

Deep Learning Markov Random Field for Semantic Segmentation

Ziwei Liu*, Xiaoxiao Li*, Ping Luo, *Member, IEEE*, Chen Change Loy, *Senior Member, IEEE*, and Xiaoou Tang, *Fellow, IEEE*

Abstract—Semantic segmentation tasks can be well modeled by Markov Random Field (MRF). This paper addresses semantic segmentation by incorporating high-order relations and mixture of label contexts into MRF. Unlike previous works that optimized MRFs using iterative algorithm, we solve MRF by proposing a Convolutional Neural Network (CNN), namely Deep Parsing Network (DPN), which enables deterministic end-to-end computation in a single forward pass. Specifically, DPN extends a contemporary CNN to model unary terms and additional layers are devised to approximate the mean field (MF) algorithm for pairwise terms. It has several appealing properties. First, different from the recent works that required many iterations of MF during back-propagation, DPN is able to achieve high performance by approximating one iteration of MF. Second, DPN represents various types of pairwise terms, making many existing models as its special cases. Furthermore, pairwise terms in DPN provide a unified framework to encode rich contextual information in high-dimensional data, such as images and videos. Third, DPN makes MF easier to be parallelized and speeded up, thus enabling efficient inference. DPN is thoroughly evaluated on standard semantic image/video segmentation benchmarks, where a single DPN model yields state-of-the-art segmentation accuracies on PASCAL VOC 2012, Cityscapes dataset and CamVid dataset.

Index Terms—Semantic Image/Video Segmentation, Markov Random Field, Convolutional Neural Network.

1 INTRODUCTION

SEMANTIC segmentation is a fundamental and long-standing problem in computer vision. It is defined as a multi-label classification problem, aiming to assign each pixel with a category label. There are two widely adopted research realms, including semantic image segmentation [1], [2], [3], [4], [5], [6], [7] and semantic video segmentation [8], [9], [10], [11], [12], [13], [14]. The former employs a static image as input, while the latter employs a video sequence. The obtained per-pixel segmentation results are extremely useful for several applications like smart editing [15], scene understanding [4] and automated driving [16].

Since pixels in natural images or videos generally exhibit strong correlation, jointly modeling label distribution in all locations is desirable. To capture these contextual information, Markov random field (MRF) and conditional random field (CRF) [17] are commonly used as classic frameworks for semantic segmentation. They model the joint distribution of labels by defining both unary term and pairwise terms. Unary term reflects the per-pixel confidence of assigning labels while pairwise terms capture the inter-pixel constraints.

Most previous studies focus on designing pairwise terms that possess strong expressive power. For example, Krähenbühl *et al.* [18] attained accurate segmentation boundary by inferring on a fully-connected graph. Vineet *et al.* [19] extended [18] by defining both high-order and long-range terms between pixels. Global or local semantic contexts between labels were also investigated by [20]. However, their performance are limited by the relatively shallow models (*e.g.* SVM or Adaboost) used as unary term. As deep learning gradually takes over in many image recognition

fields [21], researchers have also explored the possibility of designing effective deep architecture for semantic segmentation. For instance, Long *et al.* [6] transformed fully-connected layers of CNN into convolutional layers, making accurate per-pixel classification possible using contemporary CNN architectures that were pre-trained on ImageNet [22]. Chen *et al.* [23] improved [6] by feeding the outputs of CNN into a MRF with simple pairwise potentials, but it treated CNN and MRF as separate components. A recent advance was made in joint training CNN and MRF by passing the error of MRF inference backward into CNN [24]. Nonetheless, an iterative inference of MRF such as the mean field algorithm (MF) [25] is required for each training image during the back-propagation (BP). Zheng *et al.* [26] further showed that the procedure of MF inference can be represented as a Recurrent Neural Network (RNN), but their computational costs are similar to that of [24].

We observed that a direct combination of CNN and MRF as above is inefficient, posing challenges on both optimization difficulty and inference speed. Since CNN typically has millions of parameters and MRF typically has thousands of latent variables, they are cumbersome to jointly optimize and infer. Even worse, incorporating complex pairwise terms into a MRF becomes impractical, limiting the performance of the entire system. In this study, we propose a novel Deep Parsing Network (DPN), which is an end-to-end system enabling jointly training of CNN and complex pairwise terms. DPN has several appealing **properties**:

(1) DPN solves MRF with a single feed-forward pass, reducing computational cost and meanwhile maintaining high performance. Specifically, DPN models unary terms by extending the VGG-16 network (VGG₁₆) [21] pre-trained on ImageNet, while additional layers are carefully designed to model complex pairwise terms. The learning of these terms is transformed into deterministic end-to-end computation by BP, instead of embedding MF into BP as

• Z. Liu, X. Li, P. Luo, C. C. Loy and X. Tang are with the Department of Information Engineering, The Chinese University of Hong Kong.
E-mail: {l:013, lx015, pluo, ccloy, xtang}@ie.cuhk.edu.hk

• * The first two authors share first-authorship. Correspondence to: Ping Luo and Ziwei Liu.

[24], [27] did. Although MF can be represented by a RNN [26], it needs to recurrently compute the forward pass so as to achieve good performance and thus the process is time-consuming, *e.g.* each forward pass contains hundred thousand of weights. DPN approximates MF by using only one iteration of inference. This is made possible by joint learning strong unary terms and rich pairwise information.

(2) Pairwise terms determine the graphical structure. In previous studies, if the former is changed, so is the latter as well as its inference procedure. But with DPN, modifying the complexity of pairwise terms, *e.g.* range of pixels and contexts, is as simple as modifying the receptive fields of convolutions, without varying BP. Furthermore, DPN is capable of representing multiple types of pairwise terms, making many previous works [23], [26], [24] as its special cases.

(3) DPN approximates MF with convolutional and pooling operations, which can be speeded up by low-rank approximation [28] and easily parallelized [29] in a Graphical Processing Unit (GPU).

Our **contributions** are summarized as below. (1) We propose the novel DPN to jointly train VGG₁₆ for unary terms with rich pairwise information, *i.e.* mixture of label contexts and high-order relations. In comparison to existing deep models, DPN approximates MF with only *one iteration* of inference, reducing computational cost but still maintaining high performance. (2) We show that multiple types of MRFs can be represented in DPN, making many previous works such as RNN [26] and DeepLab [23] as its special cases. (3) We conduct extensive experiments to investigate which component of DPN is crucial to achieve high performance. We demonstrate the generalizability of DPN model by showing its state-of-the-art performance on several standard semantic image/video segmentation benchmarks, including PASCAL VOC 2012 [30], CityScapes dataset [16] and CamVid dataset [9].

In comparison to our earlier version of this work [7], we propose a generic deep learning framework, Deep Parsing Network (DPN) to model and solve N -Dimension (N -D) high-order Markov Random Field (MRF). Our previous study [7] only shows the possibility on 2-Dimension image segmentation problem. Specifically, we employ dynamic node linking to construct graph in N -D space, which results in a model of N -D high-order MRF. To solve this high-dimensional and high-order MRF, we reformulate the mean field (MF) update process into a feed-forward pass of Convolutional Neural Network (CNN). N -D local and global convolutional layers are designed to approximate different terms in a MF solver. Apart from the methodology, the paper was also substantially improved by providing more technical details and more extensive experimental evaluations.

2 RELATED WORK

Existing studies [1], [31], [2], [32], [33], [34], [4], [5], [6], [35], [36] on semantic segmentation focus on either constructing specific graph structure so that contextual information and long-term dependencies can be captured, or designing suitable network architecture to leverage the power of deep learning. In the following, we summarize recent research advances with respect to these two aspects.

Markov Random Field. Markov Random Field (MRF) or Conditional Random Field (CRF) has achieved great successes in semantic image segmentation, which is one of the most

challenging problems in computer vision. Researchers improved labeling accuracy by exploring rich information to define the pairwise functions, including long-range dependencies [18], [37], high-order potentials [19], [38], and semantic label contexts [3], [39], [20]. For example, Krähenbühl *et al.* [18] attained accurate segmentation boundary by inferring on a fully-connected graph. Vineet *et al.* [19] extended [18] by defining both high-order and long-range terms between pixels. Global or local semantic contexts between labels were also investigated by [20]. Although they accomplished promising results, they modeled the unary terms as SVM or Adaboost, whose learning capacity becomes a bottleneck. The learning and inference of complex pairwise terms are often expensive.

MRF and CRF have also been utilized in semantic video segmentation by extending their graph structure to spatio-temporal domain. For example, Wang *et al.* [40] unified foreground object segmentation, tracking and occlusion reasoning into a carefully designed MRF model. Optical flow based long-term trajectories [10] were also exploited to discover moving objects. Liu *et al.* [12] employed fully-connected CRF augmented with object potentials for efficient multi-class inference. However, these methods are based on hand-crafted features, thus lacking sufficient learning capacity.

Convolutional Neural Network. More recently, Convolutional Neural Network (CNN) has been leveraged as a strong unary classifier. With deep models, existing works [41], [42], [5], [6], [23], [43], [26], [24], [27] demonstrated encouraging segmentation results through using just simple definition of the pairwise function or even neglecting it. For instance, Long *et al.* [6] transformed fully-connected layers of CNN into convolutional layers, making accurate per-pixel classification possible using the contemporary CNN architectures that were pre-trained on ImageNet [22]. Chen *et al.* [23] improved [6] by feeding the outputs of CNN into a MRF with simple pairwise potentials, but it treated CNN and MRF as separated components. A recent advance was obtained by [24], which jointly trained CNN and MRF by passing the error of MRF inference backward into CNN, but iterative inference of MRF such as the mean field algorithm (MF) [25] is required for each training image during the back-propagation (BP). Zheng *et al.* [26] further showed that the procedure of MF inference can be represented as a Recurrent Neural Network (RNN), but their computational costs are similar to that of [24].

Little attempts have been made to develop unified deep learning framework for semantic video segmentation. Recent efforts in this direction include SegNet [44], which adopted an encoder-decoder architecture but did not take temporal relationships into consideration. Here we extend DPN to further include temporal voxels into the joint learning and inference process, which results in an end-to-end trainable system with rich spatio-temporal information encoded.

3 OUR APPROACH

We develop a unified framework, DPN, for modeling and solving high-order MRF. The architecture of DPN is shown in Fig. 1(a). Our model imposes no restrictions on the dimension of MRF. For example, it can be either 2-Dimension (2-D) for images segmentation, 3-Dimension (3-D) for video segmentation, or N -Dimension (N -D) for sensor data. DPN learns MRF by extending VGG₁₆ to model unary terms and additional layers are carefully designed to model pairwise terms. In the following, we describe

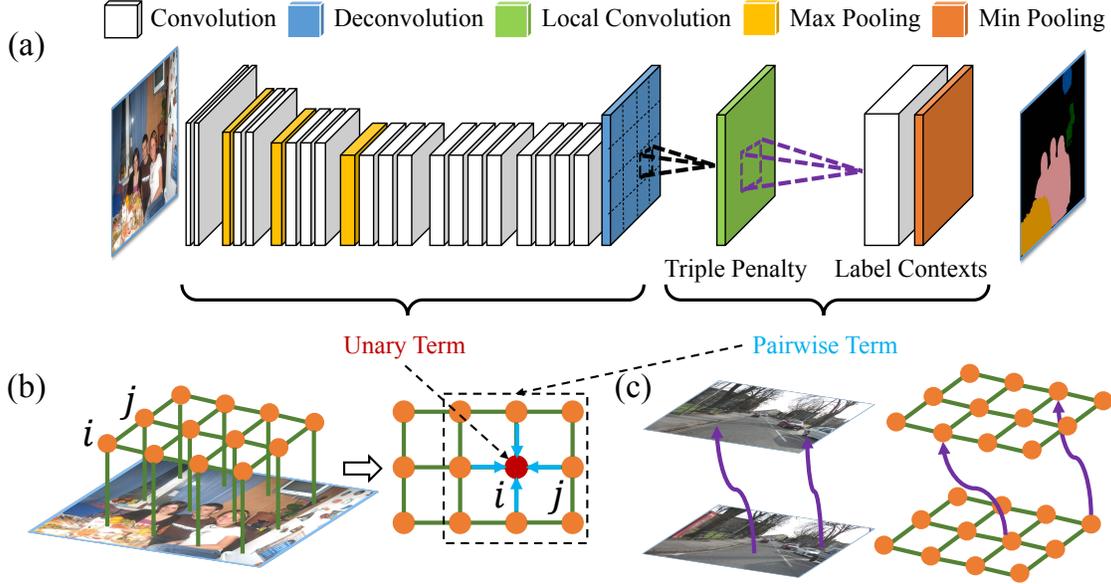


Fig. 1: (a) The network architecture of a deep parsing network (DPN). (b) DPN extends a contemporary CNN architecture to model unary terms and additional layers are carefully devised to approximate the mean field algorithm (MF) for pairwise terms. (c) DPN enables dynamic linking of nodes in Markov Random Field (MRF) by incorporating domain knowledge.

the formulation of DPN in terms of 3-D MRF. Note that it can be easily resorted to 2-D cases or extended to N -D cases by removing or adding relationships between nodes.

Markov Random Field. MRF [45] is an undirected graph where each node represents a voxel in a video, \mathbf{I} , and each edge represents relation between voxels, as shown in Fig. 1(b). Each node is associated with a binary latent variable, $y_i^u \in \{0, 1\}$, indicating whether a voxel $\mathbf{i} = [i \ t_i]$ has label u . Here, i indicates a voxel’s spatial index with respect to an image, and t_i is its temporal index with respect to a sequence. We have $\forall u \in L = \{1, 2, \dots, l\}$, representing a set of l labels. The energy function of MRF is written as

$$E(\mathbf{y}) = \sum_{\forall \mathbf{i} \in \mathcal{V}} \Phi(y_i^u) + \sum_{\forall (\mathbf{i}, \mathbf{j}) \in \mathcal{E}} \Psi(y_i^u, y_j^v), \quad (1)$$

where \mathbf{y} , \mathcal{V} , and \mathcal{E} denote a set of latent variables, nodes, and edges, respectively. $\Phi(y_i^u)$ is the unary term, measuring the cost of assigning label u to voxel \mathbf{i} . For instance, if voxel \mathbf{i} belongs to the first category other than the second one at time t_i , we should have $\Phi(y_i^1) < \Phi(y_i^2)$. Moreover, $\Psi(y_i^u, y_j^v)$ is the pairwise term that measures the penalty of assigning labels u, v to a pair of voxel (\mathbf{i}, \mathbf{j}) respectively.

Dynamic Node Linking. Traditional approaches [13], [46] usually define the edges \mathcal{E} on rectangular grid in 3-D space. However, when large motion exists, the actual temporal trajectory for certain pixel will not reside inside a rigid cube, which means the rectangular grid assumption does not hold. To better preserve the contextual information in a spatio-temporal space, we employ dynamic node linking to construct edges \mathcal{E} in DPN. Specifically, we keep the 2-D structure in the spatial domain, as illustrated in Fig. 1(c). In the temporal domain, the neighboring voxels $\mathbf{i} = [i \ t_i]$ and $\mathbf{j} = [j \ t_j]$ are defined as those lie on the same temporal trajectories $\Delta_{\mathbf{i} \rightarrow \mathbf{j}}$. This trajectory can be estimated by standard optical flow techniques [47]. The formulation of edges \mathcal{E}_t in the temporal domain is

$$(\mathbf{i}, \mathbf{j}) \in \mathcal{E}_t \iff \mathbf{j} = \mathbf{i} + \Delta_{\mathbf{i} \rightarrow \mathbf{j}}. \quad (2)$$

In this setting, adjacent nodes in the temporal space would be more likely to belong to the same category, making the label contexts more easily to be captured.

Unary and Pairwise Terms. Intuitively, the unary terms represent per-voxel classifications, while the pairwise terms represent a set of smoothness constraints. The unary term in Eqn. (1) is typically defined as

$$\Phi(y_i^u) = -\ln p(y_i^u = 1 | \mathbf{I}), \quad (3)$$

where $p(y_i^u = 1 | \mathbf{I})$ indicates the probability of the presence of label u at voxel \mathbf{i} , modeling by VGG₁₆. To simplify discussions, we abbreviate it as p_i^u . The smoothness term can be formulated as

$$\Psi(y_i^u, y_j^v) = \mu(u, v) d(\mathbf{i}, \mathbf{j}), \quad (4)$$

where the first term learns the penalty of global co-occurrence between any pair of labels. For example, the output value of $\mu(u, v)$ is large if u and v should not coexist. In Eqn.(4), the second term calculates the distances between voxels. We have

$$d(\mathbf{i}, \mathbf{j}) = d(i, t_i, j, t_j) = \omega_1 \|\mathbf{I}_i - \mathbf{I}_j\|^2 + \omega_2 \|[i \ t_i] - [j \ t_j]\|^2, \quad (5)$$

where \mathbf{I}_i indicates a feature vector such as RGB values extracted from the input video for voxel \mathbf{i} , $[i \ t_i]$ denote coordinates of voxels’ positions, and ω_1, ω_2 are the constant weights. Eqn. (4) implies that if two voxels are close and look similar, they are encouraged to have labels that are compatible. This formulation has been adopted by most of the recent deep models [23], [26], [24] for semantic image segmentation.

However, Eqn. (4) has two main drawbacks. First, its first term captures the co-occurrence frequency of two labels in the training data, but neglects the spatial context between objects. For example, ‘person’ may appear beside ‘table’, but not at its bottom. This spatial context is a mixture of patterns, as different object configurations may appear in different images, such as ‘person’ standing beside ‘table’ and ‘person’ sitting behind ‘table’. Second, it defines only the pairwise relations between pixels, missing their high-order interactions.

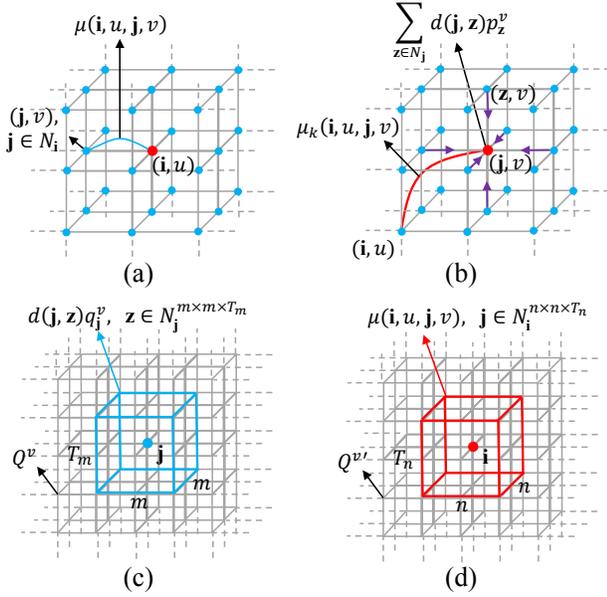


Fig. 2: (a) Illustration of the pairwise terms in DPN. (b) explains the label contexts. (c) and (d) show that the mean field update of DPN corresponds to convolutions.

To resolve these issues, we define the smoothness term by leveraging rich information between voxels, which is one of the **advantages** of DPN over existing deep models. We have

$$\Psi(y_i^u, y_j^v) = \sum_{k=1}^K \lambda_k \mu_k(i, u, j, v) \sum_{z \in \mathcal{N}_j} d(j, z) p_z^v. \quad (6)$$

The first term in Eqn. (6) learns a **mixture of local label contexts**, penalizing label assignment in a local cube, where K is the number of components in the mixture and λ_k is an indicator, determining which component is activated. We define $\lambda_k \in \{0, 1\}$ and $\sum_{k=1}^K \lambda_k = 1$. An intuitive illustration is given in Fig. 2(a), where the dots in red and blue represent a center voxel i and its neighboring voxels j , *i.e.* $j \in \mathcal{N}_i$, and (i, u) indicates assigning label u to voxel i . Here, $\mu(i, u, j, v)$ outputs labeling cost between (i, u) and (j, v) with respect to their relative positions. For instance, if u, v represent ‘person’ and ‘table’, the learned penalties of positions j that are at the bottom of voxel i should be large. The second term of Eqn. (6) basically models a **triple penalty**, which involves voxels i, j , and j ’s neighbors, implying that if (i, u) and (j, v) are compatible, then (i, u) should be also compatible with j ’s nearby pixels (z, v) , $\forall z \in \mathcal{N}_j$, as shown by the purple arrows in Fig. 2(b).

Learning parameters (*i.e.* weights of VGG₁₆ and costs of label contexts) in Eqn. (1) requires us to minimize the distances between ground-truth label map and the predicted label \mathbf{y} , which needs to be inferred subject to the smoothness constraints.

Inference Overview. Inference of Eqn. (1) can be obtained by the mean field (MF) algorithm [25], which estimates the joint distribution of MRF

$$P(\mathbf{y}) = \frac{1}{Z} \exp\{-E(\mathbf{y})\}, \quad (7)$$

by using a fully-factorized proposal distribution

$$Q(\mathbf{y}) = \prod_{i \in \mathcal{V}} \prod_{u \in \mathcal{L}} q_i^u, \quad (8)$$

where each q_i^u is a variable we need to estimate, indicating the predicted probability of assigning label u to voxel i . The Kullback-Leibler divergence between them is then calculated as

$$\begin{aligned} D_{KL}(Q \| P) &= \sum_{\mathbf{y}} Q(\mathbf{y}) \ln \left(\frac{Q(\mathbf{y})}{P(\mathbf{y})} \right) \\ &= \sum_{\mathbf{y}} Q(\mathbf{y}) E(\mathbf{y}) + \sum_{\mathbf{y}} Q(\mathbf{y}) \ln Q(\mathbf{y}) + \ln Z. \end{aligned} \quad (9)$$

Since $\ln Z$ is a constant, minimizing the Kullback-Leibler divergence between $Q(\mathbf{y})$ and $P(\mathbf{y})$ is equivalent to minimizing the former terms in Eqn. (9), which is denoted as free energy $F(Q)$ [17]. To simplify the discussion, we denote $\Phi(y_i^u)$ and $\Psi(y_i^u, y_j^v)$ as Φ_i^u and Ψ_{ij}^{uv} , respectively. And we can further substitute Eqn.(1) into $F(Q)$. We have

$$\begin{aligned} F(Q) &= \sum_{\mathbf{y}} Q(\mathbf{y}) E(\mathbf{y}) + \sum_{\mathbf{y}} Q(\mathbf{y}) \ln Q(\mathbf{y}) \\ &= \sum_{\forall i \in \mathcal{V} \forall u \in \mathcal{L}} q_i^u \Phi_i^u + \sum_{\forall i, j \in \mathcal{E} \forall u \in \mathcal{L} \forall v \in \mathcal{L}} q_i^u q_j^v \Psi_{ij}^{uv} \\ &\quad + \sum_{\forall i \in \mathcal{V} \forall u \in \mathcal{L}} q_i^u \ln q_i^u. \end{aligned} \quad (10)$$

Specifically, the first term in Eqn. (10) characterizes the cost of each voxel’s predictions, while the second term characterizes the consistencies of predictions between voxels. The last term denotes the entropy, measuring the confidences of predictions. Then, a constrained optimization problem regarding q_i^u could be formulated as

$$\begin{aligned} &\text{minimize} && F(Q) \\ & && q_i^u \\ &\text{subject to} && \sum_u q_i^u = 1, \forall i \in \mathcal{V}. \end{aligned} \quad (11)$$

To solve this minimization problem, we define $J(Q) = F(Q) + \sum_i \lambda_i (\sum_u q_i^u - 1)$ by introducing Lagrange multipliers λ_i . The final closed-form solution can be obtained by differentiating $J(Q)$ w.r.t. to q_i^u and equating the resulting expression to zero

$$q_i^u \propto \exp \left\{ - \left(\Phi_i^u + \sum_{\forall j \in \mathcal{N}_i \forall v \in \mathcal{L}} q_j^v \Psi_{ij}^{uv} \right) \right\}, \quad (12)$$

such that the predictions for each voxel is independently attained by repeatedly calculating Eqn. (12), which implies whether voxel i have label u is proportional to the estimated probabilities of all its neighboring voxels, weighted by their corresponding smoothness penalties. Substituting Eqn. (6) into Eqn. (12), we have

$$\begin{aligned} q_i^u &\propto \exp \left\{ - \Phi_i^u - \sum_{k=1}^K \lambda_k \sum_{\forall v \in \mathcal{L}} \right. \\ &\quad \left. \sum_{\forall j \in \mathcal{N}_i} \mu_k(i, u, j, v) \sum_{\forall z \in \mathcal{N}_j} d(j, z) q_j^v q_z^v \right\}, \end{aligned} \quad (13)$$

where each q_i^u is initialized by the corresponding p_i^u in Eqn. (3), which is the unary prediction of VGG₁₆. Eqn. (13) satisfies the smoothness constraints.

In the following, DPN approximates one iteration of Eqn. (13) by decomposing it into two steps. Let Q^v be a predicted label map of the v -th category. In the first step as shown in Fig. 2(c), we calculate the triple penalty term in Eqn. (13) by applying a $m \times m \times T_m$ filter on each position j , where each element of this filter equals $d(j, z) q_j^v q_z^v$, resulting in $Q^{v'}$. T_m indicates the time

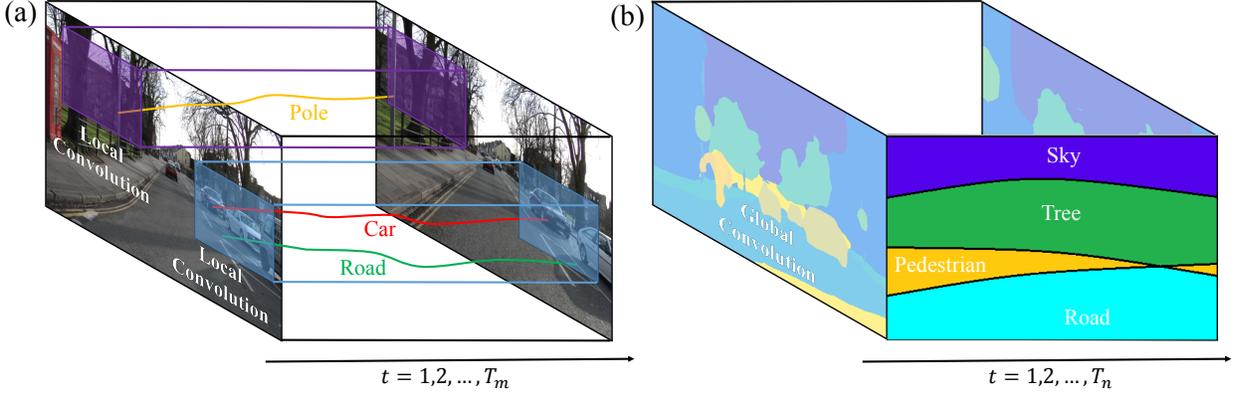


Fig. 3: Illustrative depiction of (a) triple penalty term and (b) mixture of local label contexts term.

span. Apparently, this step smoothes the prediction of voxel \mathbf{j} with respect to the distances between it and its neighborhood. In the second step as illustrated in Fig. 2(d), the labeling contexts can be obtained by convolving $Q^{v'}$ with a $n \times n \times T_n$ filter, each element of which equals $\mu_k(\mathbf{i}, u, \mathbf{j}, v)$, penalizing the triple relations as shown in Fig. 2(a).

Fig. 3 depicts the semantic meaning of the **triple penalty** and the **mixture of local label contexts** term in the spatial-temporal domain. From Fig. 3(a) we can see that the triple penalty term tracks the movement of local pixels, such as ‘pole’, ‘car’ and ‘road’. These temporal trajectories combined with local regions are subsequently used to smooth original predictions. Fig. 3(b) demonstrates that a mixture of local label contexts term not only captures label co-occurrence in a single image, but also encodes the change of label configurations along time. For example, as the observer vehicle drives, ‘tree’ will move backward and get nearer to ‘road’.

4 DEEP PARSING NETWORK

This section describes the implementation of Eqn. (13) in a Deep Parsing Network (DPN). DPN extends VGG₁₆ to model the unary term and with additional layers to approximate one iteration of MF inference as the pairwise term. The hyper-parameters of VGG₁₆ and DPN are compared in Table 1.

As listed in Table 1, the first row represents the *name* of layer and ‘ x - y ’ in the second row represents the *size* of the receptive field and the *stride* of convolution, respectively. For instance, ‘3-1’ in the convolutional layer implies that the receptive field of each filter is 3×3 and it is applied on every single pixel of an input feature map, while ‘2-2’ in the max-pooling layer indicates each feature map is pooled over every other pixel within a 2×2 local region. On the other hand, ‘3-50-1’ implies a 3D convolution of size $50 \times 50 \times 3$, where ‘3’ suggests that this 3D filter is applied on three consecutive frames. The last three rows show the number of the output feature maps, activation functions, and the size of output feature maps, respectively. T represents the number of frames in the underlying video sequence.

As summarized in Table 1(a), VGG₁₆ contains thirteen convolutional layers, five max-pooling layers, and three fully-connected layers. These layers can be partitioned into twelve groups, each of which covers one or more homogeneous layers. For example, the first group comprises two convolutional layers with 3×3 receptive field and 64 output feature maps, each of which is 224×224 .

4.1 Modeling Unary Terms

To make full use of VGG₁₆, which is pre-trained by ImageNet, we adopt all its parameters to initialize the filters of the first ten groups of DPN. To simplify the discussions, we take PASCAL VOC 2012 (VOC12) [30] as an example. Note that DPN can be easily adapted to any other semantic image segmentation dataset by modifying its hyper-parameters. VOC12 contains 21 categories and each image is rescaled to 512×512 in training. Therefore, DPN needs to predict a total of $512 \times 512 \times 21$ labels, *i.e.* one label for each pixel. To this end, we extend VGG₁₆ in two aspects.

In particular, let a_i and b_i denote the i -th group in Table 1(a) and (b), respectively. First, we **increase the resolution** of VGG₁₆ by removing its max pooling layers at a8 and a10, since most of the information is lost after pooling, *e.g.* a10 reduces the input size by 32 times, *i.e.* from 224×224 to 7×7 . As a result, the smallest size of feature map in DPN is 64×64 , keeping much more information compared with VGG₁₆. Note that the filters of b8 are initialized as the filters of a9, but the 3×3 receptive field is padded into 5×5 as shown in Fig. 4(a), where the cells in white are the original values of the a9’s filter and the cells in gray are zeros. This step is performed because a8 is not presented in DPN, therefore each filter in a9 should be convolved on every other pixel of a7. To maintain the convolution with one stride, we pad the filters with zeros. Furthermore, the feature maps in b11 are up-sampled to 512×512 by bilinear interpolation. Since DPN is trained with label maps of the entire images, the missing information in the preceding layers of b11 can be recovered by BP. The supervision signals in the interpolated pixels will guide the feature learning for full-resolution images.

Second, two fully-connected layers at a11 are transformed into two convolutional layers at b9 and b10, respectively. As shown in Table 1(a), the first ‘fc’ layer learns $7 \times 7 \times 512 \times 4096$ parameters, which can be altered to 4096 filters in b9, each of which is $25 \times 25 \times 512$. Since a8 and a10 have been removed, the 7×7 receptive field is padded into 25×25 similar as above and shown in Fig. 4 (b). The second ‘fc’ layer learns a 4096×4096 weight matrix, corresponding to 4096 filters in b10. Each filter is $1 \times 1 \times 4096$.

Overall, b11 generates the unary labeling results, producing twenty-one 512×512 feature maps, each of which represents the probabilistic label map of each category.

4.2 Modeling Smoothness Terms

The last four layers of DPN, *i.e.* from b12 to b15, are carefully designed to smooth the unary labeling results.

TABLE 1: A comparison between the network architectures of VGG₁₆ and DPN.

(a) VGG ₁₆ : 224×224×3 input image; 1×1000 output labels													
	1	2	3	4	5	6	7	8	9	10	11	12	
<i>layer</i>	2×conv	max	2×conv	max	3×conv	max	3×conv	max	3×conv	max	2×fc	fc	
<i>fi.-st.</i>	3-1	2-2	3-1	2-2	3-1	2-2	3-1	2-2	3-1	2-2	-	-	
<i>#ch.</i>	64	64	128	128	256	256	512	512	512	512	1	1	
<i>act.</i>	relu	idn	relu	soft									
<i>size</i>	224	112	112	56	56	28	28	14	14	7	4096	1000	

(b) DPN: T×512×512×3 input image; T×512×512×L output label maps															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
<i>layer</i>	2×conv	max	2×conv	max	3×conv	max	3×conv	3×conv	conv	conv	conv	lconv-3D	conv-3D	bmin	sum
<i>fi.-st.</i>	3-1	2-2	3-1	2-2	3-1	2-2	3-1	5-1	25-1	1-1	1-1	3-50-1	3-9-1	1-1	1-1
<i>#ch.</i>	64	64	128	128	256	256	512	512	4096	4096	L	L	L×k	L	L
<i>act.</i>	relu	idn	relu	idn	relu	idn	relu	relu	relu	relu	sigm	lin	lin	idn	soft
<i>size</i>	T-512	T-256	T-256	T-128	T-128	T-64	T-64	T-64	T-64	T-64	T-512	T-512	T-512	T-512	T-512

As shown in (a) and (b) respectively. Each table contains five rows, ‘layer’, ‘fi.-st.’, ‘#ch.’, ‘act.’, and ‘size’ represent the ‘name of layer’, ‘receptive field of filter’–‘stride’, ‘number of output feature maps’, ‘activation function’, and ‘size of output feature maps’, respectively. Furthermore, ‘conv’, ‘lconv-3D’, ‘conv-3D’, ‘max’, ‘bmin’, ‘fc’, and ‘sum’ represent the convolution, 3D local convolution, max pooling, block min pooling, fully connection, and summation, respectively. Moreover, ‘relu’, ‘idn’, ‘soft’, ‘sigm’, and ‘lin’ represent the activation functions, including rectified linear unit [48], identity, softmax, sigmoid, and linear, respectively. T , L , and k represent the length of frames, number of categories, and number of mixture filters.

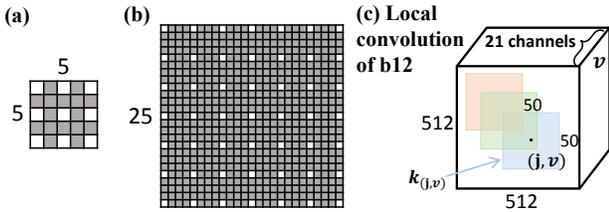


Fig. 4: (a) and (b) show the padding of the filters. (c) illustrates local convolution of b12.

• **b12.** As listed in Table 1 (b), ‘lconv’ in b12 indicates a **3D locally convolutional layer**. A counterpart of it (*i.e.* 2D locally convolutional layer) is widely used in face recognition [49], [50], [51] to capture different information from different facial positions. Similarly, distinct spatial positions of b12 have different filters, and each filter is shared across 21 input channels, as shown in Fig. 4(c). It can be formulated as

$$\mathbf{o}_{(j,v)}^{12} = \text{lin}(\mathbf{k}_{(j,v)} * \mathbf{o}_{(j,v)}^{11}), \quad (14)$$

where $\text{lin}(x) = ax + b$ representing the linear activation function, ‘*’ is the convolutional operator, and $\mathbf{k}_{(j,v)}$ is a $50 \times 50 \times 3 \times 1$ filter at position \mathbf{j} of channel v . The choice and effect of filter size will be discussed in the experiments (Sec.5.1). We have $\mathbf{k}_{(j,1)} = \mathbf{k}_{(j,2)} = \dots = \mathbf{k}_{(j,21)}$ shared across 21 channels. $\mathbf{o}_{(j,v)}^{11}$ indicates a local cube in b11, while $\mathbf{o}_{(j,v)}^{12}$ is the corresponding output of b12. Since b12 has a stride of one, the result of $\mathbf{k}_{\mathbf{j}} * \mathbf{o}_{(j,v)}^{11}$ is scalar. In summary, b12 has 512×512 different filters and produces 21 output feature maps.

Eqn. (14) implements the **triple penalty** of Eqn. (13). Recall that each output feature map of b11 indicates a probabilistic label map of a specific object appearing in the frame. As a result, Eqn. (14) suggests that the probability of object v presented at position \mathbf{j} is updated by weighted averaging over the probabilities at its nearby positions. Thus, as shown in Fig. 2(c), $\mathbf{o}_{(j,v)}^{11}$ corresponds to a cube of Q^v centered at \mathbf{j} , which has values $p_{\mathbf{z}}^v$, $\forall \mathbf{z} \in \mathcal{N}_{\mathbf{j}}^{50 \times 50 \times 3}$. Similarly, $\mathbf{k}_{(j,v)}$ is initialized by $d(\mathbf{j}, \mathbf{z}) p_{\mathbf{z}}^v$, implying each filter captures dissimilarities between positions. These filters remain fixed during BP, other than learned as in

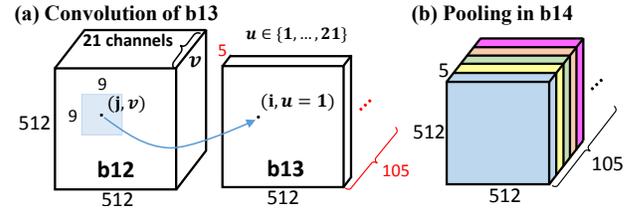


Fig. 5: (a) and (b) illustrates the convolutions of b13 and the poolings in b14.

conventional CNN¹.

• **b13.** As shown in Table 1(b) and Fig. 5(a), b13 is a **3D global convolutional layer** that generates 105 feature maps by using 105 filters of size $9 \times 9 \times 3 \times 21$. For example, the value of $(\mathbf{i}, u = 1)$ is attained by applying a $9 \times 9 \times 3 \times 21$ filter at positions $\{(\mathbf{j}, v = 1, \dots, 21)\}$. In other words, b13 learns a filter for each category to penalize the probabilistic label maps of b12, corresponding to the **local label contexts** in Eqn.(13) by assuming $K = 5$ and $n = 9$, as shown in Fig.2 (d).

• **b14.** As illustrated in Table 1 and Fig. 5(b), b14 is a block min pooling layer that pools over every 1×1 region with one stride across every 5 input channels, leading to 21 output channels, *i.e.* $105 \div 5 = 21$. Layer b14 activates the contextual pattern with the smallest penalty.

• **b15.** This layer combines both the unary and smoothness terms by summing the outputs of b11 and b14 in an element-wise manner similar to Eqn. (13),

$$\mathbf{o}_{(i,u)}^{15} = \frac{\exp \{ \ln(\mathbf{o}_{(i,u)}^{11}) - \mathbf{o}_{(i,u)}^{14} \}}{\sum_{u=1}^{21} \exp \{ \ln(\mathbf{o}_{(i,u)}^{11}) - \mathbf{o}_{(i,u)}^{14} \}}, \quad (15)$$

where probability of assigning label u to voxel \mathbf{i} is normalized over all the labels.

Relation to Previous Deep Models. Many existing deep models such as [26], [23], [24] employed Eqn. (4) as the pairwise terms,

1. Each filter in b12 actually represents a distance metric between pixels in a specific region. In VOC12, the patterns of all the training images in a specific region are heterogeneous, because of various object shapes. Therefore, we initialize each filter with Euclidean distance. Nevertheless, Eqn. (14) is a more general form than the triple penalty in Eqn. (13), *i.e.* filters in Eqn. (14) can be automatically learned from data, if the patterns in a specific region are homogeneous, such as face or human images, which have more regular shapes than images in VOC12.

which are the special cases of Eqn. (13). To see this, let $K = 1$, $j = i$ and omit t (i.e. $t_z = t_j = t_i$), the right hand side of Eqn. (13) reduces to

$$\begin{aligned} & \exp\{-\Phi_i^u - \sum_{v \in L} \lambda_1 \mu_1(i, u, i, v) \sum_{z \in \mathcal{N}_i} d(i, z) p_i^v p_z^v\} \\ &= \exp\{-\Phi_i^u - \sum_{v \in L} \mu(u, v) \sum_{z \in \mathcal{N}_i, z \neq i} d(i, z) p_z^v\}, \quad (16) \end{aligned}$$

where $\mu(u, v)$ and $d(i, z)$ represent the global label co-occurrence and pairwise pixel similarity of Eqn. (4), respectively. This is because λ_1 is a constant, $d(i, i) = 0$, and $\mu(i, u, i, v) = \mu(u, v)$. Eqn. (16) is the corresponding MF update equation of (4).

4.3 Learning Algorithms

We describe the training strategy of DPN and also its time complexity and efficient implementation.

Learning. The first ten groups of DPN are initialized by VGG₁₆², while the last four groups can be initialized randomly. DPN is then fine-tuned in an incremental manner with four stages. During fine-tuning, all these stages solve the pixelwise softmax loss [6], but updating different sets of parameters.

First, we add a loss function to b11 and fine-tune the weights from b1 to b11 without the last four groups, in order to learn the unary terms. Second, to learn the triple relations, we stack b12 on top of b11 and update its parameters (i.e. ω_1, ω_2 in the distance measure), but the weights of the preceding groups (i.e. b1~b11) are fixed. Third, b13 and b14 are stacked onto b12 and similarly, their weights are updated with all the preceding parameters fixed, so as to learn the local label contexts. Finally, all the parameters are jointly fine-tuned.

Complexity. DPN transforms Eqn. (13) into convolutions and poolings in the groups from b12 to b15, such that filtering at each pixel can be performed in a parallel manner. Assume we have f input and f' output feature maps, $N \times N$ pixels, filters with $s \times s$ receptive field, and a mini-batch with M samples. b12 takes a total $f \cdot N^2 \cdot s^2 \cdot M$ operations, b13 takes $f \cdot f' \cdot N^2 \cdot s^2 \cdot M$ operations, while both b14 and b15 require $f \cdot N^2 \cdot M$ operations. For example, when $M=10$ as in our experiment, we have $21 \times 512^2 \times 50^2 \times 10 = 1.3 \times 10^{11}$ operations in b12, which has the highest complexity in DPN. If we parallelize these operations using matrix multiplication on GPU as [29] did, the operation in b12 can be computed within 30ms. The total runtime of the last four layers of DPN is 75ms. Note that convolutions in DPN can be further speeded up by low-rank decompositions [28] of the filters and model compressions [52].

In contrast, existing works [23], [26] employ fast Gaussian filtering [53] to accelerate the direct calculation of Eqn. (13). For a mini-batch of ten 512×512 images, a recently optimized implementation [18] needs 3.5×10^{11} operations and takes 114ms on GPU to compute one iteration of (13). In DPN, simultaneous message passing is enabled over every training image, in the sense that parallelization is obtained in image level. Therefore, DPN makes (13) easier to be parallelized and speeded up.

Efficient Implementation. As mentioned in Eqn. (4), the local filters in b12 are computed by the distances between RGB values of the pixels. XY coordinates are omitted here because

2. We use the released VGG₁₆ model, which is publicly available at http://www.robots.ox.ac.uk/~vgg/research/very_deep/

TABLE 2: Summary of datasets.

name	training	validation	testing	video data
VOC12	10582	1449	1456	no
Cityscapes	2975	500	1525	no
CamVid	367	-	233	yes

they could be pre-computed. To accelerate the computation of local convolution, the lookup table-based filtering approach is employed. Specifically, we first construct a lookup table storing distances between any two pixel intensities (ranging from 0 to 255), which results in a 256×256 matrix. Then when we perform local convolution, the kernels' coefficients can be obtained efficiently by just looking up the table.

5 EXPERIMENTS

In this section, we demonstrate the effectiveness of DPN and benchmark it against other state-of-the-art semantic segmentation methods. Below, we give an overview of the dataset and evaluation metrics used in these experiments. Representative methods are also introduced. In the following experiments, we denote 2-D DPN as DPN and 3-D DPN as spatial-temporal DPN.

Dataset. We compare DPN with the state-of-the-art methods on PASCAL VOC 2012 (VOC12) [30], Cityscapes [16] and CamVid [54] datasets. VOC12 is a well-known benchmark for generic image segmentation and Cityscapes dataset focuses on parsing urban street scenes. We choose those two benchmarks to evaluate the original DPN. On the other hand, CamVid dataset is composed of several video sequences, which is suitable for the evaluation of spatial-temporal DPN. We summarize the information of all datasets we used in Table 2.

Evaluation Metrics. All existing works employed mean pixel-wise intersection-over-union (denoted as mIoU) [6] to evaluate their performance. To fully examine the effectiveness of DPN, we introduce another three metrics, including tagging accuracy (TA), localization accuracy (LA), and boundary accuracy (BA). (1) TA compares the predicted image-level tags with the ground truth tags, calculating the accuracy of multi-class image classification. (2) LA evaluates the IoU between the predicted object bounding boxes³ and the ground truth bounding boxes (denoted as bIoU), measuring the precision of object localization. (3) For those objects that have been correctly localized, we compare the predicted object boundary with the ground truth boundary, measuring the precision of semantic boundary similar to [55].

Comparisons. DPN is compared with the state-of-the-art segmentation methods, including FCN [6], Zoom-out [5], DeepLab [23], WSSL [43], BoxSup [56], Piecewise [27], RNN [26], SuperParsing [57], Dilation10 [58], ALE [59], Multiclass [12], and SegNet [44]. All these methods are based on CNNs or MRFs. They can be grouped according to different aspects: (1) **joint-train:** Piecewise and RNN; (2) **w/o joint-train:** DeepLab, WSSL, FCN, and BoxSup; (3) **pre-train on COCO:** RNN, WSSL, and BoxSup. The first and the second groups are the methods with and without joint training CNNs and MRFs, respectively. Methods in the last group also employed MS-COCO [60] to pre-train deep models. To conduct a comprehensive comparison, the performance of DPN are reported on both settings, i.e., with and without pre-training on COCO.

3. They are the bounding boxes of the predicted segmentation regions.

TABLE 3: Ablation study of hyper-parameters.

Receptive Field	baseline	10×10	50×50	100×100
mIoU (%)	63.4	63.8	64.7	64.3

(a) Comparisons between different receptive fields of b12.

Receptive Field	1×1	5×5	9×9	9×9 mixtures
mIoU (%)	64.8	66.0	66.3	66.5

(b) Comparisons between different receptive fields of b13.

Pairwise Terms	DSN [24]	DeepLab [23]	DPN
improvement (%)	2.6	3.3	5.4

(c) Comparing pairwise terms of different methods.

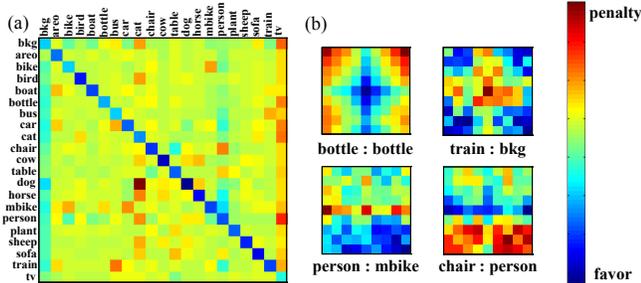


Fig. 6: Visualization of (a) learned label compatibility (b) learned contextual information. (Best viewed in color)

In the following, Sec. 5.1 investigates the effectiveness of different components of DPN on the VOC12. Sec. 5.2 evaluates the spatial-temporal DPN on the CamVid. Sec. 5.3 provides detailed analysis of the DPN system and its performance. Sec. 5.4 compares DPN with the state-of-the-art methods on the benchmarks.

5.1 Effectiveness of DPN

All the models evaluated in this section are trained on the *training set* and tested on the *validation set* of VOC12.

Triple Penalty. The receptive field of b12 indicates the range of triple relations for each pixel. We examine different settings of the receptive fields, including ‘10×10’, ‘50×50’, and ‘100×100’, as shown in Table 3(a), where ‘50×50’ achieves the best mIoU, which is slightly better than ‘100×100’. For a 512×512 image, this result implies that 50×50 neighborhood is sufficient to capture relations between pixels, while smaller or larger regions tend to under-fit or over-fit the training data. Moreover, all models of triple relations outperform the ‘baseline’ method that models dense pairwise relations, *i.e.* VGG₁₆+denseCRF [18].

Label Contexts. Receptive field of b13 indicates the range of local label context. To evaluate its effectiveness, we fix the receptive field of b12 as 50×50. As summarized in Table 3(b), ‘9×9 mixtures’ improves preceding settings by 1.7, 0.5, and 0.2 percent respectively. We observe large gaps exist between ‘1×1’ and ‘5×5’. Note that the 1×1 receptive field of b13 corresponds to learning a global label co-occurrence without considering local spatial contexts. Table 3(c) shows that the pairwise terms of DPN are more effective than DSN and DeepLab⁴.

4. The other deep models such as RNN and Piecewise did not report the exact improvements after combining unary and pairwise terms.

TABLE 4: The effectiveness of spatial-temporal DPN.

	2D	3D
Unary Term (mIoU)	59.67	59.67
+ Triple Penalty	59.93	60.16
+ Label Contexts	60.06	60.25

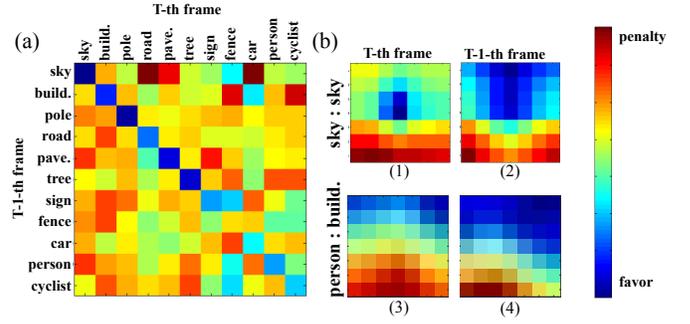


Fig. 7: Visualization of (a) 3D learned label compatibility (b) learned spatial-temporal contextual information. (Best viewed in color)

More importantly, mIoU of all the categories can be improved through increasing the size of receptive field and learning a mixture. Specifically, for each category, the improvements of the last three settings in Table 3(b) over the first one are 1.2 ± 0.2 , 1.5 ± 0.2 , and 1.7 ± 0.3 , respectively.

We also visualize the learned label compatibilities and contexts in Fig. 6(a) and (b), respectively. Fig. 6(a) is obtained by summing each filter in b13 over a 9×9 region, indicating how likely a column object would present when a row object is presented. Blue color represents high favorability. It is worth pointing out that Fig. 6(a) is non-symmetry. For example, when a ‘horse’ is presented, a ‘person’ is more likely to present than the other objects. On the other hand, when a ‘person’ presents in an image, it is more likely to find a ‘bike’ (than a ‘horse’). The ‘bkg’ (background) is compatible with all the objects. Fig. 6(b) visualizes some contextual patterns, where ‘A:B’ indicates that when ‘A’ is presented, where ‘B’ is more likely to present. For example, ‘bkg’ is around ‘train’, ‘motor bike’ is below ‘person’, and ‘person’ is sitting on ‘chair’.

5.2 Effectiveness of Spatial-temporal DPN

In this section, we quantitatively investigate the effectiveness of the pairwise term in spatial-temporal DPN. CamVid *train* are used for training and performance are reported on CamVid *test*. Images in CamVid are captured from a moving vehicle, thus exhibiting certain temporal regularization.

3D Convolution. In spatial-temporal DPN, the b12 is a 3D local convolutional layer and b13 is a 3D global convolutional layer. In Table 4, we evaluate the performance gain of each stage on DPN. According to Table 3(a), we set the receptive field of b12 as 50×50 and b13 as 9×9 in 2D setting. In 3D pairwise term, the receptive field of b12 is 50×50×3 and b13 is 7×7×3. We observe that applying 2D and 3D pairwise terms on unary term both improves the performance. Since 3D pairwise terms capture the information between successive frames, it performs slightly better than 2D pairwise terms.

Temporal Regularization. In Fig. 7, we visualize the 3D learned label compatibilities and contexts across frames.

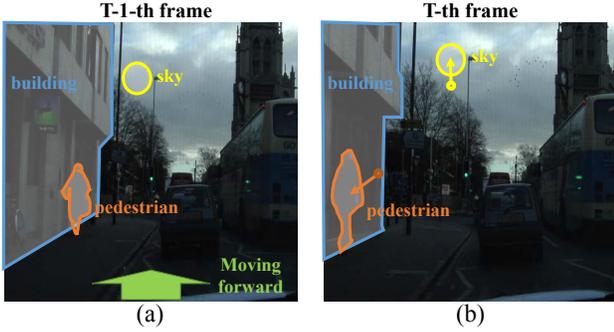


Fig. 8: (a) and (b) are neighboring video frames from CamVid dataset. When the vehicle moved forward, the ‘sky’ (blue) shift upward and the ‘pedestrian’ (orange) shift downward. **(Best viewed in color)**

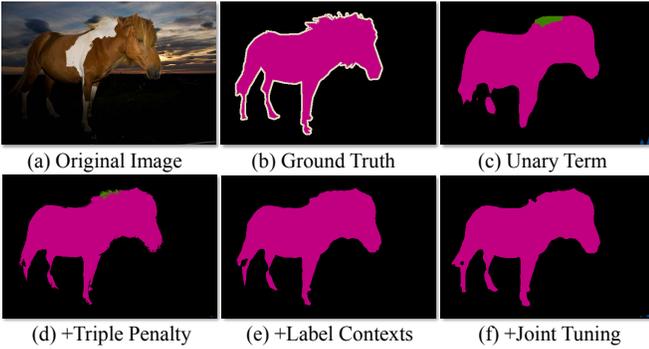


Fig. 9: Step-by-step visualization of DPN. **(Best viewed in color)**

Fig. 7(a) depicts the possibility of column object would present in the T -th frame when a row object is presented in the $T - 1$ -th frame. Blue color indicates high favorability. For instance, when ‘pavement’ is presented in the $T - 1$ -th frame, ‘road’ is more likely to present around in the T -th frame. However, if ‘sky’ is presented in the $T - 1$ -th frame, the pairwise term will penalize the probability which ‘road’ is presented in the T -th frame.

Fig. 7(b) visualizes the spatial-temporal contextual patterns. In each row, ‘A:B’ indicates how ‘A’ influences ‘B’. The figures suggest that when ‘A’ is presented in the T -th or the $T - 1$ -th frame, where ‘B’ is more likely to present in the T -th frame. In the first row, (1) represents that when ‘sky’ is presented in the T -th frame, its neighborhood is more likely to be ‘sky’. (2) means when ‘sky’ is presented in the $T - 1$ -th frame, it is more likely presented above this position in the T -th frame. As shown in Fig. 8, (a) and (b) are neighboring video frames, when the vehicle moves forward from (a) to (b), ‘sky’ region (yellow) shifts upward in the scene. So compared with the favor region in Fig. 7(b)-(1), the favor region also shifts upward in (2). In the second row of Fig. 7(b), (3) and (4) represent that when ‘pedestrian’ is presented in the T -th or the $T - 1$ -th frame, ‘building’ is more likely presented above in the T -th frame. We observe that the shifting from (3) to (4) is similar with the shifting of ‘pedestrian’ (orange) in Fig. 8 from (a) to (b).

5.3 Further Analysis

All the analysis in this section are conducted on the *validation set* of VOC12.

Incremental Learning. As discussed in Sec. 4.3, DPN is trained in an incremental manner. The right hand side of Table 5(a)

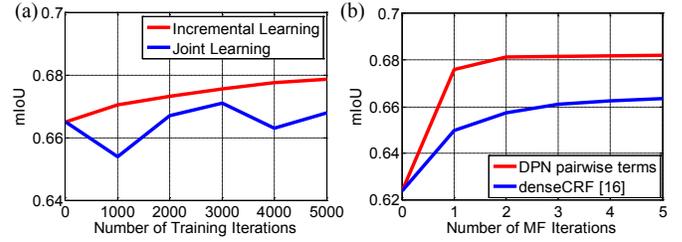


Fig. 10: Ablation study of (a) training strategy (b) required MF iterations. **(Best viewed in color)**

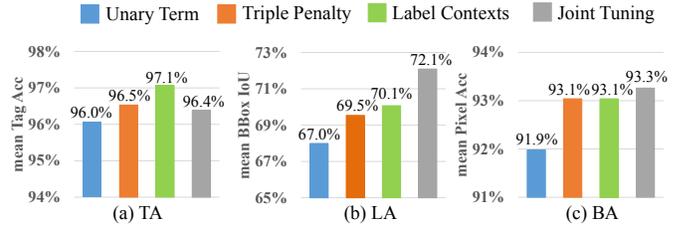


Fig. 11: Stage-wise analysis of (a) mean tagging accuracy (b) mean localization accuracy (c) mean boundary accuracy.

demonstrates that each stage leads to performance gain compared to its previous stage. For instance, ‘triple penalty’ improves ‘unary term’ by 2.3 percent, while ‘label contexts’ improves ‘triple penalty’ by 1.8 percent. More importantly, joint fine-tuning all the components (*i.e.* unary terms and pairwise terms) in DPN achieves another gain of 1.3 percent. A step-by-step visualization is provided in Fig. 9.

We also compare ‘incremental learning’ with ‘joint learning’, which fine-tunes all the components of DPN at the same time. The training curves of them are plotted in Fig. 10 (a), showing that the former leads to higher and more stable accuracies with respect to different iterations, while the latter may get stuck at local minima. This difference is easy to understand – incremental learning only introduces new parameters until all existing parameters have been fine-tuned.

One-iteration MF. DPN approximates one iteration of MF. Fig. 10(b) illustrates that DPN reaches a good accuracy with one MF iteration. A CRF [18] with dense pairwise edges needs more than 5 iterations to converge. It also has a large gap compared to DPN. Note that the existing deep models such as [23], [26], [24] required 5~10 iterations to converge as well.

Per-stage Analysis. We further evaluate DPN using three metrics. The results are given in Fig. 11. For example, (a) illustrates that the tagging accuracy can be improved in the third stage, as it captures label co-occurrence with a mixture of contextual patterns. However, TA decreases a little after the final stage. Since joint tuning maximizes segmentation accuracies by optimizing all components together, extremely small objects, which rarely occur in VOC training set, are discarded. As shown in (b), accuracies of object localization are significantly improved in the second and the final stages. This is intuitive because the unary prediction can be refined by long-range and high-order pixel relations, and joint training further improves results. (c) discloses that the second stage also captures object boundary, since it measures dissimilarities between pixels.

Per-class Analysis. Table 5(a) reports the per-class accuracies

TABLE 5: Per-class results on VOC12.

	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	Avg.
Unary Term (mIoU)	77.5	34.1	76.2	58.3	63.3	78.1	72.5	76.5	26.6	59.9	40.8	70.0	62.9	69.3	76.3	39.2	70.4	37.6	72.5	57.3	62.4
+ Triple Penalty	82.3	35.9	80.6	60.1	64.8	79.5	74.1	80.9	27.9	63.5	40.4	73.8	66.7	70.8	79.0	42.0	74.1	39.1	73.2	58.5	64.7
+ Label Contexts	83.2	35.6	82.6	61.6	65.5	80.5	74.3	82.6	29.9	67.9	47.5	75.2	70.3	71.4	79.6	42.7	77.8	40.6	75.3	59.1	66.5
+ Joint Tuning	84.8	37.5	80.7	66.3	67.5	84.2	76.4	81.5	33.8	65.8	50.4	76.8	67.1	74.9	81.1	48.3	75.9	41.8	76.6	60.4	67.8
TA (tagging Acc.)	98.8	97.9	98.4	97.7	96.1	98.6	95.2	96.8	90.1	97.5	95.7	96.7	96.3	98.1	93.3	96.1	98.7	92.2	97.4	96.3	96.4
LA (bIoU)	81.7	76.3	75.5	70.3	54.4	86.4	70.6	85.6	51.8	79.6	57.1	83.3	79.2	80.0	74.1	53.1	79.1	68.4	76.3	58.8	72.1
BA (boundary Acc.)	95.9	83.9	96.9	92.6	93.8	94.0	95.7	95.6	89.5	93.3	91.4	95.2	94.2	92.7	94.5	90.4	94.8	90.5	93.7	96.6	93.3

(a) Per-class results on VOC12 *val*.

	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mIoU
FCN [6]	76.8	34.2	68.9	49.4	60.3	75.3	74.7	77.6	21.4	62.5	46.8	71.8	63.9	76.5	73.9	45.2	72.4	37.4	70.9	55.1	62.2
Zoom-out [5]	85.6	37.3	83.2	62.5	66.0	85.1	80.7	84.9	27.2	73.2	57.5	78.1	79.2	81.1	77.1	53.6	74.0	49.2	71.7	63.3	69.6
DeepLab [23]	84.4	54.5	81.5	63.6	65.9	85.1	79.1	83.4	30.7	74.1	59.8	79.0	76.1	83.2	80.8	59.7	82.2	50.4	73.1	63.7	71.6
RNN [26]	87.5	39.0	79.7	64.2	68.3	87.6	80.8	84.4	30.4	78.2	60.4	80.5	77.8	83.1	80.6	59.5	82.8	47.8	78.3	67.1	72.0
Piecewise [27]	90.6	37.6	80.0	67.8	74.4	92.0	85.2	86.2	39.1	81.2	58.9	83.8	83.9	84.3	84.8	62.1	83.2	58.2	80.8	72.3	75.3
WSSL [†] [43]	89.2	46.7	88.5	63.5	68.4	87.0	81.2	86.3	32.6	80.7	62.4	81.0	81.3	84.3	82.1	56.2	84.6	58.3	76.2	67.2	73.9
RNN [†] [26]	90.4	55.3	88.7	68.4	69.8	88.3	82.4	85.1	32.6	78.5	64.4	79.6	81.9	86.4	81.8	58.6	82.4	53.5	77.4	70.1	74.7
BoxSup [†] [56]	89.8	38.0	89.2	68.9	68.0	89.6	83.0	87.7	34.4	83.6	67.1	81.5	83.7	85.2	83.5	58.6	84.9	55.8	81.2	70.7	75.2
Piecewise [†] [27]	94.1	40.7	84.1	67.8	75.9	93.4	84.3	88.4	42.5	86.4	64.7	85.4	89.0	85.8	86.0	67.5	90.2	63.8	80.9	73.0	78.0
DPN	87.7	59.4	78.4	64.9	70.3	89.3	83.5	86.1	31.7	79.9	62.6	81.9	80.0	83.5	82.3	60.5	83.2	53.4	77.9	65.0	74.1
DPN [†]	89.0	61.6	87.7	66.8	74.7	91.2	84.3	87.6	36.5	86.3	66.1	84.4	87.8	85.6	85.4	63.6	87.3	61.3	79.4	66.4	77.5

(b) Per-class results on VOC12 *test*. The approaches pre-trained on COCO [60] are marked with [†].

TABLE 6: Per-class results on Cityscapes.

	coarse	depth	sub	road	swalk	build.	wall	fence	pole	tlight	sign	veg.	terrain	sky	person	rider	car	truck	bus	train	mbike	bike	mIoU
Dilation10 [58]	no	no	no	97.6	79.2	89.9	37.3	47.6	53.2	58.6	65.2	91.8	69.4	93.7	78.9	55	93.3	45.5	53.4	47.7	52.2	66	67.1
Piecewise [27]	no	no	no	97.3	78.5	88.4	44.5	48.3	34.1	55.5	61.7	90.1	69.5	92.2	72.5	52.3	91	54.6	61.6	51.6	55	63.1	66.4
SiCNN+CRF	no	yes	no	96.3	76.8	88.8	40	45.4	50.1	63.3	69.6	90.6	67.1	92.2	77.6	55.9	90.1	39.2	51.3	44.4	54.4	66.1	66.3
FCN [6]	no	no	no	97.4	78.4	89.2	34.9	44.2	47.4	60.1	65	91.4	69.3	93.9	77.1	51.4	92.6	35.3	48.6	46.5	51.6	66.8	65.3
WSSL [43]	yes	no	2	97.4	78.3	88.1	47.5	44.2	29.5	44.4	55.4	89.4	67.3	92.8	71	49.3	91.4	55.9	66.6	56.7	48.1	58.1	64.8
DeepLab [23]	no	no	2	97.3	77.7	87.7	43.6	40.5	29.7	44.5	55.4	89.4	67	92.7	71.2	49.4	91.4	48.7	56.7	49.1	47.9	58.6	63.1
RNN [26]	no	no	2	96.3	73.9	88.2	47.6	41.3	35.2	49.5	59.7	90.6	66.1	93.5	70.4	34.7	90.1	39.2	57.5	55.4	43.9	54.6	62.5
DPN	no	no	no	97.5	78.5	89.5	40.4	45.9	51.1	56.8	65.3	91.5	69.4	94.5	77.5	54.2	92.5	44.5	53.4	49.9	52.1	64.8	66.8

TABLE 7: Per-class results on CamVid.

	build.	tree	sky	car	sign	road	person	fence	pole	pave.	cyclist	mIoU
ALE [59]	73.4	70.2	91.1	64.24	24.4	91.1	29.1	31	13.6	72.4	28.6	53.59
SuperParsing [57]	70.4	54.8	83.5	43.3	25.4	83.4	11.6	18.3	5.2	57.4	8.9	42.03
Tripathi <i>et al.</i> [13]	74.2	67.9	91	66.5	23.6	90.7	26.2	28.5	16.3	71.9	28.2	53.18
Liu and He [12]	66.8	66.6	90.1	62.9	21.4	85.8	28	17.8	8.3	63.5	8.5	47.2
SegNet [44]	68.7	52	87	58.5	13.4	86.2	25.3	17.9	16.0	60.5	24.8	46.4
DPN	80.6	72.6	91.2	77.8	40	90.7	43.9	28.7	15.9	71.4	47.9	60.06
Spatial-temporal DPN	80.6	73.1	91.4	77.9	40	90.8	43.9	29.2	16	71.9	47.9	60.25

of four evaluation metrics, where the first four rows represent the mIoU of four stages, while the last three rows represent TA, LA, and BA, respectively. We have several valuable observations, which motivate future researches. (1) Joint training benefits most of the categories, except animals such as ‘bird’, ‘cat’, and ‘cow’. Some instances of these categories are extremely small so that joint training discards them for smoother results. (2) Training DPN with pixelwise label maps implicitly models image-level tags, since it achieves a high averaged TA of 96.4%. (3) Object localization always helps. However, for objects with complex boundary such as ‘bike’, its mIoU is low even it can be localized, *e.g.* ‘bike’ has high LA but low BA and mIoU. (4) Failures of different categories have different factors. With these three metrics, they can be easily identified. For example, the failures of ‘chair’, ‘table’, and ‘plant’ are caused by the difficulties to accurately capture their bounding boxes and boundaries. Although ‘bottle’ and ‘tv’ are also difficult to localize, they achieve moderate mIoU because of their regular shapes. In other words, mIoU of ‘bottle’ and ‘tv’ can be significantly improved if they can be accurately localized.

5.4 Benchmarks

We evaluate the performance of DPN and spatio-temporal DPN on several standard semantic segmentation benchmarks.

5.4.1 Pascal VOC12

The VOC12 dataset is one of the most popular benchmarks for semantic image segmentation. This dataset contains 20 indoor and outdoor object categories and one background category. As previously mentioned, we employ 10,582 images for training, 1,449 images for validation, and 1,456 images for testing. Results are given Table 5 (b), we compare DPN with the best-performing methods⁵ on VOC12 test set based on two settings, *i.e.* with and without pre-training on COCO. The approaches pre-trained on COCO are marked with [†]. We evaluate DPN on several scales of the images and then average the results following [23], [27].

DPN outperforms several existing methods that were trained on VOC12, but DPN needs only one MF iteration to solve MRF,

⁵ The results of these methods were presented in either the published papers or arXiv pre-prints.

other than 10 iterations of RNN, DeepLab, and Piecewise. By averaging the results of two DPNs, we achieve 74.1% accuracy on VOC12 without outside training data. As discussed in Sec.4.3, MF iteration is the most complex step even when it is implemented as convolutions. Therefore, DPN at least reduces 10× runtime compared to previous works.

Following [26], [56], we pre-train DPN with COCO, where 20 object categories that are also presented in VOC12 are selected for training. A single DPN[†] has achieved 77.5% mIoU on VOC12 test set. The system in [27] has incorporated multi-scale pyramid training and mid-level feature refinement, which we believe is complementary to our contribution and can boost the performance of DPN as well. As shown in Table 5 (b), we observe that DPN[†] achieves competitive performances among multiple object classes.

5.4.2 Cityscapes

Cityscapes dataset focuses on street scenes segmentation. All images are captured from a moving vehicle in various seasons and cities. This dataset is highly different from VOC12. Because of the great depth of images, the street scenes have large scale variations even in a single image. It defines 19 object categories for evaluation. There are 2975 training, 500 validation and 1525 testing images with fine pixel-level annotations and extra 20000 images with cores annotations. To conduct a fair comparison, we only use the images with fine pixel-level annotations in our experiment.

As shown in Table 6, DPN achieves 66.8% on Cityscapes dataset, which is the second best method, and it is close to the first place 67.1% [58]. [58] takes a multi-scale aggregation strategy, which can be easily integrated into DPN to further boost the performance. Our method shows advantages on objects with arbitrary shape (*i.e.* ‘road’, ‘building’, ‘vegetation’, ‘terrain’ and ‘sky’) which can be easily confused with other objects. Thanks to our high-order and long-range pairwise term, DPN captures their appearance and segments them from other objects accurately.

5.4.3 CamVid

CamVid dataset consists of 367 training and 233 testing images with 11 classes annotations. All images are extracted from three video sequences at 1Hz. Similar to Cityscapes dataset, these video sequences are also captured from a moving vehicle. As shown in Table 7, DPN outperforms all existing methods, and spatial-temporal pairwise term further improves the performance to 60.25%. Considering the relatively sparse sampling rate (one frame per second) and strong unary term used, our 3-D MRF can indeed leverage temporal contextual information for joint inference.

We can observe that DPN achieves much better performance than other methods, especially on narrow and small objects (*i.e.* ‘pole’ and ‘sign’), which are very difficult in segmentation task. Since our triple penalty captures the appearance of pixels, we can predict more accurate boundary on those objects. The spatial-temporal pairwise term further improves the performance. Spatial-temporal DPN encodes the relationship between successive frames, which improve the performance on some specific categories like ‘tree’, ‘sky’ and ‘road’, because they have flexible shape and are continuous in successive frames. It is also worthwhile to note that our spatial-temporal DPN achieves best performances on 7 object categories out of 11.

5.5 Visual Quality Comparisons

In the following, we inspect visual quality of obtained label maps. Fig. 12 demonstrates the comparisons of DPN with FCN [6] and DeepLab [23]. We use the publicly released model⁶ to re-generate label maps of FCN while the results of DeepLab are extracted from their published paper. DPN generally makes more accurate predictions in both image-level and instance-level. For example, in Fig. 12 (row 2), DPN is able to discover all persons available while still recovering sharp boundaries of the aeroplane. Even in the challenging case of Fig. 12 (row 4) due to complex reflectance, the integrity of the bus is maintained by DPN. More examples of DPN label maps are shown in Fig. 14. We observe that learning local label contexts helps differentiate confusing objects and learning triple penalty facilitates the capturing of intrinsic object boundaries.

We also include some typical failure modes of DPN in Fig. 13. In the first case, the object with atypical pose is hard to be found. Further considering well-trained object detector as additional unary potential might be a potential solution to this challenge. The second case suggests that scale and illumination are also important factors that influence the performance. This problem can be partially alleviated by augmenting or adding the variance of training data.

6 CONCLUSION

We proposed Deep Parsing Network (DPN) to address semantic image/video segmentation. DPN has several appealing properties. First, DPN unifies the inference and learning of unary term and pairwise terms in a single convolutional network. No iterative inference is required during back-propagation. Second, high-order relations and mixtures of label contexts are incorporated to its pairwise terms modeling, making existing works as special cases. Third, DPN is built upon conventional operations of CNN, thus easy to be parallelized and speeded up.

DPN achieves state-of-the-art performance on VOC12, Cityscapes and CamVid datasets. Multiple valuable facts about semantic segmentation are revealed through extensive experiments, such as the interplay between tagging, localization and boundary accuracies along different processing stages. Future directions include investigating the generalizability of DPN to more challenging scenarios with large number of object classes and substantial appearance/scale variations.

ACKNOWLEDGMENT

This work is supported by SenseTime Group Limited and the General Research Fund sponsored by the Research Grants Council of the Hong Kong SAR (CUHK 416713, 14241716, 14224316, 14209217).

REFERENCES

- [1] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *TPAMI*, vol. 22, no. 8, 2000.
- [2] P. F. Felzenszwalb and D. P. Huttenlocher, “Efficient belief propagation for early vision,” *IJCV*, vol. 70, no. 1, 2006.
- [3] C. Liu, J. Yuen, and A. Torralba, “Nonparametric scene parsing via label transfer,” *TPAMI*, vol. 33, no. 12, 2011.
- [4] C. Farabet, C. Couprie, L. Najman, and Y. LeCun, “Learning hierarchical features for scene labeling,” *TPAMI*, vol. 35, no. 8, 2013.

6. <http://dl.caffe.berkeleyvision.org/fcn-8s-pascal.caffemodel>

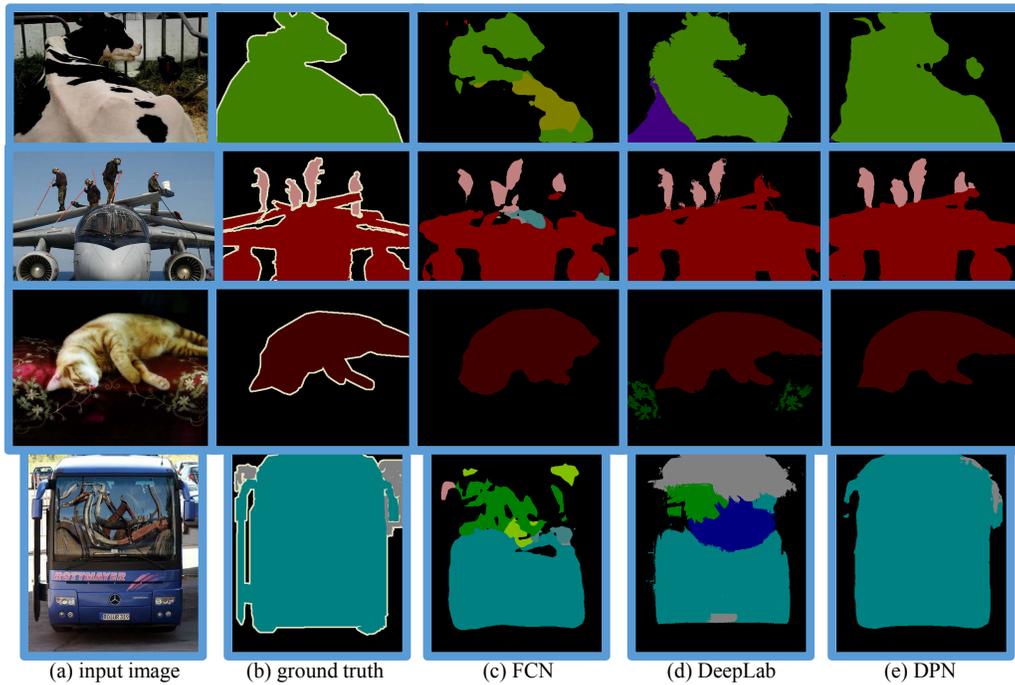


Fig. 12: Visual quality comparison of different semantic image segmentation methods: (a) input image (b) ground truth (c) FCN [6] (d) DeepLab [23] and (e) DPN.

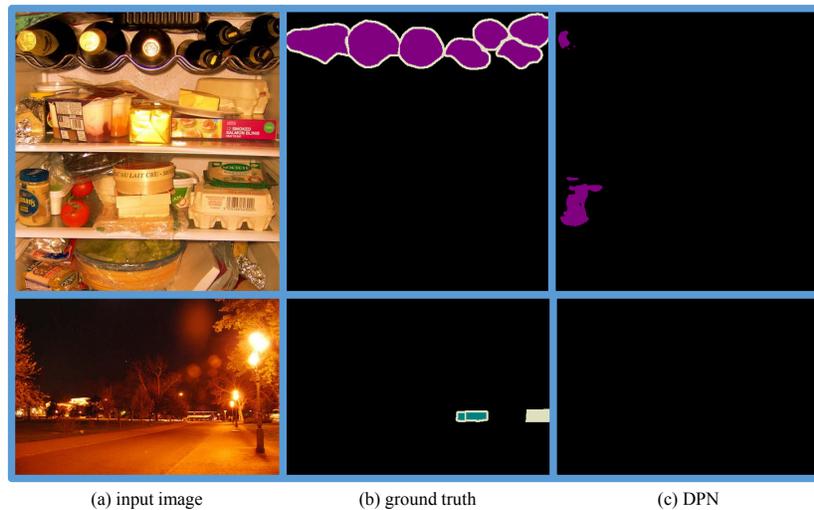


Fig. 13: Failure cases: (a) input image (b) ground truth (c) DPN.

- [5] M. Mostajabi, P. Yadollahpour, and G. Shakhnarovich, "Feedforward semantic segmentation with zoom-out features," in *CVPR*, 2015.
- [6] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *CVPR*, 2015.
- [7] Z. Liu, X. Li, P. Luo, C.-C. Loy, and X. Tang, "Semantic image segmentation via deep parsing network," in *ICCV*, 2015.
- [8] C. Gu and M.-C. Lee, "Semiautomatic segmentation and tracking of semantic video objects," *TCSVT*, vol. 8, no. 5, 1998.
- [9] G. J. Brostow, J. Fauqueur, and R. Cipolla, "Semantic object classes in video: A high-definition ground truth database," *Pattern Recognition Letters*, 2008.
- [10] J. Lezama, K. Alahari, J. Sivic, and I. Laptev, "Track to the future: Spatio-temporal video segmentation with long-range motion cues," in *CVPR*, 2011.
- [11] A. Jain, S. Chatterjee, and R. Vidal, "Coarse-to-fine semantic video segmentation using supervoxel trees," in *ICCV*, 2013.
- [12] B. Liu and X. He, "Multiclass semantic video segmentation with object-level active inference," in *CVPR*, 2015.
- [13] S. Tripathi, S. Belongie, Y. Hwang, and T. Nguyen, "Semantic video segmentation: Exploring inference efficiency," in *ISOC*, 2015.
- [14] A. Kundu, V. Vineet, and V. Koltun, "Feature space optimization for semantic video segmentation," in *CVPR*, 2016.
- [15] P. Isola and C. Liu, "Scene collaging: Analysis and synthesis of natural images with semantic layers," in *ICCV*, 2013.
- [16] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *CVPR*, 2016.
- [17] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, 1999.
- [18] P. Krähenbühl and V. Koltun, "Efficient inference in fully connected crfs with gaussian edge potentials," in *NIPS*, 2011.
- [19] V. Vineet, J. Warrell, and P. H. Torr, "Filter-based mean-field inference for random fields with higher-order terms and product label-spaces," in *ECCV*, 2012.

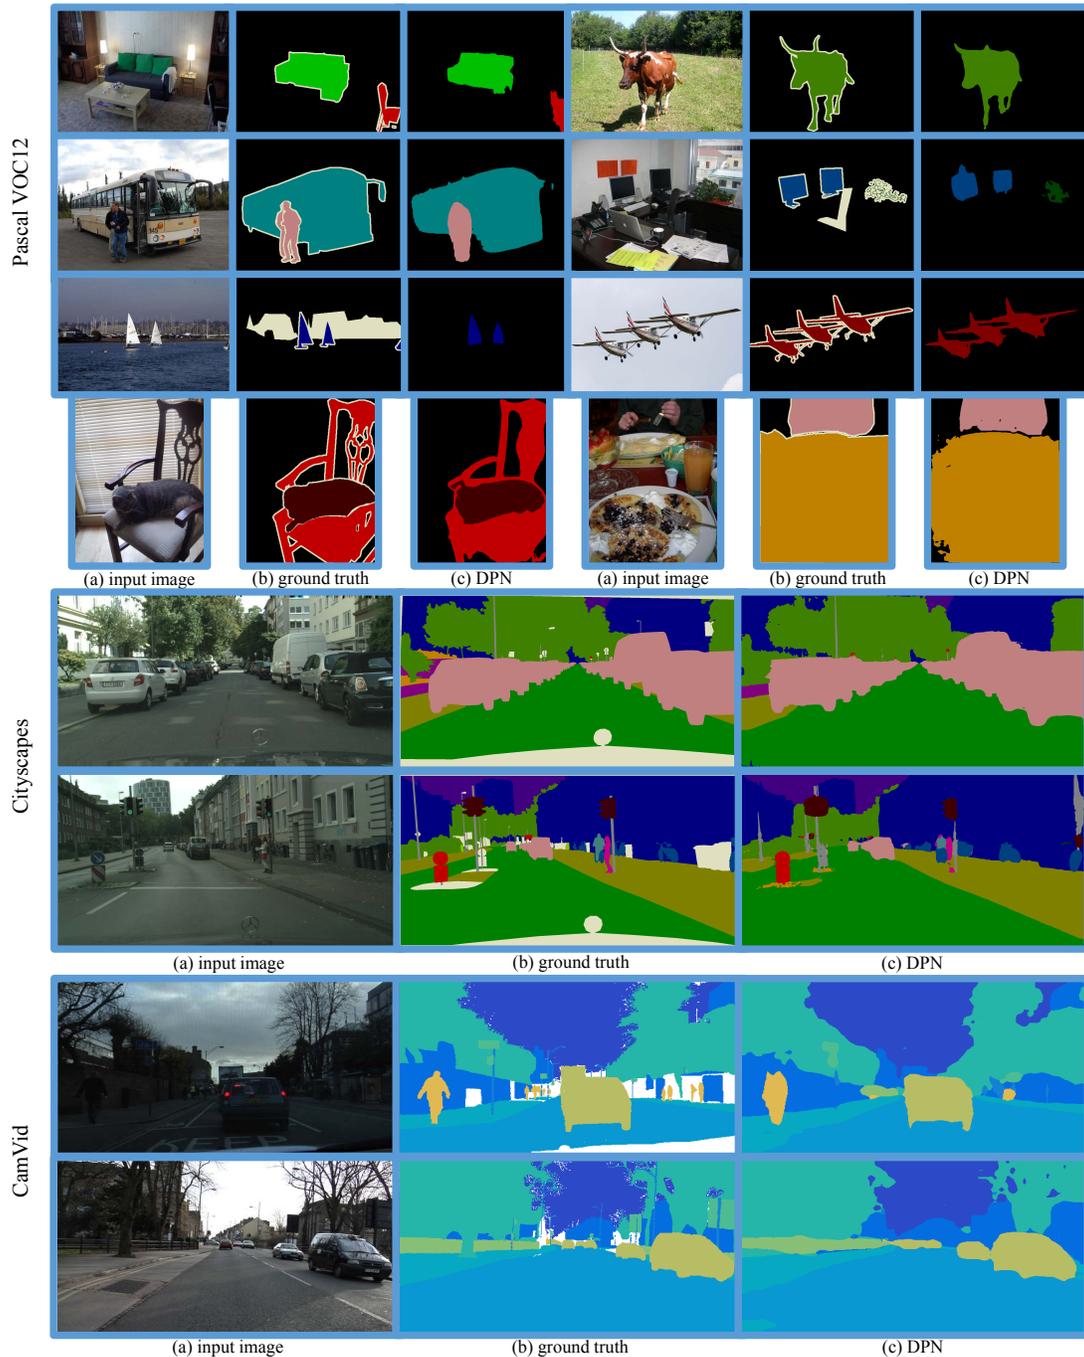


Fig. 14: Visual quality of DPN label maps: (a) input image (b) ground truth (white labels indicating ambiguous regions) and (c) DPN.

- [20] J. Yang, B. Price, S. Cohen, and M.-H. Yang, "Context driven scene parsing with attention to rare classes," in *CVPR*, 2014.
- [21] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [23] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," in *ICLR*, 2015.
- [24] A. G. Schwing and R. Urtasun, "Fully connected deep structured networks," *arXiv:1503.02351v1*, 2015.
- [25] M. Opper, O. Winther *et al.*, "From naive mean field theory to the tap equations," 2001.
- [26] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr, "Conditional random fields as recurrent neural networks," in *ICCV*, 2015.
- [27] G. Lin, C. Shen, I. Reid, and A. Hengel, "Efficient piecewise training of deep structured models for semantic segmentation," in *CVPR*, 2016.
- [28] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," in *BMVC*, 2014.
- [29] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, "cudnn: Efficient primitives for deep learning," in *NIPS Deep Learning Workshop*, 2014.
- [30] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *IJCV*, vol. 88, no. 2, 2010.
- [31] X. Ren and J. Malik, "Learning a classification model for segmentation," in *ICCV*, 2003.
- [32] M. Szummer, P. Kohli, and D. Hoiem, "Learning crfs using graph cuts," in *ECCV*, 2008.
- [33] B. Fulkerson, A. Vedaldi, and S. Soatto, "Class segmentation and object localization with superpixel neighborhoods," in *ICCV*, 2009.

- [34] P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *TPAMI*, vol. 33, no. 5, 2011.
- [35] G. Wang, P. Luo, L. Lin, and X. Wang, "Learning object interactions and descriptions for semantic image segmentation," in *CVPR*, 2017.
- [36] P. Luo, G. Wang, L. Lin, and X. Wang, "Deep dual learning for semantic image segmentation," in *ICCV*, 2017.
- [37] P. Krähenbühl and V. Koltun, "Parameter learning and convergent inference for dense random fields," in *ICML*, 2013.
- [38] V. Vineet, G. Sheasby, J. Warrell, and P. H. Torr, "Posefield: An efficient mean-field based method for joint estimation of human pose, segmentation, and depth," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, 2013.
- [39] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, and A. Yuille, "The role of context for object detection and semantic segmentation in the wild," in *CVPR*, 2014.
- [40] C. Wang, M. de La Gorce, and N. Paragios, "Segmentation, ordering and multi-object tracking using graphical models," in *ICCV*, 2009.
- [41] P. Luo, X. Wang, and X. Tang, "Hierarchical face parsing via deep learning," in *CVPR*, 2012.
- [42] P. Luo, X. Wang, and X. Tang, "Pedestrian parsing via deep decompositional network," in *ICCV*, 2013.
- [43] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille, "Weakly-and semi-supervised learning of a dcnn for semantic image segmentation," in *ICCV*, 2015.
- [44] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *arXiv preprint arXiv:1511.00561*, 2015.
- [45] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *IJCV*, vol. 40, no. 1, 2000.
- [46] Z. Liu, L. Yuan, X. Tang, M. Uyttendaele, and J. Sun, "Fast burst images denoising," *TOG*, vol. 33, no. 6, 2014.
- [47] C. Liu, J. Yuen, and A. Torralba, "Sift flow: Dense correspondence across scenes and its applications," *TPAMI*, vol. 33, no. 5, 2011.
- [48] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012.
- [49] Y. Sun, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in *NIPS*, 2014.
- [50] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *CVPR*, 2014.
- [51] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *ICCV*, 2015.
- [52] G. E. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *NIPS Deep Learning Workshop*, 2014.
- [53] A. Adams, J. Baek, and M. A. Davis, "Fast high-dimensional filtering using the permutohedral lattice," in *Eurographics*, 2010.
- [54] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla, "Segmentation and recognition using structure from motion point clouds," in *ECCV*, 2008.
- [55] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors," in *ICCV*, 2011.
- [56] J. Dai, K. He, and J. Sun, "Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation," in *ICCV*, 2015.
- [57] J. Tighe and S. Lazebnik, "Superparsing: scalable nonparametric image parsing with superpixels," in *ECCV*, 2010.
- [58] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," in *ICLR*, 2016.
- [59] C. Russell, P. Kohli, P. H. Torr *et al.*, "Associative hierarchical crfs for object class image segmentation," in *ICCV*, 2009.
- [60] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *ECCV*, 2014.



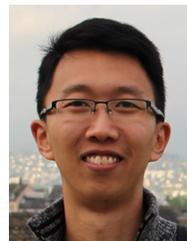
Ziwei Liu received his B.E. degree in Department of Electronics and Information Engineering in 2013, from Huazhong University of Science and Technology, Wuhan, China. He is currently a PhD student at the Department of Information Engineering, The Chinese University of Hong Kong. His research interests include computer vision, machine learning and computational photography, with focus on human-centric visual understanding and deep learning.



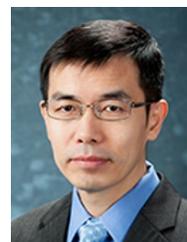
Xiaoxiao Li received his B.E. degree in Department of Computer Science and Technology from Tsinghua University in 2014. He is currently a PhD student at the Department of Information Engineering, The Chinese University of Hong Kong. His research interests include computer vision and machine learning, especially image segmentation and deep learning.



Ping Luo is a postdoctoral researcher at the Department of Information Engineering, The Chinese University of Hong Kong, where he received his Ph.D. in 2014. His research interests include deep learning, computer vision, and computer graphics, focusing on face analysis, pedestrian analysis, and large-scale object recognition and detection.



Chen Change Loy (M'06-M'10-SM'17) received the PhD degree in computer science from the Queen Mary University of London in 2010. He is currently a research assistant professor in the Department of Information Engineering, Chinese University of Hong Kong. Previously, he was a postdoctoral researcher at Queen Mary University of London and Vision Semantics Ltd. His research interests include computer vision and pattern recognition, with focus on face analysis, deep learning, and visual surveillance. He serves as an Associate Editor of IET Computer Vision Journal. He is a senior member of the IEEE.



Xiaoou Tang (S'93-M'96-SM'02-F'09) received the B.S. degree from the University of Science and Technology of China, Hefei, in 1990, and the M.S. degree from the University of Rochester, Rochester, NY, in 1991. He received the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, in 1996. He is a Professor and the Chairman chair of the Department of Information Engineering, Chinese University of Hong Kong. He worked as the group manager of the Visual Computing Group at the Microsoft Research Asia from 2005 to 2008. His research interests include computer vision, pattern recognition, and video processing. He received the Best Paper Award at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2009 and Outstanding Student Paper Award at the AAAI 2015. He was a program chair of the IEEE International Conference on Computer Vision (ICCV) 2009 and served as an associate editor of the IEEE Transactions on Pattern Analysis and Machine Intelligence and International Journal of Computer Vision. He is now the Editor-in-Chief of the International Journal of Computer Vision. He is a fellow of the IEEE.