

A Recurrent Encoder-Decoder Network for Sequential Face Alignment

Xi Peng^{1(✉)}, Rogerio S. Feris², Xiaoyu Wang³, and Dimitris N. Metaxas¹

¹ Rutgers University, Piscataway, USA
{xipeng.cs,dnm}@rutgers.edu

² IBM T. J. Watson Research Center, Yorktown Heights, USA
rsferis@us.ibm.com

³ Snapchat Research, Venice, CA, USA
xiaoyu.wang@snapchat.com

Abstract. We propose a novel recurrent encoder-decoder network model for real-time video-based face alignment. Our proposed model predicts 2D facial point maps regularized by a regression loss, while uniquely exploiting recurrent learning at both spatial and temporal dimensions. At the spatial level, we add a feedback loop connection between the combined output response map and the input, in order to enable iterative coarse-to-fine face alignment using a *single network model*. At the temporal level, we first decouple the features in the bottleneck of the network into *temporal-variant factors*, such as pose and expression, and *temporal-invariant factors*, such as identity information. Temporal recurrent learning is then applied to the decoupled temporal-variant features, yielding better generalization and significantly more accurate results at test time. We perform a comprehensive experimental analysis, showing the importance of each component of our proposed model, as well as superior results over the state-of-the-art in standard datasets.

Keywords: Recurrent learning · Encoder-decoder · Face alignment

1 Introduction

Face landmark detection plays a fundamental role in many computer vision tasks, such as face recognition, expression analysis, and 3D face modeling. In the past few years, many methods have been proposed to address this problem, with significant progress being made towards systems that work in real-world conditions (“in the wild”).

Regression-based approaches [6, 50] have achieved impressive results by cascading discriminative regression functions that directly map facial appearance to landmark coordinates. In this framework, deep convolutional neural networks have proven effective as a choice for feature extraction and non-linear regression modeling [21, 54, 55]. Although these methods can achieve very reliable results in

standard benchmark datasets, they still suffer from limited performance in challenging scenarios, e.g., involving large face pose variations and heavy occlusions.

A promising direction to address these challenges is to consider video-based face alignment (i.e., sequential face landmark detection) [39], leveraging temporal information as an additional constraint [47]. Despite the long history of research in rigid and non-rigid face tracking [5, 10, 32, 33], current efforts have mostly focused on face alignment in still images [37, 45, 54, 57]. In fact, most methods often perform video-based landmark detection by independently applying models trained on still images in each frame in a tracking-by-detection manner [48], with notable exceptions such as [1, 36], which explore incremental learning based on previous frames. How to effectively model long-term temporal constraints while handling large face pose variations and occlusions is an open research problem for video-based face alignment.

In this work, we address this problem by proposing a novel recurrent encoder-decoder deep neural network model (see Fig. 1). The encoding module projects image pixels into a low-dimensional feature space, whereas the decoding module maps features in this space to 2D facial point maps, which are further regularized by a regression loss. In order to handle large face pose variations, we introduce a feedback loop connection between the aggregated 2D facial point maps and the input. The intuition is similar to cascading multiple regression functions [50, 54] for iterative coarse-to-fine face alignment, but in our approach the iterations are modeled jointly with shared parameters, using a single network model.

For more effective temporal modeling, we first decouple the features in the bottleneck of the network into temporal-variant factors, such as pose and expression, and temporal-invariant factors, such as identity. More specifically, we split the features into two components, where one component is used to learn face recognition using identity labels, and recurrent temporal learning is applied to the other component, which encodes temporal-variant factors only. We show in our experiments that recurrent learning in both spatial and temporal dimensions is crucial to improve performance of sequential face landmark detection.

In summary, our work makes the following **contributions**:

- We propose a novel recurrent encoder-decoder network model for real-time sequential face landmark detection. To the best of our knowledge, this is the first time a recurrent model is investigated to perform video-based facial landmark detection.
- Our proposed *spatial recurrent learning* enables a novel iterative coarse-to-fine face alignment using a single network model. This is critical to handle large face pose changes and a more effective alternative than cascading multiple network models in terms of accuracy and memory footprint.
- Different from traditional methods, we apply *temporal recurrent learning* to temporal-variant features which are decoupled from temporal-invariant features in the bottleneck of the network, achieving better generalization and more accurate results.

- We provide a detailed experimental analysis of each component of our model, as well as insights about key contributing factors to achieve superior performance over the state-of-the-art. The project page is public available.¹

2 Related Work

Face alignment has been advanced in last decades. Remarkably, regression based methods [1, 2, 6, 17, 41, 45, 49, 50, 54, 57, 58] significantly boost the generalization performance of face landmark detection, compared to algorithms based on statistical models such as Active shape models [9, 29] and Active appearance models [12]. A regression-based approach directly regresses landmark locations where features extracted from face images serve as regressors. Landmark models are learned either in an independent manner, or in a joint fashion [6]. This paper performs landmark detection via both a classification model and a regression model. Different from most of the previous methods, this work deals with face alignment in a video. It jointly optimizes detection output by utilizing multiple observations from the same person.

Learning cascade-like regression models show superior performance on the face alignment task [41, 50, 54]. Supervised descent method [50] learns cascades of regression models based on SIFT feature. Sun *et al.* [41] proposed to use three levels of neural networks to predict landmark locations. Zhang *et al.* [54] studied the problem via cascades of stacked auto-encoders which gradually refine the landmark position with higher resolution inputs. Compared to these efforts which explicitly define cascade structures, our method learns a spatial recurrent model which implicitly incorporates the cascade structure with shared parameters. It is also more “end-to-end” compared to previous works that handcraftly divide the learning process into multiple stages.

Recurrent neural networks (RNNs) are widely employed in the literature of speech recognition [28] and natural language processing [27]. They are also recently used in computer vision. For example, in the task of image captioning [18] and video captioning [52], RNNs are employed for text generation. Veeriah *et al.* [46] use RNNs to learn complex time-series representations via high-order derivatives of states for action recognition. Benefiting from the deep architecture, RNNs are naturally good alternatives to Conditional Random Fields (CRFs) [56] which are popular in image segmentation.

Encoder and decoder networks are well studied in machine translation [7] where the encoder learns the intermediate representation and the decoder generates the translation from the representation. It is also investigated in speech recognition [26] and computer vision [3, 14]. Yang *et al.* [51] proposed to decouple identity units and pose units in the bottleneck of the network for 3D view synthesis. However, how to fully utilize the decoupled units for correspondence regularization [25] is still unexplored. In this work, we employ the encoder to learn a joint representation for identity, pose, expression as well as landmarks.

¹ <https://sites.google.com/site/xipengshomepage/project/face-alignment>.

The decoder translates the representation to landmark heatmaps. Our spatial recurrent model loops the whole encoder-decoder framework.

3 Recurrent Encoder-Decoder Network

In this section, we first give an overview of our approach. Then we describe the novelty of our work in detail: spatial and temporal recurrent learning, supervised identity disentangling, and constrained shape prediction.

3.1 Method Overview

Our task is to locate L landmarks in sequential images using an end-to-end deep neural network. Figure 1 shows the overview of our approach. We consider f_* as potential nonlinear and multi-layered functions. The input of the network are the image $\mathbf{x} \in \mathbb{R}^{w \times h \times 3}$ and the landmark label map $\mathbf{z} \in \mathbb{R}^{w \times h \times 1}$. Each pixel in \mathbf{z} is a discrete label $\{0, \dots, L\}$ that marks the presence of the corresponding landmark, where 0 denotes a non-landmark area.

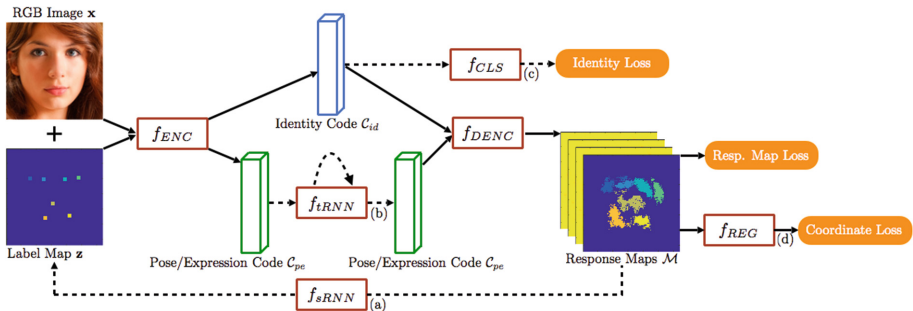


Fig. 1. Overview of the recurrent encoder-decoder network: (a) spatial recurrent learning (Sect. 3.2); (b) temporal recurrent learning (Sect. 3.3); (c) supervised identity disentangling (Sect. 3.4); and (d) constrained shape prediction (Sect. 3.5). f_{ENC} , f_{DENC} , f_{sRNN} , f_{tRNN} , f_{CLS} , f_{REG} are potentially nonlinear and multi-layered mappings.

The *encoder* (f_{ENC}) performs a sequence of convolution, pooling and batch normalization [15] to extract a representation code from inputs:

$$\mathcal{C} = f_{ENC}(\mathbf{x}, \mathbf{z}; \theta_{ENC}), \mathcal{C} \in \mathbb{R}^{w_c \times h_c \times d_c}, \quad (1)$$

where \mathcal{C} represents the encoded features. θ_{ENC} denotes encoder parameters. Symmetrically, the *decoder* (f_{DENC}) performs a sequence of unpooling, convolution and batch normalization to upsample the representation codes to a multi-channel response map:

$$\mathcal{M} = f_{DENC}(\mathcal{C}; \theta_{DENC}), \mathcal{M} \in \mathbb{R}^{w \times h \times (L+1)}, \quad (2)$$

where θ_{DENC} denotes the decoder parameters. The first channel of \mathcal{M} represents the background, while the rest L channels of \mathcal{M} present pixel-wise confidence of the corresponding landmarks. The $(L + 1)$ -channel response map is crucial to preserve the landmark unity, compared with a 2-channel setup (landmark *v.s.* non-landmark).

The encoder-decoder framework plays an important role in our task. **First**, it is convenient to perform *spatial recurrent learning* (f_{sRNN}) since \mathcal{M} has the same dimension (but different number of channels) as \mathbf{x} . The output of the decoder can be directly fed back into the encoder to provide pixel-wise spatial cues for the next recurrent step. **Second**, we can decouple \mathcal{C} in the bottleneck of the network into temporal-variant and -invariant factors. The former is further exploited in *temporal recurrent learning* (f_{tRNN}) for robust alignment, while the latter is used in *supervised identity disentangling* (f_{CLS}) to facilitate the network training. **Third**, \mathcal{M} can be further regularized in *constrained shape prediction* (f_{REG}) to directly output landmark coordinates. The details of each module are explained in following subsections.

3.2 Spatial Recurrent Learning

The purpose of spatial recurrent learning is to pinpoint landmark locations in a coarse-to-fine manner. Unlike existing approaches [41, 54] that employ multiple networks in cascade, we accomplish the coarse-to-fine search in a single network in which the parameters are jointly learned in successive recurrent steps.

Given an image \mathbf{x} and initial guess of the shape \mathbf{z}_0 , we refine the shape prediction iteratively $\{\mathbf{z}^1, \dots, \mathbf{z}^k\}$ by feeding back the previous prediction:

$$\mathbf{z}^k = f_{sRNN}(\mathcal{M}^{k-1}) = f_{sRNN}(f_{DENC}(f_{ENC}(\mathbf{x}, \mathbf{z}^{k-1}))), k = 1, \dots, K, \quad (3)$$

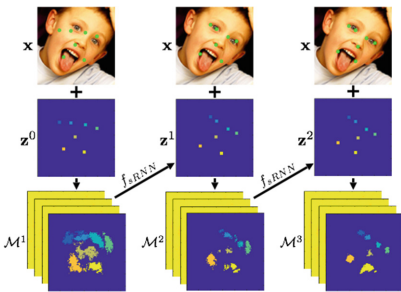


Fig. 2. An unrolled illustration of *spatial recurrent learning*. The response map is pretty coarse when the initial guess is far away from the ground truth if large pose and expression exist. It eventually gets refined in the successive recurrent steps.

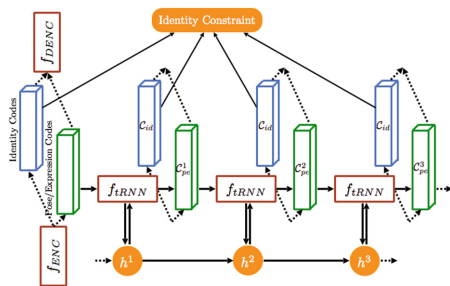


Fig. 3. An unrolled illustration of *temporal recurrent learning*. C_{id} encodes temporal-invariant factor which subjects to the same identity constraint. C_{pe} encodes temporal-variant factors which is further modeled in f_{tRNN} .

where we omit network parameters θ_{ENC} and θ_{DENC} for concise expression. The network parameters are learned by recurrently minimizing the classification loss between the annotation and the response map output by the encoder-decoder:

$$\operatorname{argmin}_{\theta_{ENC}, \theta_{DENC}} \sum_{k=1}^K \sum_{l=0}^L \ell(\mathcal{M}_l^*, f_{DENC}(f_{ENC}(\mathbf{x}, \mathbf{z}^k)))_l, \quad (4)$$

where k counts iterations and l counts landmarks. $\mathcal{M}_l^* \in \mathbb{R}^{w \times h \times 1}$ is the ground truth of the response map for the l -th landmark. As shown in Fig. 2, our recurrent model progressively improves the prediction accuracy when a face exhibits challenging pose or expression. The whole process is learned end-to-end during training.

3.3 Temporal Recurrent Learning

The recurrent learning is performed at both the spatial and temporal dimensions. Given T successive frames $\{\mathbf{x}^t; t = 1, \dots, T\}$, the encoder extracts a sequence of representation codes $\{\mathcal{C}^t; t = 1, \dots, T\}$. We can decouple \mathcal{C} as: identity code \mathcal{C}_{id} that is *temporal-invariant* since all frames are subject to the same identity constraint; and pose/expression code \mathcal{C}_{pe} that is *temporal-variant* since pose and expression changes over time [34]. We exploit the temporal consistence of \mathcal{C}_{pe} via the proposed temporal recurrent learning.

Figure 3 shows the unrolled illustration of the proposed temporal recurrent learning. More specifically, we aim to achieve a nonlinear mapping f_{tRNN} , which simultaneously tracks the latent state $\{h^t; t = 1, \dots, T\}$ and updates \mathcal{C}_{pe} at time t :

$$h^t = p(\mathcal{C}_{pe}^t, h^{t-1}; \theta_{tRNN}), \mathcal{C}_{pe}^{t'} = q(h^t; \theta_{tRNN}), t = 1, \dots, T \quad (5)$$

where $p(\cdot)$ and $q(\cdot)$ are functions of f_{tRNN} . $\mathcal{C}_{pe}^{t'}$ is the update of \mathcal{C}_{pe}^t . θ_{tRNN} corresponds to mapping parameters which are learned in the end-to-end task using the same classification loss as Eq. 4 but unrolled at the temporal dimension:

$$\operatorname{argmin}_{\theta_{ENC}, \theta_{DENC}, \theta_{tRNN}} \sum_{t=1}^T \sum_{l=0}^L \ell_{tRNN}(\mathcal{M}_l^{t*}, f_{DENC}(\mathcal{C}_{id}^t, \mathcal{C}_{pe}^t)_l), \quad (6)$$

where t counts time steps and l counts landmarks. Note that both spatial and temporal recurrent learning are performed to jointly learn θ_{ENC} , θ_{DENC} and θ_{tRNN} in the same task according to Eqs. 4 and 6.

The temporal recurrent learning memorize the motion patterns of pose and expression variations from offline training data. It can significantly improve the fitting accuracy and robustness when large variations and partial occlusions exist.

3.4 Supervised Identity Disentangling

There is no guarantee that temporal-invariant and -variant factors can be completely decoupled in the bottleneck by simply splitting the representation codes into two parts. More supervised information is required to achieve the decoupling. To address this issue, we propose to apply a face recognition task on the identity code, in addition to the temporal recurrent learning applied on pose/expression code.

The supervised identity disentangling is formulated as an N -way classification problem. N is the number of unique individuals present in the training sequences. In general, the classification network f_{CLS} associates the identity code \mathcal{C}_{id} with a vector indicating the score of each identity. Classification loss is used to learn the mapping parameters:

$$\operatorname{argmin}_{\theta_{CLS}} \sum_{m=1}^M \ell_{CLS}(\mathbf{e}^*, f_{CLS}(\mathcal{C}_{id}; \theta_{CLS})), \quad (7)$$

where m counts the number of training images in a mini batch. \mathbf{e}^* is the one-hot identity annotation vector with a 1 for the correct identity and all 0s for others.

It has been shown in [55] that learning the face alignment task together with correlated tasks, *e.g.* head pose, can improve the fitting performance. We have the similar observation when adding face recognition task to the alignment task. More specifically, we found that supervised identity disentangling can significantly improve the generalization as well as fitting accuracy at test time. In this case, the factors are better decoupled, which facilitates f_{tRNN} to better handle temporal variations.

3.5 Constrained Shape Prediction

The response map output by the encoder-decoder may have a few false high responses when distractions exist in the background. Although this issue is significantly alleviated by spatial recurrent learning, it still impairs the fitting accuracy in challenging conditions. Besides, the response map uses separate channels to depict each landmark. The spatial dependencies among landmarks are not well explored. To overcome these limitations, we append nonlinear mappings after the encoder-decoder to learn the shape constraint for shape prediction.

f_{REG} takes the response map as the input and outputs landmark coordinates $\mathbf{y} \in \mathbb{R}^{2L \times 1}$. Regression loss is used to learn the mapping parameters:

$$\operatorname{argmin}_{\theta_{REG}} \sum_{n=1}^N \ell_{REG}(\mathbf{y}^*, f_{REG}(\mathcal{M}; \theta_{REG})), \quad (8)$$

where \mathbf{y}^* is the ground truth of landmark coordinates. All coordinates are normalized by subtracting a mean shape calculated from training images. The summation accumulates loss within a mini batch to avoid gradient jiggling.

4 Network Architecture and Implementation Details

All modules are embedded in a unified framework that can be trained end-to-end. Next we provide more details about how we guarantee efficient training convergence and robust performance at test time.

4.1 f_{ENC} and f_{DENC}

Figure 4 illustrates the detailed configuration of the encoder-decoder. The encoder is designed based on a variant of the VGG-16 network [19, 40]. It has 13 convolutional layers with constant 3×3 filters which correspond to the first 13 convolutional layers in VGG-16. We can therefore initialize the training process from weights trained on large datasets for object classification. We remove all fully connected layers in favor of *fully convolutional networks* (FCNs) [24] and output two $4 \times 4 \times 256$ feature maps in the bottleneck. This strategy not only reduces the number of parameters from 117M to 14.8M [3], but also preserves spatial information in high-resolution feature maps instead of fully-connected feature vectors, which is crucial for our landmark localization task.

There are 5 max-pooling layers with 2×2 pooling windows and a constant stride of 2 in the encoder to halve the resolution of feature maps after each convolutional stage. Although max-pooling can help to achieve translation invariance, it inevitably results in a considerable loss of spatial information especially when several max-pooling layers are applied in succession. To solve this issue, we use a 2-bit code to record the index of the maximum activation selected in a 2×2 pooling window [53]. As illustrated in Fig. 4, the memorized index is then used in the corresponding unpooling layer to place each activation back to its original location. This strategy is particularly useful for the decoder to recover the input

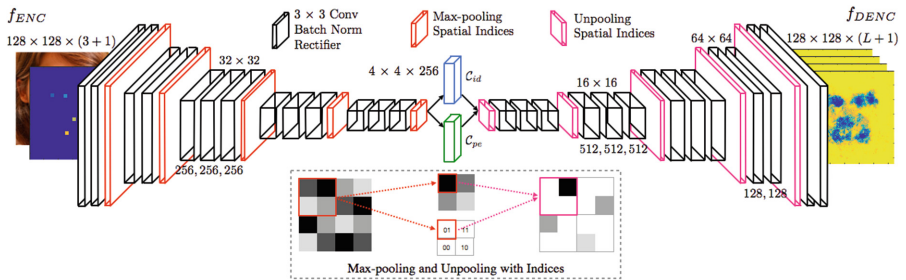


Fig. 4. Architecture of f_{ENC} and f_{DENC} . The input of the encoder is the concatenation of 3-channel image and 1-channel label map. The decoder is exactly symmetrical to the encoder except the output is a $(L+1)$ -channel response map. The representation code is split into C_{id} and C_{pe} in the bottleneck, where each one is a $4 \times 4 \times 256$ feature map. 3×3 kernels are used in all convolutional layers. 2×2 max-pooling or unpooling windows are applied in all pooling layers. The corresponding max-pooling and unpooling share pooling indices with a 2-bit switch for each 2×2 pooling window.

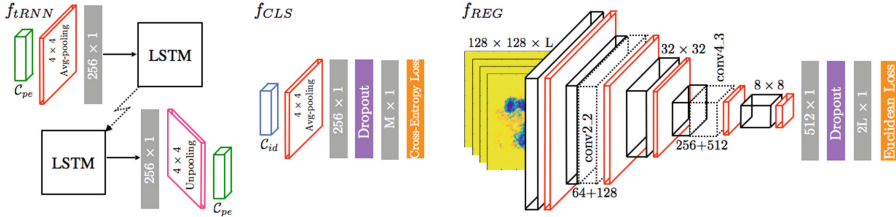


Fig. 5. Architecture of f_{tRNN} , f_{CLS} and f_{REG} . In f_{tRNN} , pooling and unpooling with spatial indices are applied to cut down the input and output complexity of LSTM module. In f_{REG} , intermediate feature maps from the encoder, *i.e.* conv2.2 and conv4.3, are concatenated to incorporate both global and local features.

structure from the highly compressed feature map. Besides, it is much more efficient to store the spatial indices than to memorize the entire feature map in float precision as proposed in FCNs [24].

The decoder is symmetrical to the encoder with a mirrored configuration but replacing all max-pooling layers with corresponding unpooling layers. The final output of the decoder is a $(L+1)$ -channel response map which is fed to a softmax classifier to predict pixel-wise confidence. We find that batch normalization [15] can significantly boost the training speed as it can effectively reduce internal shift within a mini batch. Therefore, batch normalization and rectified linear unit (ReLU) [30] are applied after each convolutional layer.

4.2 f_{sRNN} and f_{tRNN}

As shown in Figs. 1 and 2, f_{sRNN} maps the $(L+1)$ -channel response map \mathcal{M} to a single-channel label map \mathbf{z} . This mapping can be achieved efficiently in two steps. First, we merge \mathcal{M} to a single map with $(L+1)$ clusters. The value of the map at location (i, j) is set to the channel index of \mathcal{M} that has the largest confidence:

$$m_{ij} = \underset{l}{\operatorname{argmax}}(\mathcal{M}_{ij})_l, \text{ where } l = 0, \dots, L. \quad (9)$$

The second step is to generate a label map from the clustering. We label each landmark with a small square centered at the corresponding clustering center with varied sizes. The sizes are set to 7-pixel, 5-pixel, and 3-pixel for the three recurrent steps, respectively, in order to provide the spatial feedback in a coarse-to-fine manner.

We employ Long Short-Term Memory (LSTM) [13, 31] networks to model f_{tRNN} . 256 hidden units are used in the LSTM layer we empirically set $T = 10$. The prediction loss is calculated at each time step and then accumulated after T steps for backpropagation. Directly feeding C_{pe}^t into the LSTM layer leads to a low training rate as it needs $4 \times 4 \times 256 = 4096$ neurons for both the input and output. We apply 4×4 pooling and unpooling to compress C_{pe} to a 256×1 vector as illustrated in Fig. 5.

4.3 f_{CLS} and f_{REG}

To facilitate the decoupling in the bottleneck, we use a classification network to predict identity labels from \mathcal{C}_{id} . f_{CLS} takes \mathcal{C}_{id} as input and applies 4×4 average pooling to obtain a $256d$ feature vector for identity representation. Instead of using a very long feature vector in former face recognition network [43], *e.g.* $4096d$, we use a more compact vector, *e.g.* $256d$, to reduce the computational cost without losing recognition accuracy [38, 42]. To avoid overfitting, 0.4 dropout is applied, followed with a fully connected layer with M neurons to predict the entity using the cross-entropy loss.

The regression network takes $128 \times 128 \times (L + 1)$ response map as input to directly predict $2L \times 1$ normalized landmark coordinates. The network architecture is similar to the encoder but using fewer feature maps in each convolutional layer: 64-64-256-256-512. The dimension of feature maps is halved after each 2×2 max-pooling layer except the last 8×8 pooling layer to achieve a $512d$ feature vector. Similar to the classification network, 0.4 dropout is applied. A fully connected layer with $2L \times 1$ neurons is used to output landmark coordinates, which is used to compute the Euclidean loss.

We experienced suboptimal performance with the designed f_{REG} at the beginning. The reason is the response map is highly abstract and missing detailed information of the input image. To address this issue, we incorporate feature maps from the encoder to boost the regression accuracy. More specifically, we concatenate feature maps from both shallow layer ($conv_{2-2}$) and deep layer ($conv_{4-3}$) to the corresponding layers in f_{REG} to utilize both global and local features. Figure 5 illustrates the idea. Both $conv_{2-2}$ and $conv_{4-3}$ are learned in the encoder-decoder and remain unchanged in f_{REG} .

5 Experiments

In this section, we first demonstrate the effectiveness of each component in our framework, followed with performance comparison against the state-of-the-arts on both controlled and unconstrained datasets.

5.1 Datasets and Settings

Datasets. We conduct our experiments on widely used benchmark datasets as listed in Table 1. These datasets present challenges in multiple aspects such as large pose, extensive expression variation, severe occlusion and dynamic illumination.

We generated 7-landmark annotation for all datasets to locate eye corners, nose tip and mouth corners. Besides, we followed [37] for unified 68-landmark annotation for Helen, LFPW, Talking Face (TF), Face Movie (FM) and 300-VW. Moreover, we manually labeled the identity for each video in TF, FM and 300-VW. The landmark annotation of LFW is given by [23].

AFLW and 300-VW have the largest number of labeled images. They are also more challenging than others due to the extensive variations. Therefore,

Table 1. The *image* and *video datasets* used in training and evaluation. LFW, TF, FM and 300-VW have both landmark and identity annotation. AFLW and 300-VW are split into two sets for both training and evaluation.

	AFLW [20]	LFW [23]	Helen [22]	LFPW [4]	TF [11]	FM [36]	300-VW [39]
In-the-wild	Y	Y	Y	Y	N	Y	Y
Image	21,080	12,007	2,330	1,035	500	2,150	114,000
Video	-	-	-	-	5	6	114
Landmark	21	7	194	68	68	68	68
Identity	-	5,371	-	-	1	6	105
Training	16,864	12,007	2,330	1,035	0	0	90,000
Evaluation	4,216	0	0	0	500	2150	24,000

we used them for both training and evaluation. More specifically, 80% of the images in AFLW and 90 out of 114 videos in 300-VW were used for training, and the rest were used for evaluation. We sampled videos to roughly cover the three different scenarios defined in [8], *i.e.* “Scenario 1”, “Scenario 2” and “Scenario 3”, corresponding to well-lit, mild unconstrained and completely unconstrained conditions, respectively.

We performed data augmentation by sampling ten variations from each image in the image training datasets. The sampling was achieved by random perturbation of scale (0.9 to 1.1), rotation ($\pm 15^\circ$), translation (7 pixels), as well as horizontal flip. To generate sequential training data, we randomly sampled 100 clips from each training video, where each clip has 10 frames. It is worthy mentioning that no augmentation is applied on video training data to preserve the temporal consistency in the successive frames.

Training. Our approach is capable of end-to-end training on the video datasets. However, there are only 105 different identities in 300-VW. To make full use of all annotated datasets, we conducted the training through three steps. In each step, we optimized the network parameters by using *stochastic gradient descent* (SGD) with 0.9 momentum. The learning rate started at 0.01 and decayed 20% after every 10 epochs.

In the first step, we trained the network without f_{CLS} and f_{tRNN} using AFLW, Helen and LFPW. We initialized f_{ENC} using pre-trained weights in VGG-16 [40], and left other modules with Gaussian initialization [16]. The training was performed for 30 epochs. In the second step, we added f_{CLS} and fine-tuned other modules using LFW. The training was performed for 20 epochs. In the third step, we added f_{tRNN} and fine-tuned the entire network using 300-VW. The mini-batch size was set to 5 clips that had no identity overlap to avoid oscillations of the identity loss. For each training clip, we performed temporal recurrent learning for another 50 epochs in both forward and backward direction to double the training data.

Evaluation. To avoid overfitting, we ensure that the training and testing videos do not have identity overlap on the 300-VW (16 videos share 7 identities).

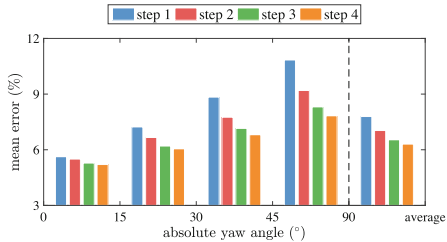


Fig. 6. Mean errors after each spatial recurrent step on the validation set of AFLW [20]. The fitting improvement is more significant on faces with large head poses (45° – 90°) than near frontal faces (0° – 15°). Three-step recurrent learning achieve a good trade-off between fitting accuracy and efficiency, as the fourth step has very limited improvement.




Fig. 7. Examples of three-step spatial recurrent learning. Successive recurrent steps are not necessary in easy cases (first row), but is crucial in challenging cases such as large pose and intense expression (rest of rows). The response clusters shrink and converge in successive recurrent steps, which moves landmarks toward ground truth step by step.

Table 2. Mean error comparison between the proposed spatial recurrent learning and the widely used cascade learning on large pose ($> 30^\circ$) set of AFLW. Each network in cascade has exactly the same architecture as the recurrent version but not sharing weight among cascades. The recurrent learning beats the cascade variant in terms of fitting accuracy and efficiency.

	Left eye (%)	Right eye (%)	Nose (%)	Mouth (%)	Mean (%)	Std (%)	Failure (%)	Time (ms)	Memory (MB)
Cascade	8.32	7.37	9.60	7.52	8.07	6.36	25.5	31.7	88.9
Recurrent	7.59	7.21	8.76	6.45	7.33	3.94	17.3	28.5	29.6

We used normalized *root mean square error* (RMSE) [37] for fitting accuracy evaluation. A prediction with larger than 10% mean error was reported as a failure [39, 44].

5.2 Validation of Spatial Recurrent Learning

We validate the proposed spatial recurrent learning on the validation set of AFLW. To better investigate the benefits of spatial recurrent learning, we partitioned the validation set into four image groups according to the absolute value of yaw angle [35]: 0° – 15° , 15° – 30° , 30° – 45° and 45° – 90° .

First, we trained a 4-step recurrent model and reported the mean error after each step in Fig. 6. From which, we had the following observations: **(1)** The fitting errors decrease in the successive recurrent steps. **(2)** The improvement of fitting accuracy is much more significant on faces with large head poses than

Table 3. Mean error comparison between the proposed temporal recurrent learning and the variant without f_{tRNN} on the validation set of 300-VW [37]. The temporal recurrent learning significantly improves the tracking accuracy (smaller mean error) and robustness (smaller std and lower failure rate), especially on the validation set in challenging settings.

	Common			Challenging			Full		
	Mean (%)	Std (%)	Failure (%)	Mean (%)	Std (%)	Failure (%)	Mean (%)	Std (%)	Failure (%)
w/o f_{tRNN}	4.52	2.24	3.48	6.27	5.33	13.3	5.83	3.42	6.43
f_{tRNN}	4.21	1.85	1.71	5.64	3.28	5.40	5.25	2.15	2.82

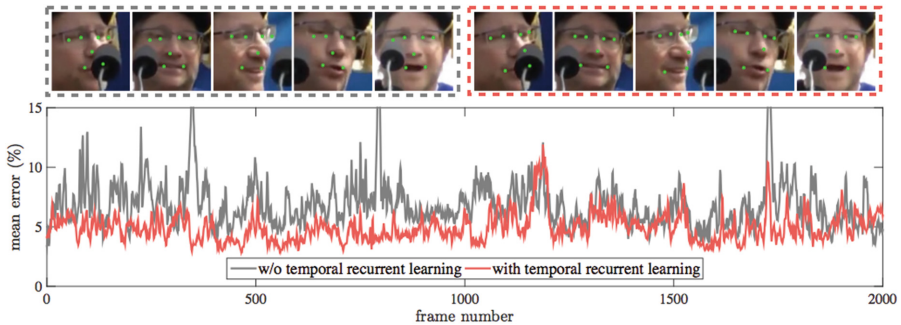


Fig. 8. Examples of validation results in challenging settings. The tracked subject undergoes intensive pose and expression variations as well as severe partial occlusions. The proposed temporal recurrent learning has substantial improvement in terms of tracking accuracy and robustness, especially for landmarks on nose tips and mouth corners.

near frontal faces, *e.g.* 23.3% improvement on 45° – 90° set and 6.10% improvement on 0° – 15° pose set. **(3)** The improvement is saturated after the first three recurrent steps as the fourth step has very limited improvement. These observations validate the proposed spatial recurrent learning to improve the fitting accuracy especially in challenging cases such as large pose. Besides, we set the number of recurrent steps to 3 in the following experiments, as it achieves a good trade-off between fitting accuracy and efficiency. Figure 7 shows examples of recurrent learning. The response clusters shrink and converge in successive recurrent steps, which moves landmarks from initial to ground truth step by step.

Second, it is reasonable to compare the proposed spatial recurrent learning with the widely used cascade learning such as [41, 54]. For a fair comparison, we implemented a three-step cascade variant of our approach. Each network in the cascade has exactly the same architecture as the spatial recurrent version but there is no weight sharing among cascades. We fully trained the cascade networks

using the same training set and validated the performance on the large pose ($> 30^\circ$) set of AFLW. The comparison is presented in Table 2. We can see that the spatial recurrent learning can significantly improve the fitting performance. The underlying reason is the recurrent network learns the tep-by-step fitting strategy jointly, while the cascade networks learn each step independently. It can better handle the challenging case where the initial guess is usually far away from the ground truth. Moreover, a single network with shared weights can instantly reduce the memory usage to one third of the cascaded implementation.

5.3 Validation of Temporal Recurrent Learning

In this section, we validate the proposed temporal recurrent learning on the validation set of 300-VW. To better study the performance under different settings, we split the validation set into two groups: 9 videos in common settings that roughly match “Scenario 1”, and 15 videos in challenging settings that roughly match “Scenario 2” and “Scenario 3”. The common, challenging and full sets were used in the following evaluation.

We implemented a variant of our approach that turns off the temporal recurrent learning f_{tRNN} . It was also pre-trained on the image training set and fine-tuned on the video training set. Since there was no temporal recurrent learning, we used frames instead of clips to conduct the fine-tuning which was performed for the same 50 epochs. We showed the result with and without temporal recurrent learning in Table 3.

For videos in common settings, the temporal recurrent learning achieves 6.8% and 17.4% improvement in terms of mean error and standard deviation respectively, while the failure rate is remarkably reduced by 50.8%. Temporal modeling produces better prediction by taking consideration of history observations.

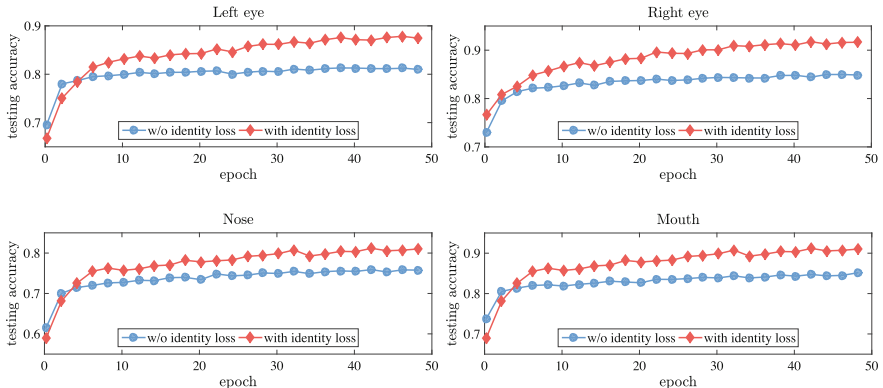


Fig. 9. Testing accuracy of different facial components with respect to the number of training epochs. The proposed supervised identity disentangling helps to achieve a more complete factor decoupling in the bottleneck of the encoder-decoder, which yields better generalization capability and more accurate testing results.

It may implicitly learn to model the motion dynamics in the hidden units from the training clips.

For videos in challenging settings, the temporal recurrent learning won with even bigger margin. Without f_{tRNN} , it is hard to capture the drastic motion or changes in consecutive frames, which inevitably results in higher mean error, std and failure rate. Figure 8 shows an example where the subject exhibits intensive pose and expression variations as well as severe partial occlusions. The curve showed our recurrent model obviously reduced landmark errors, especially for landmarks on nose tip and mouth corners. The less oscillating error also suggests that f_{tRNN} significantly improves the prediction stability over frames.

5.4 Benefits of Supervised Identity Disentangling

The supervised identity disentangling is proposed to better decouple the temporal-invariant and temporal-variant factors in the bottleneck of the encoder-decoder. This facilitates the temporal recurrent training, yielding better generalization and more accurate fittings at test time.

To study the effectiveness of the identity network, we removed f_{CLS} and follow the exact training steps. The testing accuracy comparison on the 300-VW dataset is shown in Fig. 9. The accuracy was calculated as the ratio of pixels that were correctly classified in the corresponding channel(s) of the response map.

The validation results of different facial components show similar trends: (1) The network demonstrates better generalization capability by using additional identity cues, which results in a more efficient training. For instance, after only 10 training epochs, the validation accuracy for landmarks located at the left eye reaches 0.84 with identity loss compared to 0.8 without identity loss. (2) The supervised identity information can substantially boost the testing accuracy. There is an approximately 9% improvement by using the additional identity loss. It worth mentioning that, at the very beginning of the training (< 5 epochs), the network has inferior testing accuracy with supervised identity disentangling.

Table 4. Mean error comparison with state-of-the-art methods on multiple video validation sets. The top performance in each dataset is highlighted. Our approach achieves the best fitting accuracy on both controlled and unconstrained datasets.

	7 landmarks				68 landmarks		
	TF [11]	FM [36]	300-VW [39]		TF [11]	FM [36]	300VW [39]
	%	%	Challenging		%	%	Challenging
DRMF [2]	4.43	8.53	9.16	ESR [6]	3.49	6.74	7.09
ESR [6]	3.81	7.58	7.83	SDM [50]	3.80	7.38	7.25
SDM [50]	4.01	7.49	7.65	CFAN [54]	3.31	6.47	6.64
IFA [1]	3.45	6.39	6.78	TCDCN [55]	3.45	6.92	7.59
DCNC [41]	3.67	6.16	6.43	PIEFA [36]	3.24	6.07	6.37
OURS	3.32	5.43	5.64	OURS	3.17	6.18	6.25

It is because the suddenly added identity loss perturbs the backpropagation process. However, the testing accuracy with identity loss increases rapidly and outperforms the one without identity loss after only a few more training epochs.

5.5 Comparison with State-of-the-Art Methods

We compared our framework with both traditional approaches and deep learning based approaches. The methods with hand-crafted features include: **(1)** DRMF [2], **(2)** ESR [6], **(3)** SDM [50], **(4)** IFA [1], and **(5)** PIEFA [36]. The deep learning based methods include: **(1)** DCNC [41], **(2)** CFAN [54], and **(3)** TCDCN [55]. All these methods were recently proposed and reported state-of-the-art performance. For fair comparison, we evaluated these methods in a tracking protocol: fitting result of current frame was used as the initial shape (DRMF, SDM and IFA) or the bounding box (ESR and PIEFA) in the next frame. The comparison was performed on both controlled, *e.g.* Talking Face (TF) [11], and in-the-wild datasets, *e.g.* Face Movie (FM) [36] and 300-VW [39].

We report the evaluation results for both 7 and 68 landmark setups in Table 4. Our approach achieves state-of-the-art performance under both settings. It outperforms others with a substantial margin on all datasets under 7-landmark evaluation. The performance gain is more significant on the challenging datasets (FM and 300-VW) than controlled dataset (TF). The performance of our approach degrades slightly under 68-landmark evaluation. It is a reasonable degradation considering training images (3k) that have 68-landmark annotation are much less than the ones that have 7-landmark annotation (30k). Although the training set of 300-VW contains 90k frames, the variations are limited as only 105 different identities are present. Our alignment model runs fairly fast, it takes around 30ms to process an image using a Tesla K40 GPU accelerator.

6 Future Work

In this paper, we proposed a novel recurrent encoder-decoder network for real-time sequential face alignment. Intensive experiments demonstrated the effectiveness of our framework and its superior performance. It decouples temporal-invariant and -variant factors in the bottleneck of the network, and exploits recurrent learning at both spatial and temporal dimensions.

The proposed method provides a general framework that can be further applied to other localization-sensitive tasks, such as human pose estimation, object detection, scene classification, etc. In the future, we plan to further exploit the proposed recurrent encoder-decoder network for boarder impact.

References

1. Asthana, A., Zafeiriou, S., Cheng, S., Pantic, M.: Incremental face alignment in the wild. In: CVPR (2014)
2. Asthana, A., Zafeiriou, S., Cheng, S., Pantic, M.: Robust discriminative response map fitting with constrained local models. In: CVPR, pp. 3444–3451 (2013)

3. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: a deep convolutional encoder-decoder architecture for image segmentation. *CoRR* (2015)
4. Belhumeur, P.N., Jacobs, D.W., Kriegman, D.J., Kumar, N.: Localizing parts of faces using a consensus of exemplars. In: *CVPR* (2011)
5. Black, M., Yacoob, Y.: Tracking and recognizing rigid and non-rigid facial motions using local parametric models of image motion. In: *CVPR*, pp. 374–381 (1995)
6. Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. *IJCV* **107**(2), 177–190 (2014)
7. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: encoder-decoder approaches. *CoRR* abs/1409.1259 (2014)
8. Chrysos, G.G., Antonakos, E., Zafeiriou, S., Snape, P.: Offline deformable face tracking in arbitrary videos. In: *ICCVW*, pp. 954–962 (2015)
9. Cootes, T.F., Taylor, C.J.: Active shape models-smart snakes. In: *BMVC* (1992)
10. Decarlo, D., Metaxas, D.: Optical flow constraints on deformable models with applications to face tracking. *IJCV* **38**(2), 99–127 (2000)
11. FGNet: talking face video. Technical report (2004). http://www-prima.inrialpes.fr/FGnet/data/01-TalkingFace/talking_face.html
12. Gao, X., Su, Y., Li, X., Tao, D.: A review of active appearance models. *IEEE Trans. Syst. Man Cybern.* **40**(2), 145–158 (2010)
13. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
14. Hong, S., Noh, H., Han, B.: Decoupled deep neural network for semi-supervised semantic segmentation. *CoRR* abs/1506.04924 (2015)
15. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. *CoRR* abs/1502.03167 (2015)
16. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: convolutional architecture for fast feature embedding. In: *ACMM*, pp. 675–678 (2014)
17. Jourabloo, A., Liu, X.: Large-pose face alignment via cnn-based dense 3D model fitting. In: *CVPR* (2016)
18. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: *CVPR*, June 2015
19. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian segnet: model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *CoRR* abs/1511.02680 (2015)
20. Koestinger, M., Wohlhart, P., Roth, P.M., Bischof, H.: Annotated facial landmarks in the wild: a large-scale, real-world database for facial landmark localization. In: *Workshop on Benchmarking Facial Image Analysis Technologies* (2011)
21. Lai, H., Xiao, S., Cui, Z., Pan, Y., Xu, C., Yan, S.: Deep cascaded regression for face alignment (2015). [arXiv:1510.09083v2](https://arxiv.org/abs/1510.09083v2)
22. 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. LNCS*, vol. 7574, pp. 679–692. Springer, Heidelberg (2012). doi:10.1007/978-3-642-33712-3_49
23. Learned-Miller, G.: Labeled faces in the wild: updates and new reporting procedures. Technical report. UM-CS-2014-003, University of Massachusetts, Amherst (2014)
24. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. *CoRR* abs/1411.4038 (2014)

25. Long, J.L., Zhang, N., Darrell, T.: Do convnets learn correspondence? In: NIPS, pp. 1601–1609 (2014)
26. Lu, L., Zhang, X., Cho, K., Renals, S.: A study of the recurrent neural network encoder-decoder for large vocabulary speech recognition. In: INTERSPEECH (2015)
27. Mikolov, T., Joulin, A., Chopra, S., Mathieu, M., Ranzato, M.: Learning longer memory in recurrent neural networks. CoRR abs/1412.7753 (2014)
28. Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., Khudanpur, S.: Recurrent neural network based language model. In: INTERSPEECH, pp. 1045–1048 (2010)
29. Milborrow, S., Nicolls, F.: Locating facial features with an extended active shape model. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008. LNCS, vol. 5305, pp. 504–513. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88693-8_37](https://doi.org/10.1007/978-3-540-88693-8_37)
30. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: ICML, pp. 807–814 (2010)
31. Oh, J., Guo, X., Lee, H., Lewis, R.L., Singh, S.: Action-conditional video prediction using deep networks in atari games. In: NIPS, pp. 2845–2853 (2015)
32. Oliver, N., Pentland, A., Berard, F.: Lafter: lips and face real time tracker. In: CVPR, pp. 123–129 (1997)
33. Patras, I., Pantic, M.: Particle filtering with factorized likelihoods for tracking facial features. In: Proceedings of Automatic Face and Gesture Recognition, pp. 97–102 (2004)
34. Peng, X., Huang, J., Hu, Q., Zhang, S., Elgammal, A., Metaxas, D.: From circle to 3-sphere: head pose estimation by instance parameterization. CVIU **136**, 92–102 (2015)
35. Peng, X., Huang, J., Hu, Q., Zhang, S., Metaxas, D.N.: Three-dimensional head pose estimation in-the-wild. In: FG, vol. 1, pp. 1–6 (2015)
36. Peng, X., Zhang, S., Yang, Y., Metaxas, D.N.: Piefa: personalized incremental and ensemble face alignment. In: ICCV (2015)
37. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: 300 faces in-the-wild challenge: the first facial landmark localization challenge. In: ICCVW (2013)
38. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: a unified embedding for face recognition and clustering. In: CVPR, pp. 815–823 (2015)
39. Shen, J., Zafeiriou, S., Chrysos, G., Kossaiji, J., Tzimiropoulos, G., Pantic, M.: The first facial landmark tracking in-the-wild challenge: benchmark and results. In: ICCVW (2015)
40. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014)
41. Sun, Y., Wang, X., Tang, X.: Deep convolutional network cascade for facial point detection. In: CVPR, pp. 3476–3483 (2013)
42. Sun, Y., Wang, X., Tang, X.: Deeply learned face representations are sparse, selective, and robust. In: CVPR, pp. 2892–2900 (2015)
43. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: closing the gap to human-level performance in face verification. In: CVPR (2014)
44. Tang, M., Peng, X.: Robust tracking with discriminative ranking lists. TIP **21**(7), 3273–3281 (2012)
45. Tzimiropoulos, G.: Project-out cascaded regression with an application to face alignment. In: CVPR, pp. 3659–3667 (2015)
46. Veeriah, V., Zhuang, N., Qi, G.J.: Differential recurrent neural networks for action recognition. In: ICCV, December 2015
47. Wang, J., Cheng, Y., Feris, R.S.: Walk and learn: facial attribute representation learning from egocentric video and contextual data. In: CVPR (2016)

48. Wang, X., Yang, M., Zhu, S., Lin, Y.: Regionlets for generic object detection. *TPAMI* **37**(10), 2071–2084 (2015)
49. Wu, Y., Ji, Q.: Constrained joint cascade regression framework for simultaneous facial action unit recognition and facial landmark detection. In: *CVPR* (2016)
50. Xuehan-Xiong, D., la Torre, F.: Supervised descent method and its application to face alignment. In: *CVPR* (2013)
51. Yang, J., Reed, S., Yang, M.H., Lee, H.: Weakly-supervised disentangling with recurrent transformations for 3D view synthesis. In: *NIPS* (2015)
52. Yao, L., Torabi, A., Cho, K., Ballas, N., Pal, C., Larochelle, H., Courville, A.: Describing videos by exploiting temporal structure. In: *ICCV*, December 2015
53. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8689, pp. 818–833. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1_53](https://doi.org/10.1007/978-3-319-10590-1_53)
54. 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*. LNCS, vol. 8690, pp. 1–16. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10605-2_1](https://doi.org/10.1007/978-3-319-10605-2_1)
55. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: Facial landmark detection by deep multi-task learning. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8694, pp. 94–108. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10599-4_7](https://doi.org/10.1007/978-3-319-10599-4_7)
56. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.H.S.: Conditional random fields as recurrent neural networks. In: *ICCV*, December 2015
57. Zhu, S., Li, C., Loy, C.C., Tang, X.: Face alignment by coarse-to-fine shape searching. In: *CVPR*, pp. 4998–5006 (2015)
58. Zhu, X., Lei, Z., Liu, X., Shi, H., Li, S.Z.: Face alignment across large poses: a 3D solution. In: *CVPR* (2016)