1.007	1.007	1.025
Afr	Neu	1.108	1.114	1.002	0.875	0.886	0.933	0.974
Ang	Neu	0.964	1.013	0.864	0.852	0.835	0.858	0.882
Dis	Neu	1.010	0.970	0.882	0.862	0.843	0.864	0.884
Hap	Neu	0.868	0.833	0.748	0.748	0.738	0.759	0.782
Sad	Neu	0.913	0.921	0.784	0.760	0.754	0.772	0.790
Sur	Neu	1.039	1.048	0.911	0.949	0.917	0.911	0.928
Mean MSE		1.021	1.034	0.907	0.874	0.862	0.879	0.901

MR_r refers to masked regression with receptive fields of size $r \times r$ pixels
 Bold represents the minimum mean-squared-error among the different methods compared and which method is better

Table 4 Comparison of training times in seconds averaged over 12 different expression mappings

MR	ℓ_1	ℓ_0	ℓ_2
0.010	16.782	0.237	0.115

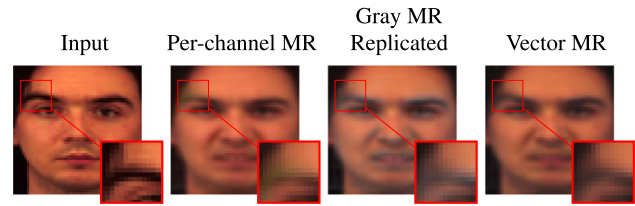


Fig. 10 Options for performing MR on color images. Column 2 color leakage due to learning separate transformations for each color channel. The eyebrow has developed a greenish tinge. Column 3: this can be avoided by using weights learned from a gray-scale mapping and replicating them on each color channel. However, it leads to some loss of colorfulness. Column 4: Best results are achieved by learning a single mapping between color vectors (Color figure online)

Sparsity comparison In addition to better performance and faster training, the ratio of the number of non-zero weights learned via the closest competitor (ℓ_1 -regression) to those learned via MR was 1.94 after averaging over 12 expression transformations over RGB images. In other words, masked regression was almost twice as sparse as ℓ_1 -regression.

5.3 Experiments on Non-Frontal Faces

We learned a neutral to happy mapping for non-frontal faces via the proposed MR technique. Results on a few test images are shown in Fig. 11 for 45° and 90° poses from the KDEF dataset. Training was performed on 56 image pairs while validation and test sets contained 7 image pairs each. It can

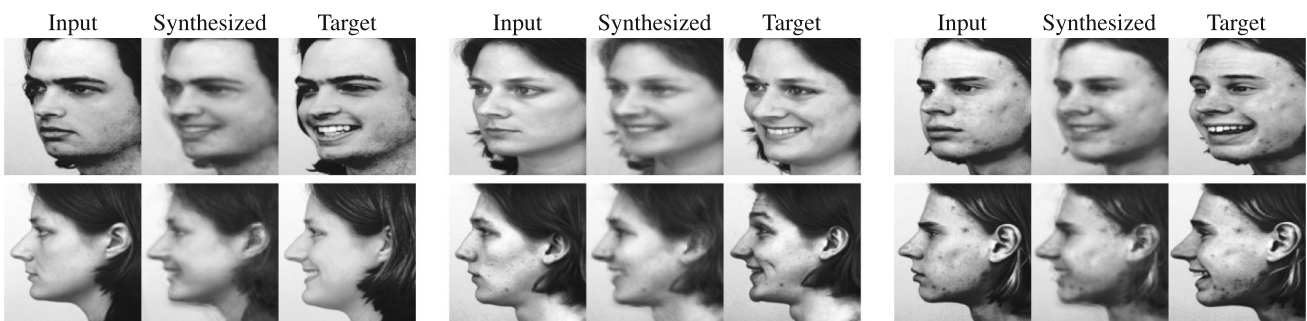


Fig. 11 Synthesis of neutral to happy expressions on non-frontal faces learned via MR. Top: 45° and Bottom: 90° rotation

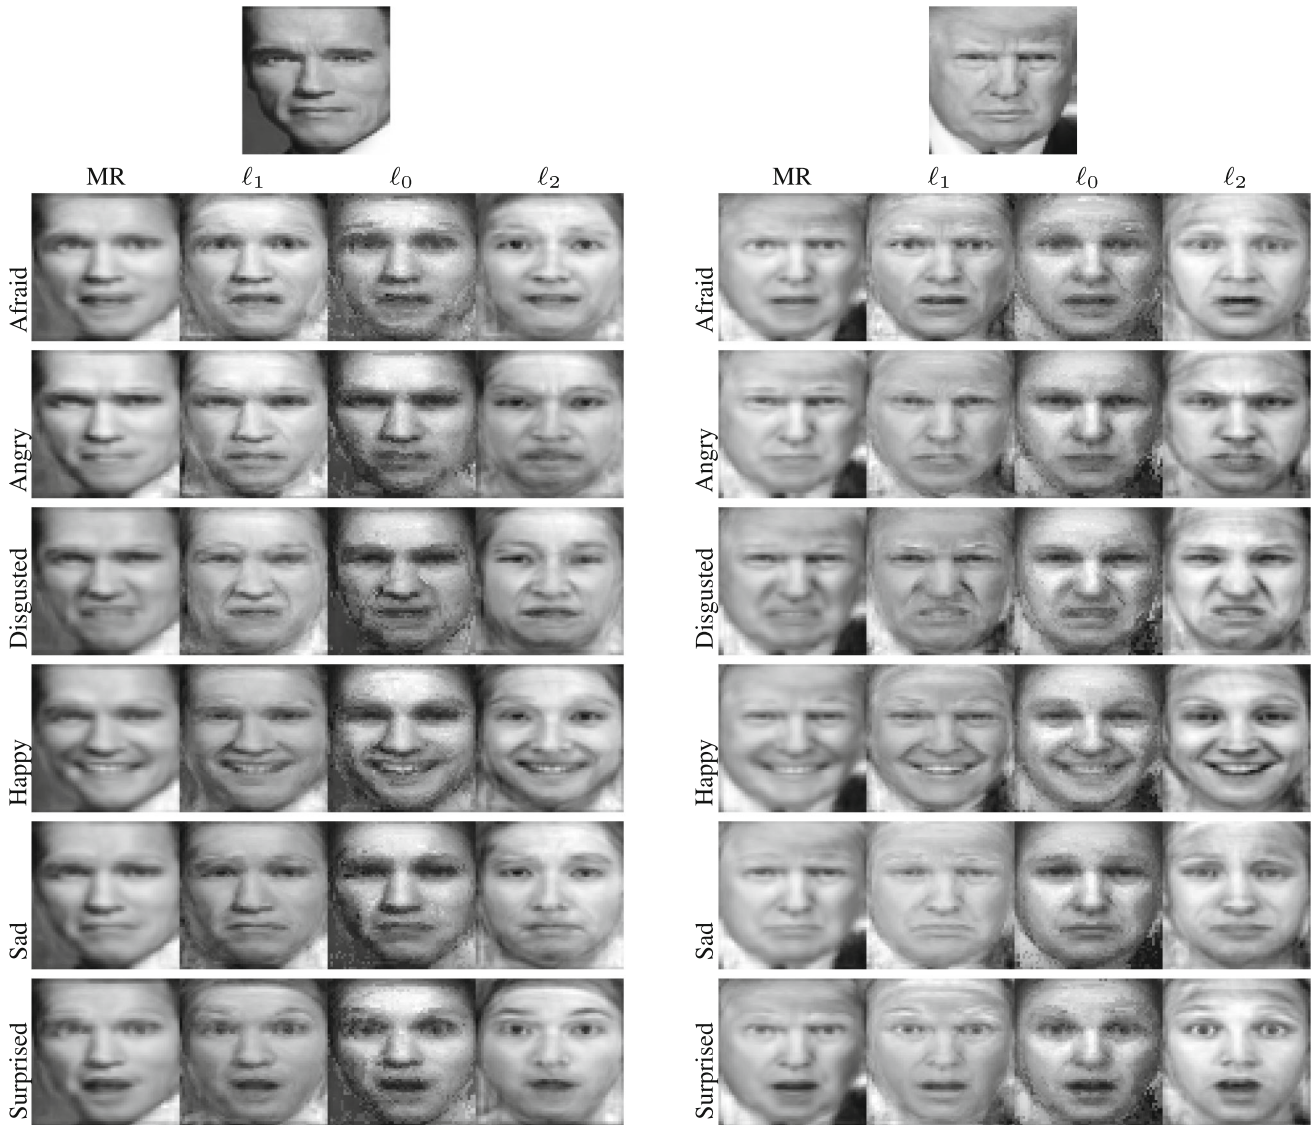


Fig. 12 Comparison of different regression methods on out-of-dataset face images downloaded from the Internet. The proposed masked regression (MR) generalized better than the compared methods. Despite

being trained on frontal faces only, MR did not enforce a frontal bias over inputs that were not entirely frontal faces, while competing methods introduced a frontal bias



Fig. 13 MR successfully generalized over pencil sketches. Left: a pencil sketch containing occlusion of the face, Right: an atypical sketch drawn by appropriate placements of English words. Compared meth-

ods demonstrated significant bias towards the training data by changing the pose, identity or facial details of the input faces. In contrast, MR was able to handle occlusion because it learns localized expression mappings

be seen that MR learns to change only the relevant portions of the input. Very small details (such as long hair visible near the mouth profile in 45° poses) are left unaffected as long as they have no role to play in the expression mapping.

5.4 Generalization Over Out-of-Dataset Images

Since masked regression uses so few parameters, it should be expected to generalize better than competing approaches. To check this, some specific and some arbitrary images were downloaded from the Internet. The intensity distributions of these images were significantly different from the datasets used for training, validation and testing.

5.4.1 Photographs

Figure 12 demonstrates that masked regression generalizes well over photographs taken in unconstrained settings of persons not belonging to any of the training datasets. The closest competing technique in this instance was once again ℓ_1 -regression which was sometimes able to produce identity preserving expression mappings but generally produced hallucination artifacts. It can also be noted for test faces that are not entirely frontal, MR does not enforce a strong frontal prior on the generated expression. The same cannot be said about competing methods that introduce a frontal bias learned from training data consisting of only frontal faces.

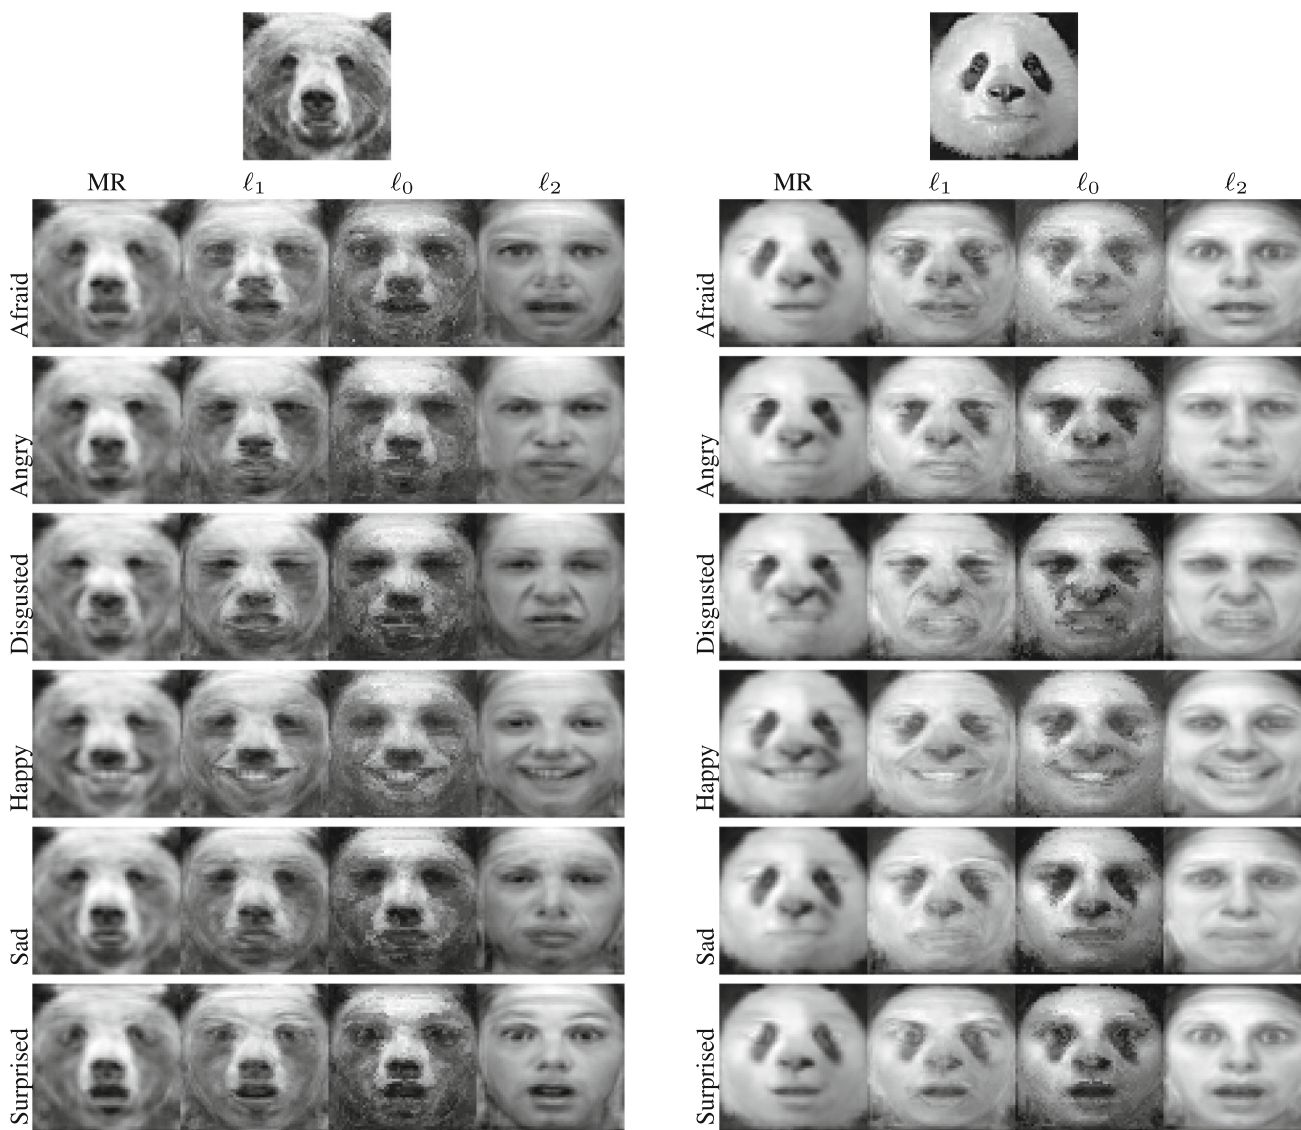


Fig. 14 Synthesized expressions for animal faces using the proposed algorithm. Since training was performed entirely on real human faces, these results demonstrate the strength of masked regression in learning essential attributes of expressions and generalizing them to non-human faces as well

5.4.2 Pencil Sketches and Animal Faces

Figure 13 shows the results of different regression methods on pencil sketches. Masked regression successfully generalized over pencil sketches containing occlusion of the face and an atypical sketch drawn by appropriate placements of English words. Competing methods demonstrated significant bias towards the training data by changing the pose, identity or facial details of the input face. In contrast, MR was able to handle occlusion since it focuses on learning localized expression mappings instead of global mappings.

Figure 14 shows the results of generating expressions for animal faces using the proposed algorithm. Since training was performed entirely on real human faces, these results

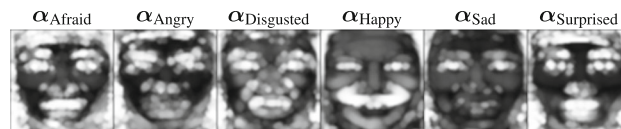


Fig. 15 Visualization of the α maps showing importance of different facial regions in generating 6 different expressions. The α maps are derived automatically as explained in Sect. 6

demonstrate the strength of masked regression in learning essential attributes of happy expressions and generalizing them to non-human faces as well.



Fig. 16 Visualization of the blur refinement algorithm as explained in Sect. 6. Image details from input x and expression details from MR output y are used to yield a refined expression image. The refined results show better recovery of facial hair, illumination effects and subtle facial features. An overall improvement in identity preservation can also be observed. The input image in the last row is made by combining different letter strokes. In the refined result, many of the letter strokes are also recovered (zoom in for better view) (Color figure online)

6 Blur Refinement Algorithm

In Fig. 4, a comparison of different regression techniques reveals blurriness in the synthesized expression images. In case of MR, this is due to the fact that for weights learned by minimizing sum-squared-error, predictions at test time are conditional means of the target variable (Bishop 2006, p. 46). Blurring can be reduced by determining the role α_{ij} of each output pixel in generating an expression. If a pixel has no role in expression generation, then its output value can be replaced by the corresponding value in the input image. This refinement of results can be written as a linear combination of input and output images. That is,

$$y' = (1 - \alpha) \circ x + \alpha \circ y \tag{17}$$

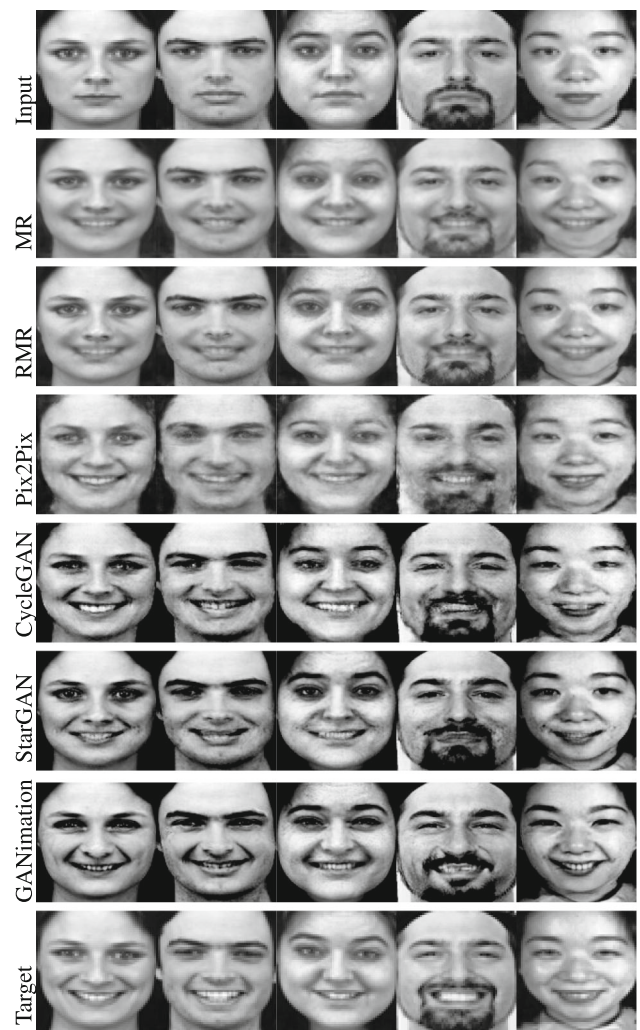


Fig. 17 Neutral to happy mappings synthesized by the proposed MR and RMR, Pix2Pix, CycleGAN, StarGAN and GANimation. Results are shown for unseen test images belonging to the same datasets that were used for training. Results produced by GANs were sharp but occasionally contained some artifacts. MR results were a bit smooth while RMR was able to produce convincing expressions with more facial details

where x , y and y' are the the input, output and refined images respectively and the α map contains per-pixel importances used for blending the input and output. We refer to refinement of MR results via Eq. (17) as Refined Masked Regression (RMR). We compute the importance image α as follows. First, we compute the ℓ_1 -norm of the receptive field (including bias) of each output pixel to obtain an image s of absolute receptive field sums. Let μ and σ denote the mean and standard deviation of image s . We standardize the sums in s and compute their absolute values as $z = |\frac{s-\mu}{\sigma}|$. These z values indicate how different a receptive field is from the average receptive field in terms of standard deviation. Then we perform morphological dilation with a disk shaped structuring element and rescale the result between 0 and 1. The dilation expands the influence of atypical receptive fields to surround-

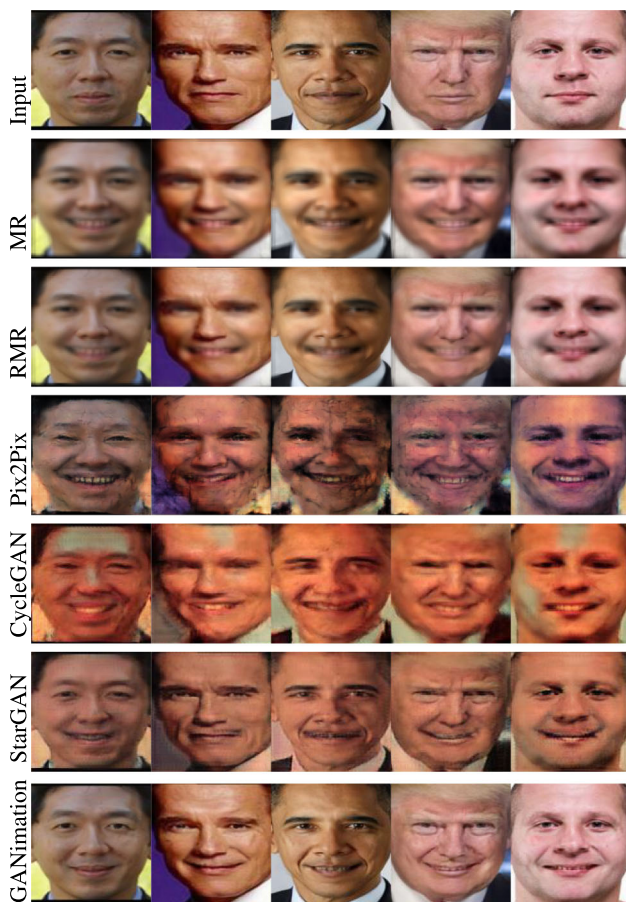


Fig. 18 Results of neutral to happy mappings on out-of-dataset human face photographs downloaded from the Internet. GANs fail to generalize well when test and training distributions are significantly different. In contrast, expressions synthesized by MR and RMR were satisfactory. Among the compared GANs, GANimation produced better results (Color figure online)

ing pixels. Then we pass the result through a smoothed-out step-function so that pixels with values greater than a threshold are moved towards 1 and the rest are moved towards 0. The smooth step-function that we use in our experiments is the logistic sigmoid function $(1 + \exp(-k(z - \tau)))^{-1}$ with $k = 10$ and threshold $\tau = 0.2$. After scaling the result between 0 and 1 again, we convolve with a Gaussian filter to obtain a smooth α map. All parameters related to dilation and smoothing are set adaptively with respect to image size.

This procedure of computing the α -map will make the synthesized output more important for pixels with receptive fields that are different from the average receptive field in terms of ℓ_1 -norm. Figure 15 shows the α maps corresponding to 6 expressions. It can be seen that eyes have a dominant role in all expressions. The mouth and cheeks have an important role in generating happy expressions. The forehead is important for afraid, angry and surprised expressions.

In the refined image, the input image contributes more in regions that do not play a major role in expression genera-

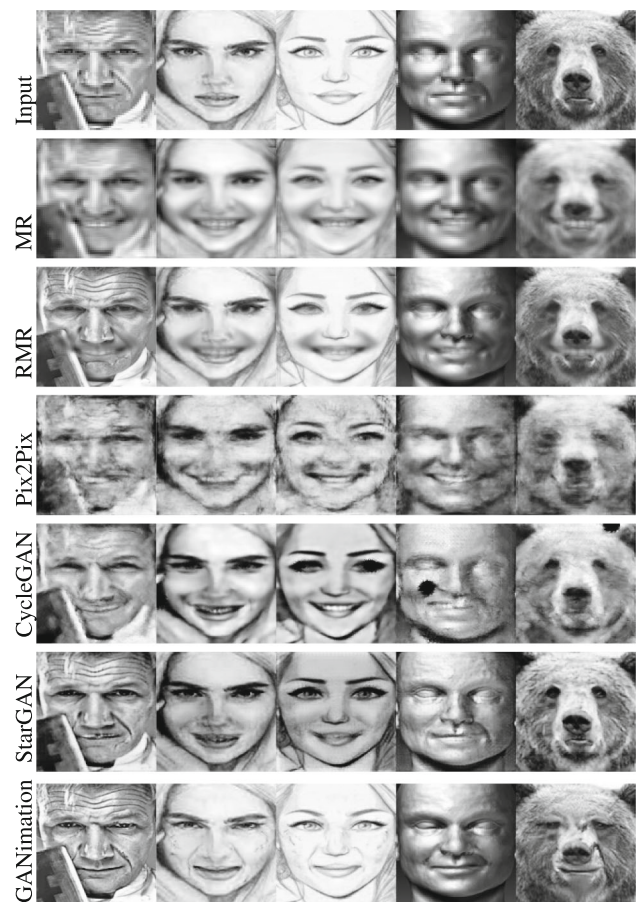


Fig. 19 Results of neutral to happy mappings. In most cases, GANs trained on real human photographs failed to generalize well. In contrast, MR and RMR also trained on real human photographs generated quite satisfactory happy expressions. First three columns are pencil sketches, last column is an animal face. CycleGAN was able to produce good results in some sketches while Pix2Pix and StarGAN showed more degraded performance. GANimation results depend heavily on reliable extraction of action units from a target face. The fourth column shows a 2D projection of a computer generated 3D model for which only MR and RMR were able to induce a convincing and artefact-free happy expression. The GANs were not able to induce expression in the animal face shown in the last column

tion. In contrast, in regions with a stronger role in expression generation the output of MR contributes more. This best-of-both-worlds solution adaptively copies sharp face details from the input and expression details from the output as shown in Fig. 16. In the rest of the paper, we refer to blur refined MR results as RMR.

7 Comparison with Generative Adversarial Networks

Recently, Generative Adversarial Networks (GANs) have induced tremendous interest in image-to-image translation tasks. We compare our results with four state-of-the-art

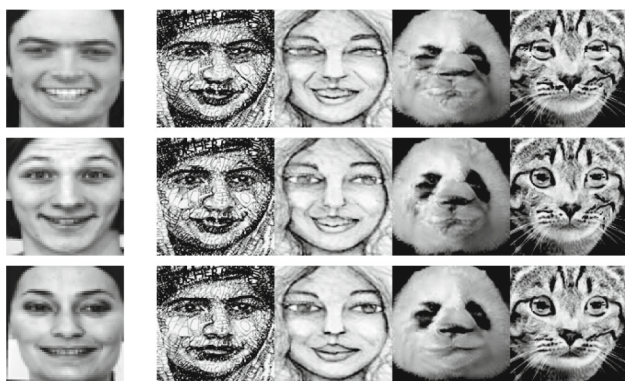


Fig. 20 Using three different target faces (left), GANimation failed to synthesize a happy expression over the two pencil sketches and the two animal faces. The eyes of the cat were transformed into human-like eyes (compare with sixth column of Fig. 1)

GANs, including Pix2Pix (Isola et al. 2016), CycleGAN (Zhu et al. 2017), StarGAN (Choi et al. 2018) and GANimation (Pumarola et al. 2019). We trained each of the first three GANs on the same dataset as used by MR and other algorithms as discussed in Sect. 5. We trained Pix2Pix for 100 epochs (in 5 h) on the same machine as used for other experiments. The CycleGAN was trained for 100 epochs in 48 h and the StarGAN was trained for 1000 epochs in 120 h. As reported in Table 4, training times for MR were less than a second. We used a pre-trained GANimation model that was trained for 30 epochs on the EmotionNet dataset (Benitez-Quiroz et al. 2016) which is much larger than our training set. Figure 17 demonstrates that GANs may generate quite good results as long as the testing images come from a distribution similar to the training images. However, for input images with features uncommon in the training set, such as facial hair in row number 4, the proposed MR and RMR methods were successful in inserting a reasonable looking smile. In addition, MR and RMR seem to better preserve the outer profile of faces.

In contrast, GANs produce sharper images, though sometimes, the outer profile is not well preserved (last row). For MR hidden details such as teeth are learned as the bias while GANs generate teeth as part of the samples from the learned distribution. In some cases, the generated teeth are quite good, while in other cases the teeth may degenerate and get mixed up with lips and other facial features. RMR retains expression details of MR while presenting better facial details similar to GANs.

Performance on out-of-dataset images

The performance of GANs and MRs is compared on out-of-dataset images downloaded from the Internet as discussed in Sect. 5.4.

We observe that in some cases, for testing images coming from different distributions, GANs were not able to generate convincing results as shown in Fig. 18. In contrast,

generalization of MR and RMR on out-of-dataset human photographs is better.

We further compare the generalization of GANs and MR algorithms on pencil-sketches of human faces in Fig. 19. Both GANs and MR algorithms were trained on the same real human face photographs as described in Sect. 5. Once again we observe that MR algorithms were able to produce better smiles. The gray color distribution of input sketches is also better preserved by the MR algorithms compared to GANs. Among the four compared GANs, CycleGAN produced better smiles on sketch images.

The performance of GANs and MR algorithms is also compared by generating happy expressions in animal faces. While GANs and MR algorithms were trained on the same real human face photographs, GANs were not able to synthesize a happy expression on any animal as demonstrated in Figs. 1 and 19. In contrast, MR and RMR were able to synthesize quite convincing happy expressions in animal faces. These experiments reveal the generalization strength of MR algorithms on images coming from distributions that are significantly different from the distribution of training datasets. Since GANimation results depend heavily on reliable extraction of action units from target faces, we used three different target faces in order to perform a fair comparison. Figure 20 shows that even using multiple targets, GANimation could not generalize well for pencil sketches and animal faces. It also produced human-like artefacts in animal faces. For example, the eyes of the cat were transformed into human-like eyes. In contrast, our proposed method preserved the cat's original features (see third row of Fig. 1).

Figure 21 compares the proposed method with the expression transfer results of GANimation (Pumarola et al. 2019). Input images were taken from their paper. The proposed method compared favorably against GANimation in terms of expression synthesis but GANimation results are sharper, irrespective of whether the expression was adequately transferred or not.

To quantitatively validate the out-of-dataset generalization of the proposed method, we used the EmoPy¹ expression recognition classifier pre-trained on the CK+ (Lucey et al. 2010) and FER+ (Barsoum et al. 2016) datasets to find the expression recognition accuracy for images synthesized by different methods. Table 5 shows the drop in expression recognition accuracy when test set images are replaced by out-of-dataset images. GAN based approaches suffered a larger drop in performance when tested on out-of-dataset images.

¹ <https://github.com/thoughtworksarts/EmoPy>.



Fig. 21 Direct comparison with expression transfer work in GANimation (Pumarola et al. 2019). Input images were taken from their paper. The proposed method compared favorably against GANimation in terms

of expression synthesis but GANimation results are sharper, irrespective of whether the expression was transferred or not (Color figure online)

Table 5 Drop in expression recognition accuracy (in percentage points) when changing from test set images to out-of-dataset images

Pix2Pix	CycleGAN	StarGAN	GANimation	MR
35.72	16.39	21.43	20.74	12.39

8 Conclusion

In this work masked regression has been introduced for facial expression synthesis using local receptive fields. Masked regression corresponds to a constrained version of ridge

regression. An efficient closed form solution for obtaining the global minimum for this problem is proposed. Despite being simple, the proposed algorithm has shown excellent learning ability on very small datasets. Compared to the existing learning based solutions, the proposed method is easier to implement and faster to train and has better generalization despite using small training datasets. The number of parameters in the learned model is also significantly smaller than competing methods. These properties are quite useful for learning high-dimensional to high-dimensional mappings as required for facial expression synthesis. Experiments performed on three publicly available datasets have shown the

superiority of the proposed method over approaches based on regression, sparse regression, kernelized regression and basis learning for both grayscale as well as color images.

Receptive fields learned via masked regression have a very intuitive interpretation which is further exploited to refine the output images.

Beyond the basic Masked Regression (MR) algorithm, an advanced Refined MR (RMR) algorithm is also proposed to reduce the blurring effects. Evaluations are also performed on out-of-dataset human photographs, pencil sketches, and animal faces. Results demonstrate that MR and RMR successfully synthesize the required expressions despite significant variations in the distribution of the test images compared to the training datasets. Comparisons are also performed on four state-of-the-art GANs including Pix2Pix, CycleGAN, StarGAN and GANimation. These GANs are able to generate photo-realistic expressions as long as testing and training distributions are similar. For the cases of out-of-dataset human photographs, pencil sketches and animal faces, these GANs exhibited degraded performance. In contrast, the proposed algorithm was able to generate quite satisfactory expressions in these cases as well. Therefore, the proposed algorithms generalize well compared to the current state-of-the-art facial expression synthesis methods.

As a future research direction, we suggest integration of the proposed MR and RMR algorithms within current-state-of-the-art GANs such as CycleGAN and StarGAN so that the resulting algorithm generalizes well on the out-of-dataset images and at the same time should be able to synthesize photo-realistic images. In addition, redundancy among different facial expressions can be exploited by learning a single weight matrix for all expressions. This is exploited by both StarGAN and GANimation to increase their training set from just source and target expressions to all available expressions. The proposed MR method can be extended in a similar fashion. Another future research direction is to explore generation of expressions with varying intensity levels. Expression intensity may be handled by learning discrete expression mappings corresponding to targets with different intensities. A continuous expression intensity map may be obtained by interpolating between discrete intensity levels.

References

- Barsoum, E., Zhang, C., Ferrer, C.C., & Zhang, Z. (2016) Training deep networks for facial expression recognition with crowd-sourced label distribution. In *Proceedings of the 18th ACM international conference on multimodal interaction, ACM* (pp 279–283).
- Belhumeur, P. N., Hespanha, J. P., & Kriegman, D. J. (1997). Eigenfaces vs Fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 711–720.
- Bermano, A. H., Bradley, D., Beeler, T., Zund, F., Nowrouzezahrai, D., Baran, I., et al. (2014). Facial performance enhancement using dynamic shape space analysis. *ACM Transactions on Graphics*, 33(2), 13:1–13:12.
- Bishop, C. M. (2006). *Pattern recognition and machine learning (information science and statistics)*. Berlin: Springer.
- Blanz, V., Vetter, T., et al. (1999). A morphable model for the synthesis of 3d faces. *SIGGRAPH*, 99, 187–194.
- Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., & Choo, J. (2018) StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8789–8797).
- Coates, A., & Ng, A.Y. (2011) Selecting receptive fields in deep networks. In *Advances in neural information processing systems* (pp. 2528–2536).
- Cootes, T. F., Edwards, G. J., Taylor, C. J., et al. (2001). Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6), 681–685.
- Costigan, T., Prasad, M., & McDonnell, R. (2014) Facial retargeting using neural networks. In *Proceedings of the seventh international conference on motion in games, ACM* (pp. 31–38).
- De La Hunty, M., Asthana, A., & Goecke, R. (2010) Linear facial expression transfer with active appearance models. In *2010 20th international conference on pattern recognition, IEEE* (pp. 3789–3792).
- Deng, Z., & Noh, J. (2008) Computer facial animation: A survey. In *Data driven 3D facial animation* (pp. 1–28). Berlin: Springer.
- Ekman, P., Friesen, W. V., & Ellsworth, P. (2013). *Emotion in the human face: Guidelines for research and an integration of findings*. Amsterdam: Elsevier.
- Elaiwat, S., Bennamoun, M., & Boussaid, F. (2016). A spatio-temporal RBMbased model for facial expression recognition. *Pattern Recognition*, 49, 152–161.
- Fabian Benitez-Quiroz, C., Srinivasan, R., & Martinez, A.M. (2016) Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5562–5570).
- Georgakis, C., Panagakis, Y., & Pantic, M. (2016). Discriminant incoherent component analysis. *IEEE Transactions on Image Processing*, 25(5), 2021–2034.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014) Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672–2680).
- Havaldar, P. (2006) Sony pictures imageworks. In *ACM SIGGRAPH 2006 Courses* (p. 5). New York: ACM.
- Huang, D., & De la Torre, F. (2010) Bilinear kernel reduced rank regression for facial expression synthesis. In *European conference on computer vision* (pp. 364–377). Berlin: Springer.
- Isola, P., Zhu, J.Y., Zhou, T., & Efros, A.A. (2017) Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1125–1134).
- Jain, A. K., & Li, S. Z. (2011). *Handbook of face recognition*. Berlin: Springer.
- Jampour, M., Mauthner, T., & Bischof, H. (2015) Multi-view facial expressions recognition using local linear regression of sparse codes. In *Proceedings of the 20th computer vision winter workshop*.
- Kim, T., Cha, M., Kim, H., Lee, J.K., & Kim, J. (2017) Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th international conference on machine learning* (Vol. 70, pp. 1857–1865).
- Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., & Matas, J. (2018) Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8183–8192).

- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., & Wang, Z. et al. (2017). Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4681–4690).
- Lee, C.S., & Elgammal, A. (2006). Nonlinear shape and appearance models for facial expression analysis and synthesis. In *18th international conference on pattern recognition* (Vol. 1, pp. 497–502). IEEE.
- Lee, H. S., & Kim, D. (2008). Tensor-based aam with continuous variation estimation: Application to variation-robust face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(6), 1102–1116.
- Lin, J. R., & Ic, Lin. (2011). Multi-layered expression synthesis. *Journal of Information Science and Engineering*, 27(1), 337–351.
- Liu, M. Y., & Tuzel, O. (2016). Coupled generative adversarial networks. In *Advances in neural information processing systems* (pp. 469–477).
- Liu, M. Y., Breuel, T., & Kautz, J. (2017). Unsupervised image-to-image translation networks. In *Advances in neural information processing systems* (pp. 700–708).
- Liu, S., Huang, D. Y., Lin, W., Dong, M., Li, H., & Ong, E. P. (2014). Emotional facial expression transfer based on temporal restricted Boltzmann machines. In *2014 Asia-Pacific Signal and information processing association annual summit and conference (APSIPA)*, (pp. 1–7).
- Liu, Z., Shan, Y., & Zhang, Z. (2001). Expressive expression mapping with ratio images. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 271–276). ACM.
- Lucey, P., Cohn, J.F., Kanade, T., Saragih, J., Ambadar, Z., & Matthews, I. (2010). The extended Cohn-Kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE computer society conference on computer vision and pattern recognition-workshops* (pp. 94–101). IEEE.
- Lundqvist, D., Flykt, A., & Öhman, A. (1998). *The karolinska directed emotional faces-KDEF, CD ROM*. Stockholm: Department of Clinical Neuroscience, Psychology section, Karolinska Institutet.
- Lyons, M., Akamatsu, S., Kamachi, M., & Gyoba, J. (1998). Coding facial expressions with Gabor wavelets. In *Proceedings third IEEE international conference on automatic face and gesture recognition* (pp. 200–205). IEEE.
- Mirza, M., Osindero, S. (2014). Conditional generative adversarial nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784)
- Nhan Duong, C., Luu, K., Gia Quach, K., & Bui, T.D. (2016). Longitudinal face modeling via temporal deep restricted boltzmann machines. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5772–5780).
- Pantic, M., & Rothkrantz, L. J. M. (2000). Automatic analysis of facial expressions: The state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12), 1424–1445.
- Patel, N. M., & Zaveri, M. (2010). Parametric facial expression synthesis and animation. *International Journal of Computer Applications*, 3, 34–40.
- Pati, Y.C., Rezaifar, R., & Krishnaprasad, P.S. (1993). Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar conference on signals, systems and computers* (Vol. 1, pp. 40–44).
- Pighin, F., & Lewis, J. (2006). Performance-driven facial animation. In Pighin, F., Hecker, J., Lischinski, D., Szeliski, R., & Salesin, D.H. (eds). *Synthesizing realistic facial expressions from photographs, ACM SIGGRAPH 2006 Courses*, ACM (p. 19).
- Pumarola, A., Agudo, A., Martinez, A., Sanfeliu, A., & Moreno-Noguer, F. (2019). GANimation: One-shot anatomically consistent facial animation. *International Journal of Computer Vision (IJCV)*. <https://doi.org/10.1007/s11263-019-01210-3>.
- Rizzo, A. A., Neumann, U., Enciso, R., Fidaleo, D., & Noh, J. (2004). Performance-driven facial animation: Basic research on human judgments of emotional state in facial avatars. *CyberPsychology & Behavior*, 4(4), 471–487.
- Saragih, J.M., Lucey, S., & Cohn, J.F. (2011). Real-time avatar animation from a single image. In *2011 IEEE international conference on automatic face & gesture recognition and workshops (FG 2011)* (pp. 117–124). IEEE.
- Savran, A., Alyüz, N., Dibeklioglu, H., Çeliktutan, O., Gökberk, B., Sankur, B., & Akarun, L. (2008). Bosphorus database for 3D face analysis. In *European workshop on biometrics and identity management* (pp. 47–56). Springer.
- Shen, W., & Liu, R. (2017). Learning residual images for face attribute manipulation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4030–4038).
- Susskind, J.M., Anderson, A.K., Hinton, G.E., & Movellan, J.R. (2008). Generating facial expressions with deep belief nets. In: Or, J. (ed.) *Affective computing* (chapter 10, pp. 421–440). InTech.
- Suwajanakorn, S., Seitz, S.M., & Kemelmacher-Shlizerman, I. (2015). What makes tom hanks look like tom hanks. In *Proceedings of the IEEE international conference on computer vision* (pp. 3952–3960).
- Tenenbaum, J. B., & Freeman, W. T. (2000). Separating style and content with bilinear models. *Neural Computation*, 12(6), 1247–1283.
- Thies, J., Zollhofer, M., Stamminger, M., Theobalt, C., & Nießner, M. (2016). Face2Face: Real-time face capture and reenactment of RGB videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2387–2395).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)*, 58, 267–288.
- Tropp, J. A., & Gilbert, A. C. (2007). Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12), 4655–4666.
- Wang, H., et al (2003). Facial expression decomposition. In *Proceedings ninth IEEE international conference on computer vision* (pp. 958–965). IEEE.
- Wei, W., Tian, C., Maybank, S. J., & Zhang, Y. (2016). Facial expression transfer method based on frequency analysis. *Pattern Recognition*, 49, 115–128.
- Yi, Z., Zhang, H., Tan, P., & Gong, M. (2017). DualGAN: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision* (pp. 2849–2857).
- Zeiler, M.D., Taylor, G.W., Sigal, L., Matthews, I., & Fergus, R. (2011). Facial expression transfer with input-output temporal restricted boltzmann machines. In *Advances in neural information processing systems* (pp. 1629–1637).
- Zeng, Z., Pantic, M., Roisman, G. I., & Huang, T. S. (2009). A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(1), 39–58.
- Zhang, G., Kan, M., Shan, S., & Chen, X. (2018). Generative adversarial network with spatial attention for face attribute editing. In *Proceedings of the European conference on computer vision* (pp. 417–432).
- Zhang, Q., Liu, Z., Quo, G., Terzopoulos, D., & Shum, H. Y. (2006). Geometrydriven photorealistic facial expression synthesis. *IEEE Transactions on Visualization and Computer Graphics*, 12(1), 48–60.
- Zhang, Y., & Wei, W. (2012). A realistic dynamic facial expression transfer method. *Neurocomputing*, 89, 21–29.

- Zhu, J.Y., Krähenbühl, P., Shechtman, E., & Efros, A.A. (2016) Generative visual manipulation on the natural image manifold. In *European conference on computer vision* (pp. 597–613). Springer.
- Zhu, J.Y., Park, T., Isola, P., & Efros, A.A. (2017) Unpaired image-to-image translation using cycle-consistent adversarial networks.

In *Proceedings of the IEEE international conference on computer vision* (pp. 2223–2232).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.