# **Face Parsing via Recurrent Propagation**

Sifei Liu<sup>1</sup> sliu32@ucmerced.edu Jianping Shi<sup>2</sup> shijianping@sensetime.com Ji Liang<sup>3</sup> liangji20040249@gmail.com Ming-Hsuan Yang<sup>1</sup> mhyang@ucmerced.edu <sup>1</sup> University of California, Merced <sup>2</sup> SenseTime, Inc. <sup>3</sup> Momenta, Inc.

#### Abstract

Face parsing is an important problem in computer vision that finds numerous applications including recognition and editing. Recently, deep convolutional neural networks (CNNs) have been applied to image parsing and segmentation with the state-of-the-art performance. In this paper, we propose a face parsing algorithm that combines hierarchical representations learned by a CNN, and accurate label propagations achieved by a spatially variant recurrent neural network (RNN). The RNN-based propagation approach enables efficient inference over a global space with the guidance of semantic edges generated by a local convolutional model. Since the convolutional architecture can be shallow and the spatial RNN can have few parameters, the framework is much faster and more light-weighted than the state-of-the-art CNNs for the same task. We apply the proposed model to coarse-grained and fine-grained face parsing. For fine-grained face parsing, we develop a two-stage approach by first identifying the main regions and then segmenting the detail components, which achieves better performance in terms of accuracy and efficiency. With a single GPU, the proposed algorithm parses face images accurately at 300 frames per second, which facilitates real-time applications.

## **1** Introduction

The recent years have witnessed significant progress in object segmentation and image parsing using deep CNNs [1, 6, 8, 13, 12, 13]. With end-to-end nonlinear classifiers and hierarchical features, CNN-based face parsing methods [16, 26] achieve the state-of-the-art performance than approaches based on hand-crafted features [12]. The main issues with existing CNN-based face parsing are the heavy computational load and large memory requirement. Both issues can be alleviated by using shallow or light-weighted convolutional structures, but at the expense of parsing accuracy.

In this work, we propose a face parsing algorithm in which a spatially variant recurrent module is incorporated for global propagation of label information. A straightforward combination of CNN and RNN is to take each activation in a CNN feature map as the input to a hidden recurrent node in a two-dimensional (2D) spatial sequence and use the recurrent structure to learn the propagation weight matrix in an end-to-end fashion [II, III]. These models either utilize a spatial RNN [III], or a stacked long short-term memory (LSTM) [III]. In contrast, the proposed recurrent structure exploits the strength of both models in which we apply a simple structure similar to a typical RNN but maintains the capability of spatially variant propagation of an LSTM. Specifically, the proposed model uses a spatially variant gate map to control the propagation strength over different locations in the label space. For face parsing, this gate is naturally associated with the semantic boundary. A gate allows propagation between pixels in a label-consistent region or stops it otherwise. We show that this gate can be obtained via a relatively shallow CNN that focuses on learning low and mid-level image features. The RNN module, controlled by the gate, can utilize rich redundant information by propagating the predicted labels to their neighboring pixels in the label-consistent region. Compared to a deep CNN face parser with similar performance, the propagation layer requires a small amount of model parameters and significantly reduces the computational cost. As a result, we construct a model that is hundreds of times faster and smaller than deep CNN-based methods [III], III] for face parsing without loss of accuracy.

We validate the proposed algorithm on both coarse-grained (parsing an image with major regions including skin, hair and background) and fine-grained (parsing an image with detailed facial components such as eyes, eyebrows, nose and mouth) face parsing. Both are of critical importance for real-world applications in face processing, e.g., coarse-grained face parsing for style transfer [**D**] and fine-grained face parsing for virtual makeup. Parsing only the main classes is generally easier under the same settings due to the complexity of solutions and more balanced distributions of training samples. We show that the proposed model can parse all faces of an image in one shot, and significantly outperform the state-of-the-art methods in terms of accuracy and speed.

One issue with applying a single network to fine-grained face parsing is the performance on small facial components. This is due to the extremely unbalanced sample distributions and image size of these regions. We design a two-stage method to parse these components efficiently. We train the model for the main classes in the first stage and then focus on the others with relatively simpler sub-networks. Specifically, the sub-networks in the second stage take a cropped facial region as input. In contrast to a face component may occupy a small amount of pixels with a whole image, the distributions of the pixels for a cropped region are more balanced. We show that by dividing the second face parsing problem into several sub-tasks, the overall network complexity is significantly reduced.

The contributions of this work are summarized as follows. First, a light-weighted network is proposed for pixel-wise face parsing by combining a shallow CNN and a spatially variant RNN, which significantly reduces the computational load of deep CNN. Second, We show that when parsing a face image with multiple detailed components, dividing the problem into several sub-tasks is significantly more efficient than using one single model, with even better accuracy. Experimental results on numerous datasets demonstrate the efficiency and effectiveness of the proposed face parsing algorithm against the state-of-the-art methods.

### 2 Related Work

**Face Parsing.** Face parsing considered in this work assigns dense semantic labels to all pixels in an image. Typically, only one face image is assumed to be detected in an input frame. Several approaches have been developed based on graphical models [12], exemplars [22] and convolution networks [16, 20]. The face parsing method developed by Luo *et al.* [20]

hierarchically combines several separately trained deep models. Liu *et al.* [[13] develop a unified model to generate complete labels of facial regions in one single pipeline. In [[13] Yamashita *et al.* propose a weight cost function to deal with unbalanced samples for parsing small regions. Closest to this work is the multi-objective CNN [[13] which introduces additional supervision of semantic edges and achieves substantial improvements for coarse-grained and fine-grained face parsing. However it is less effective for parsing detailed facial components and computationally expensive. We show these issues can be largely resolved by the proposed methods.

**Recurrent Neural network.** Recurrent networks  $[\square, \square, \square, \square]$  have been shown to be effective for modeling long term dependencies in sequential data (e.g., speech). For image data, we can apply one-dimensional (1D) RNN to multiple dimensions in row/column-wise manner  $[\square, \square, \square]$ ,  $[\square]$  or multi-dimensional RNN (MDRNN)  $[\square]$  such that each neural node can receive informations from multiple directions  $[\square, \square]$  (as opposed to one direction in the conventional RNN). In addition, there are other variants that leverage these two models, e.g., the grid LSTM  $[\square]$ ,  $[\square]$ . The proposed model belongs to the first category, which is easier for parallelization.

The proposed recurrent model is closely related to two recent image parsing methods  $[\square, \square]$  that utilize linear recurrent formulation to associate the adjacent pixels in either the semantic label space or the low-level image space. In  $[\square]$  Chen *et al.* propose the recurrent model with the concept of domain transform, where object edges learned on top of the intermediate layers of a fully convolutional network (FCN  $[\square]$ ) are used to regularize the transforms between adjacent pixels. Liu *et al.*  $[\square]$  further extend the recurrent structure to high-order recursive filters to model more variations in the low-level space. We compare different RNN frameworks with respect to the spatially invariant  $[\square]$  and variant (as proposed in this work) recurrent modules. In addition, we show that the formulation with recurrent propagation can substantially simplify the network structures with lower computational loads, which are not well exploited in the existing methods.

## **3** Proposed Algorithm

Most CNN-based face parsing algorithms [**II**], **KI**] apply deep networks with a large number of parameters, which entail heavy computational loads. On the other hand, shallow models can be executed efficiently but not able to model global data dependency. In this work, we use a shallow CNN with a combination of spatially variant recurrent propagation module to model image data effectively and efficiently.

Our model contains a shallow CNN and a spatial RNN, as shown in Figure 1. First, the CNN takes a color image as its input and learns a coarse pixel-wise label score map (Figure 1(b)). Second, the coarse label result is fed to a spatial recurrent unit for global propagation. Specifically, the spatial propagation is controlled by a gate map (Figure 1(c)), which is referred to as a recurrent gate in the rest of the paper. Each pixel in the map, formulated as a scalar weight coefficient to the recurrent term, controls the connection between two adjacent hidden nodes at the corresponding location. Since a gate map can be supervised by the ground truth semantic boundaries from labeled annotations, it enables the recurrent propagation to be discriminative between semantically consistent and inconsistent regions, with respect to the specific input image.

We first briefly review conventional RNNs and describe how we extend it to the 2D



Figure 1: Proposed parsing network architecture by combining a CNN and a spatial RNN. The CNN generates a coarse label map (b) and a recurrent gate (c), which are fed into 4 RNNs with different directions to generate a more accurate result (d). The network structure is shown where the notation for Conv1 " $5 \times 5 \times 16/1$ " means convolution layer with  $5 \times 5$  kernel, 16 channels and stride 1. The face image in (d) is further segmented with detailed labels in the second stage (see text and Figure 2).

space for image data, before introducing the recurrent gates. We then discuss how to train the hybrid model in an end-to-end fashion.

#### 3.1 Recurrent Neural Networks

The conventional RNN is developed to process 1D sequential data where each hidden node represents a single character, frame, pixel and is connected to its adjacent neighbor. The hidden node *i*, denoted as  $h_i \in H$  receives two inputs: an external input  $x_i \in X$  and its previous activation  $h_{i-1}$  from one step back. The summation of these two inputs is then non-linearly mapped via a function  $\theta(\cdot)$  as the activation of the current step:

$$h_i = \theta(a_i), \quad a_i = \omega_x x_i + (\omega_h h_{i-1} + b). \tag{1}$$

In this formulation,  $x_i$  and  $h_i$  can have different dimensions, where the input transition matrix  $\omega_x$  aligns  $x_i$  to have the same dimension as  $h_i$ . In addition, b is a bias or offset term to model data points centered at a point other than the origin. For simplicity, we set  $x_i$  and  $h_i$  to have the same dimension, and remove the  $\omega_x$  so that only the recurrent state transition matrix needs to be learned.

To extend the 1D RNN in (1) to 2D images, we consider each row/column as 1D sequence, and then adopt an approach similar to the bidirectional recurrent neural network for processing temporal sequences [III]. First, the 1D sequential RNN is processed respectively along left-to-right, top-to-bottom, and their reverse ways. Taking the left-to-right direction for a 2D feature/label map as an example, the 1D sequential RNN scans each row from left to right. As a result, four hidden activation maps are generated.

The four hidden activation maps can be grouped either in parallel or cascade, as introduced in [ $\square$ ]. The four maps share the same input X with the parallel method, while in the cascade way, each RNN takes the output from a previous RNN as its input. We adopt the parallel method and integrate the maps by selecting the optimal direction based on the maximum response at each location. This is carried out by a node-wise max pooling operation that can effectively select the maximally responded direction as the desired information to be propagated and reject noisy information from other directions. We note that the fourdirectional RNNs with parallel integration can be executed simultaneously with multiple GPUs for further acceleration as they are independent. The backward pass is also an extension of the back propagation through time (BPTT) method used in RNNs [23]. Due to space limitation, we only present the derivative with respect to  $a_i$ :

$$\delta_i = \theta'(a_i) \cdot (\xi_i + \omega_h \delta_{i+1}), \qquad (2)$$

where  $\omega_h$  is a square weight matrix and all the others are 1D vectors. We denote  $\xi$  the influence from the output layer on top of the proposed spatial RNN, and the second term in (2) the influence from the next hidden node. The derivatives are passed back in reverse order against the feedforward process, with four distinct directions computed respectively [29].

#### 3.2 Spatially Variant Recurrent Network

The fundamental problem of the RNN in (1) is that the hidden state transition matrix  $\omega_h$  is spatially invariant. As such, it tends to propagate any pixel to its adjacent ones with a group of fixed weights. However, the label space is spatially variant with respect to different locations. The propagation between pixels that share the same label should be distinguished from those between pixels with different labels on the semantic boundaries.

To this end, we propose a spatially variant recurrent network with gate maps  $g_i \in G$  as an additional input to the spatial RNN. Each  $g_i$  is an additional coefficient that controls the strength of connections between nodes to guide the recurrent propagations. Intuitively, strong connections (e.g.,  $g_i$  is close to 1) should be enforced between nodes in the label-consistent region. On the other hand, weak connections (e.g.,  $g_i$  is close to 0) should be assigned to the nodes belonging to semantically different categories, so that they can be successfully separated.

To reformulate the framework, we have two types of inputs to a RNN, i.e., an external input X, and a spatially variant gate G. Given a hidden node  $h_i$ , the spatially controllable recurrent propagation is:

$$a_i = x_i + g_i \cdot (\omega_h h_{i-1} + b). \tag{3}$$

The propagation of the hidden activation at i-1 to i is controlled by dot product with  $g_i$ . We use the identity function  $\theta(x) = x$  as the activation (also used by [**B**, **LD**]), since experimentally it achieves better performance. To maintain the stability of the linearized formulation, the absolute value of  $g_i$ , and norm of  $\omega_h$  are both normalized to be within one during parameter update in order to prevent the hidden activation in H to be increased exponentially.

Similar to the sequential RNN, the BPTT algorithm is adopted to adjust X and G in the spatially variant RNN. The derivatives with respect to  $a_i$  and  $g_i$ , denoted as  $\delta_i$  and  $\varepsilon_i$  are:

$$\delta_i = \xi_i + g_i \cdot \omega \delta_{i+1}, \qquad \varepsilon_i = \delta_i \cdot (\omega_h h_{i+1} + b). \tag{4}$$

In addition, the derivative from RNN with respect to  $x_i$  is equal to  $\delta_i$ .

#### 3.3 Hybrid Model of CNN and RNN

In the proposed framework, we apply a CNN that provides label representation X and spatially variant gate representation G to the spatial RNN (see Figure 1). With the effective propagation of RNN, the CNN can be relatively shallow as revealed in the experimental analyses. Taking the three-class face parsing as an example, the main part of CNN is equipped with only three convolutional layers, two max pooling (down-sampling) as well as deconvolutional (upsampling) layers, as shown in Figure 1, and at most 32 channels for each layer.

The proposed network is significantly smaller than most existing CNN-based face parsing models based on 6 convolutional layers with 2 fully-connected layers [16], or 16 layers [26] from VGG [23].

To connect with the spatial RNN, the feature maps with 16 channels generated from the first deconvolutional layer (Deconv6 in Figure 1) are equally split into two components (each with 8 channels), where one is for pixel-wise labels and the other is for the recurrent gate, with equal width and height. They are then fed to four recurrent layers with different directions as X and G respectively, where each pixel *i* in the hidden layers is processed by combining  $x_i$  and  $g_i$  based on (3).

The hybrid network contains three different loss layers. At the top of the CNN, both X and G are supervised with the softmax cross entropy loss. The labeling representations are transferred by a convolutional layer to be directly supervised by the ground truth labels (see Figure 1(b)). The gate representations are transferred by a  $1 \times 1$  convolutional layer to have a single channel output, which is supervised by the boundary between different categories (see Figure 1(c)). Finally, the output of RNN with 8 channels are transferred to 3 channels, upsampled to the original image scale, and supervised by the ground truth labels (see Figure 1(d)). All the losses encourage the CNN to learn better label candidates as well as guidances to the propagation. Specifically, the ground truth boundaries are obtained from the annotated labels, by setting its boundary pixels to zeros and all the others to one. For example, with a pixel *i* that is located on a boundary of two categories, the ground truth value is set to zero, which can encourage the  $g_i$  to "cut off" the connection between different classes, and vise-versa.

### 4 Sub-networks for the Detailed Components

As discussed in Section 1 and revealed in the experiments, a single network does not perform well on small facial components. One problem is that some facial components amount to small percentages of the entire dataset, e.g., the eye regions in Figure 1(a) occupy less than 1% of the whole image. It is difficult to parse such components in one stage due to unbalanced labeled data. The work of [[[6]] applies a simple strategy by sampling with an equal number of input patches. However, the performance on small facial components is not satisfactory compared to categories with more pixels, e.g., skin. The other problem is the limited resolution of facial components. With a larger input image, more details of the components can be learned. However, it requires deeper or larger models to adapt to the enlarged receptive fields. For a single model, it is a trade-off between effectiveness and efficiency.

We decompose a unified face segmentation network into a two-stage framework. In practice, parsing major classes with either frontal, canonical face or multiple random faces can be handled using the first stage only. For parsing 11 classes in the HELEN dataset, each component can be labeled independently first and then combined with the major ones.

**First Stage Face-Hair-Background Network.** The first stage network classifies an image into skin, hair and background regions using the combination of CNN and RNN, as introduced in Section 3. Since there are only three labels with relatively equal distribution, we do not need to balance the samples. As these classes do not contain detailed structures such as facial components, the input resolution does not need to be high. Similar to [III], a face image is detected and roughly aligned to the center using [II], with a resolution of  $128 \times 128$  pixels. The result of the label has the same resolution as the input image.



Figure 2: The second stage network operates on the cropped region, i.e., left and right eyes, nose, and mouth, to parse accurate facial components. The final parsing result in (d) is the combination of segments from two stages.

Second Stage Facial Component Networks. We locate the facial components for high resolution faces image through 5 detected key points (eye centers, nose tip, and mouth corners) [23], and crop the patches accordingly. We train three simple and efficient networks to segment eye and eyebrow, nose, and mouth regions, respectively. Figure 2(b) shows the structure of eye/eyebrow network. It contains five convolution layers, two max-pooling layers, and two deconvolution layers, with an input size of  $64 \times 64$ . Similar network structures are used for the nose as well as the mouth, and the input image size is  $64 \times 64$  and  $32 \times 64$ , respectively. Since each image is cropped around each facial component, it does not include many pixels from the skin region. Therefore, the sample distribution is balanced for network training. The final parsing result is composed of the accurate facial component segments in the second stage and the coarse segments in the first stage. Since the segmentation task in the second stage is easier, we do not apply the spatial RNN for efficiency reason.

## 5 Experimental Results

We carry out experiments on images containing one or multiple faces. For single face parsing, we evaluate our method on the LFW-PL [12] and HELEN [12] datasets. In addition, we develop a Multi-Face dataset to evaluate parsing numerous faces in one image. All experiments are conducted on a Nvidia GeForce GTX TITAN X GPU.

### 5.1 Datasets and Settings

**LFW-PL and HELEN Datasets.** The LFW part label (LFW-PL) dataset contains 2,927 face images. Each face image is annotated as skin, hair or background using superpixels, and roughly aligned to the center [II]. The HELEN dataset contains 2,330 face images with manually labeled facial components including eyes, eyebrows, nose, lips, etc. For both datasets, the most centered face in each image is annotated. We adopt the same setting of data splits as [III] and resize each image and its corresponding label to  $128 \times 128$ . For the HELEN, dataset, the hair region is trained as one category in the first stage of our algorithm but is not evaluated for fair comparisons with the existing method [III].

**Multi-Face Dataset.** We collect a Multi-Face dataset where each image contains multiple faces. It contains 9,645 images in unconstrained environments with pixel-wise labels including skin, hair, and background. This dataset is divided into a training set of 9,045 images, a test set of 200 images, and a validation set of 200 images. We rescale each image and its corresponding label according to the length of the long side to maintain the aspect ratio. Each one is zero padded to result in a 512 × 512 image where all faces appear clearly.



(a) (b) (c) (d) (e) (f) Figure 3: Face parsing results on the LFW-PL dataset. (a) input image. (b) ground-truth annotations. (c) results from [III]. (d) results from CNN-S. (e) results from CNN with dense CRF. (f) results by RNN-G. More results are presented in the supplementary material.



(a) (b) (c) (d) (e) (f) Figure 4: Parsing results on the Multi-Face dataset. (a) input image. (b) results by the baseline CNN. (c) results by the standard RNN. (d) results from RNN-G. (e) the ground truth. (f) a visualized version of RNN-G. Our method is able to effectively and efficiently parse multiple faces in the cluttered background. More results are presented in the supplementary material.

**Network Implementation.** Our network structures are described in Figure 1 and 2. We use the first stage model (see Figure 1) to parse images in the LFW-PL and Multi-Face datasets, and the facial skin and hair regions in the Helen dataset. In addition, we use the second stage model (see Figure 2) to parse facial components of images in the HELEN dataset.

For fair comparison with the previous work, we align the input images according to [II] in the HELEN dataset. The faces in the LFW-PL dataset do not need additional processing since the released images are already coarsely aligned. On the other hand, we directly use the  $512 \times 512$  images as the network inputs, and do not preprocess any face for the Multi-Face dataset. We quantitatively evaluate and compare our model using per-pixel accuracy and F-measure for each class in all experiments.

In the first stage, the boundaries in Figure 1(c) are balanced with the ratio of positive/negative number of pixels set to 1 : 5 such that a sufficient number of boundary samples can be drawn. The training images augmented by random affine and mirror transformations for increasing the variation of training samples. The network for the Multi-Face dataset has two more  $3 \times 3 \times 16$  convolutional units (with max-pooling) and one more deconvolutional layer to adapt to the input size. The results are evaluated with the resolution of  $256 \times 256$ . The boundary loss sampling and training image augmentation strategies are uniformly applied to all experiments. For the second stage model, we crop the facial components based on the 5 facial key points from [ $\square$ ] for training and tests. We include at least additional 20% height/width of the total foreground height/width in the cropped images during training and maintain the aspect ratio.

#### 5.2 Coarse-grained Face Parsing

Face parsing with 3 classes are carried out using the first stage model on the LFW-PL and the Multi-Face datasets, respectively. We compare the proposed method, denoted as RNN-G with: (a) shallow CNN part only (CNN-S). (b) shallow CNN with the RNN module replaced

Table 1: Quantitative results on the LFW-PL dataset. "F" denotes f-score, "bg" denotes background, "-" denotes not available.

(%)	GLOC [	MO [	CNN-S	CNN-deep	CNN-CRF [	RNN [🛂]	RNN-G
F-skin	-	93.93	90.47	91.63	91.25	93.72	97.55
F-hair	-	80.70	76.09	78.30	75.21	81.21	83.43
F-bg	-	97.10	95.42	95.95	99.58	97.15	94.37
Accuracy	94.95	95.09	92.44	93.27	92.59	94.85	95.46
Time (ms)	254 (CPU)	$\sim 110$	< 1	$\sim 2$	$\sim 7$	$\sim 2$	$\sim 2$

#### Table 2: Quantitative results on the Multi-Face dataset.

(%)	CNN-deep	CNN-CRF	RNN	Det+RNN-G single	RNN-G
F-skin	75.56	77.84	73.33	81.02	87.36
F-hair	64.62	61.53	62.85	55.35	73.09
F-background	96.5	97.08	96.18	97.10	98.19
AC	93.39	94.5	92.78	94.42	96.35



(a) (b) (c) (d) (a) (b) (c) (d) Figure 5: Face parsing results on the HELEN [22] dataset. (a) input image. (b) ground-truth annotations. (c) results from [16]. (We roughly crop the results for better visual comparisons.) (d) our results with 11-class pixel-wise parsing. More results are presented in the supplementary material.

by two  $3 \times 3$  convolutional layers with 32 channels as a baseline network, denoted as CNN-Deep. We increase the number of the output channels of Deconv6 (Figure 1) from 8 to 16 to ensure that the shallow model can converge. (c) a combination of the shallow CNN and the post processing with a dense CRF, denoted as CNN-CRF, which is commonly used in recent semantic segmentation tasks [I]. (d) a standard RNN in (1) (similar to [I]) with the same CNN. We note that both [I], [I] do note have experiments on shallow networks.

We show two more baseline methods  $[\Box_2, \Box_3]$  evaluated on the LFW-PL dataset. Specifically, we adjust  $[\Box_3]$  by using only one-time feedforward with  $2 \times$  bilinear upsampling layer for fair comparisons in speed and accuracy. For the Multi-Face dataset, we use the models to parse all faces in images without using any detector. This is computationally efficient and useful for numerous applications without the need of instance-level information. We note the label distribution of the Multi-Face dataset with respect to different categories are significantly unbalanced since the vast majority of pixels belong to the background regions. Thus, we apply a data sampling strategy at each loss layer by maintaining the number of sampled background pixels as 5 times of the total number of pixels for skin and hair regions.

Table 1 and 2 show the results with similar trends on these two datasets. Overall, the shallow CNN, i.e., CNN-S, has limited performance. There is no significant improvement gain by simply adding more layers (CNN-Deep) or adding an additional dense CRF module (CNN-CRF). The standard RNN without the spatially variant gate performs better, but still worse than the proposed method. With the spatially variant gate, the RNN-G model performs significantly better than the baseline CNN-S, CNN-Deep and RNN models. The results demonstrate the effectiveness of the proposed spatially variant RNN structure. The proposed models operate at 500 fps for a  $128 \times 128$  single face image and 200 fps for a  $512 \times 512$  image with multiple faces.

Figure 3 and 4 show some parsing results on the two datasets. The proposed RNN-G

Table 3:	Quantitative	evaluation	results o	n the l	HELEN	dataset.	We d	lenote	the up	per and	lower	lips
as "U-lip"	and "L-lip",	and overall	mouth p	art as	"mouth"	, respec	tively	. See [	16, 24	] for mo	ore deta	ails.

Methods	eyes	brows	nose	in mouth	U-lip	L-lip	mouth	skin	overall
Liu et al. [🗳]	77.0	64.0	84.3	60.1	65.0	61.8	74.2	88.6	73.8
Smith <i>et al</i> . [23]	78.5	72.2	92.2	71.3	65.1	70.0	85.7	88.2	80.4
Liu <i>et al</i> . [🍱]	76.8	71.3	90.9	80.8	62.3	69.4	84.1	91.0	84.7
Ours 1-stage	63.3	53.7	87.5	65.7	54.0	72.6	80.6	91.1	78.8
Ours 2-stage	86.8	77.0	93.0	79.2	74.3	81.7	89.1	92.1	88.6

model performs favorably against the CNN-S, CNN-CRF, standard RNN, and the method using nonparametric prior and graph cut inference [II]. For Multi-Face dataset, we evaluate the alternative method using a face detector [II] and the single face parser trained on the LFW-PL dataset, which operates at 37 fps on average (depending on the number of detected faces). The RNN-G model performs favorably in the cluttered background against all alternative methods in terms of accuracy and efficiency.

#### 5.3 Fine-grained Face Parsing

In the HELEN dataset, we evaluate the parsing results following the settings in [II], where the second stage network is utilized to improve parsing results. Since the second stage takes less than 1 ms, the overall run-time for parsing a face with 11 classes can operate at 300 fps on a single GPU.

Table 3 and Figure 5 show the quantitative and qualitative parsing results. We first show that by using a single stage, the unified model cannot handle detailed facial parts even with the spatially variant RNN module. Our two-stage network performs favorably against the state-of-the-art methods, and the one stage network model, on all categories. It is worth noting that the overall F-measure achieved by the RNN-G model is 0.886, which amounts to 20% reduction in error rate from the state-of-the-art method [13]. These experimental results demonstrate that the two-stage network structure with the spatially variant gate is effective for accurate and efficient face parsing.

### 6 Conclusions

In this paper, we propose a pixel-level face parsing network by combining a shallow CNN and a spatially variant RNN. The recurrent propagation infers globally over the entire image with the guidance of a local model, which reduces the computational load of deep CNNs. We develop a two-stage approach for accurate parsing of the detailed facial component. Experimental results on the HELEN [22], LFW-PL [12] and the proposed Multi-Face datasets demonstrate the efficiency and effectiveness of the proposed face parsing algorithm against the state-of-the-art methods.

## Acknowledgment

This work is supported in part by SenseTime Inc. and the NSF CAREER Grant #1149783, gifts from NEC and Nvidia.

### References

- [1] Wonmin Byeon, Thomas M Breuel, Federico Raue, and Marcus Liwicki. Scene labeling with lstm recurrent neural networks. In *CVPR*, 2015.
- [2] Alex J. Champandard. Semantic style transfer and turning two-bit doodles into fine artworks. *arXiv preprint arXiv:1603.01768*, 2016.
- [3] Liang-Chieh Chen, Jonathan T. Barron, George Papandreou, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. arXiv preprint arXiv:1511.03328, 2015.
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015.
- [5] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259, 2014.
- [6] Grangier David, Bottou Léon, and Collobert Ronan. Deep convolutional networks for scene parsing. In *ICML Deep Learning Workshop*, 2009.
- [7] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In CVPR, 2015.
- [8] Clément Farabet, Camille Couprie, Laurent Najman, and Yann LeCun. Learning hierarchical features for scene labeling. *IEEE PAMI*, 35(8):1915–1929, 2013.
- [9] Alan Graves, Abdel rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, pages 6645–6649. IEEE, 2013.
- [10] Alex Graves, Santiago FernÃandez, and JÃijrgen Schmidhuber. Multi-dimensional recurrent neural networks. In *ICANN*, pages 549–558, 2007.
- [11] Karol Gregor, Ivo Danihelka, Alex Graves, and Daan Wierstra. DRAW: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- [12] Andrew Kae, Kihyuk Sohn, Honglak Lee, and Erik Learned-Miller. Augmenting CRFs with Boltzmann machine shape priors for image labeling. In *CVPR*, 2013.
- [13] Nal Kalchbrenner, Ivo Danihelka, and Alex Graves. Grid long short-term memory. *arXiv preprint arXiv:1507.01526*, 2015.
- [14] Xiaodan Liang, Xiaohui Shen, Donglai Xiang, Jiashi Feng, Liang Lin, and Shuicheng Yan. Semantic object parsing with local-global long short-term memory. arXiv preprint arXiv:1511.04510, 2015.
- [15] Ce Liu, Jenny Yuen, and Antonio Torralba. Nonparametric scene parsing via label transfer. PAMI, 2011.
- [16] Sifei Liu, Jimei Yang, Chang Huang, and Ming-Hsuan Yang. Multi-objective convolutional learning for face labeling. In CVPR, 2015.
- [17] Sifei Liu, Jinshan Pan, and Ming-Hsuan Yang. Learning recursive filters for low-level vision via a hybrid neural network. In *ECCV*, 2016.

#### 12 S. LIU, J. SHI, J. LIANG, M.H. YANG: FACE PARSING VIA RECURRENT PROPAGATION

- [18] Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015.
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *arXiv preprint arXiv:1411.4038*, 2014.
- [20] Ping Luo, Xiaogang Wang, and Xiaoou Tang. Hierarchical face parsing via deep learning. In CVPR. IEEE, 2012.
- [21] Hongwei Qin, Junjie Yan, Xiu Li, and Xiaolin Hu. Joint training of cascaded cnn for face detection. In CVPR, 2016.
- [22] Hannes Schulz and Sven Behnke. Learning object-class segmentation with convolutional neural networks. In *ESANN*, volume 3, page 1, 2012.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [24] Brandon Smith, Li Zhang, Jonathan Brandt, Zhe Lin, and Jianchao Yang. Exemplar-based face parsing. In CVPR, 2013.
- [25] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep convolutional network cascade for facial point detection. In CVPR, 2013.
- [26] Stavros Tsogkas, Iasonas Kokkinos, George Papandreou, and Andrea Vedaldi. Semantic part segmentation with deep learning. *arXiv preprint arXiv:1505.02438*, 2015.
- [27] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. arXiv preprint arXiv:1601.06759, 2016.
- [28] Francesco Visin, Kyle Kastner, Kyunghyun Cho, Matteo Matteucci, Aaron Courville, and Yoshua Bengio. Renet: A recurrent neural network based alternative to convolutional networks. arXiv preprint arXiv:1505.00393, 2015.
- [29] Ronald J Williams and David Zipser. Gradient-based learning algorithms for recurrent networks and their computational complexity. *Back-propagation: Theory, architectures and applications*, pages 433–486, 1995.
- [30] Takayoshi Yamashita, Takaya Nakamura, Hiroshi Fukui, Yuji Yamauchi, and Hironobu Fujiyoshi. Cost-alleviative learning for deep convolutional neural network-based facial part labeling. *IPSJ Transactions on Computer Vision and Applications*, 2015. doi: 10.2197/ipsjtcva.7.99.
- [31] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015.
- [32] Zhen Zuo, Bing Shuai, Gang Wang, Xiao Liu, Xingxing Wang, Bing Wang, and Yushi Chen. Convolutional recurrent neural networks: Learning spatial dependencies for image representation. In *CVPR Workshops*, 2015.