

Global Deconvolutional Networks for Semantic Segmentation

Vladimir Nekrasov
nekrasowladimir@unist.ac.kr

Janghoon Ju
janghoon.ju@unist.ac.kr

Jaesik Choi
jaesik@unist.ac.kr

Ulsan National Institute of Science and
Technology
50 UNIST, Ulsan, 44919 Korea

Abstract

Semantic image segmentation is a principal problem in computer vision, where the aim is to correctly classify each individual pixel of an image into a semantic label. Its widespread use in many areas, including medical imaging and autonomous driving, has fostered extensive research in recent years. Empirical improvements in tackling this task have primarily been motivated by successful exploitation of Convolutional Neural Networks (CNNs) pre-trained for image classification and object recognition. However, the pixel-wise labelling with CNNs has its own unique challenges: (1) an accurate deconvolution, or upsampling, of low-resolution output into a higher-resolution segmentation mask and (2) an inclusion of global information, or context, within locally extracted features. To address these issues, we propose a novel architecture to conduct the equivalent of the deconvolution operation globally and acquire dense predictions. We demonstrate that it leads to improved performance of state-of-the-art semantic segmentation models on the PASCAL VOC 2012 benchmark, reaching 74.0% mean IU accuracy on the test set.

1 Introduction

Convolutional Neural Networks [26] have become an essential part of deep learning models [27] designed to tackle a wide range of computer vision tasks including image classification and recognition [24, 42, 44, 46, 50], image captioning [20, 47, 48], object detection [10, 11, 37]. Recent advances in computing technologies with efficient utilisation of Graphical Processing Units (GPUs), as well as availability of large-scale datasets [6, 29] have been among primary factors in such a rapid rise in CNN popularity.

An adaptation of convolutional network models [31], pre-trained for the image classification task, has fostered extensive research on the exploitation of CNNs in semantic image segmentation - a problem of marking (or classifying) each pixel of the image with one of the given semantic labels. Among important applications of this problem are road scene understanding [2, 3, 45], biomedical imaging [5, 38], aerial imaging [21, 32].

Recent breakthrough methods in the area have efficiently and effectively combined neural networks with probabilistic graphical models, such as Conditional Random Fields (CRFs)

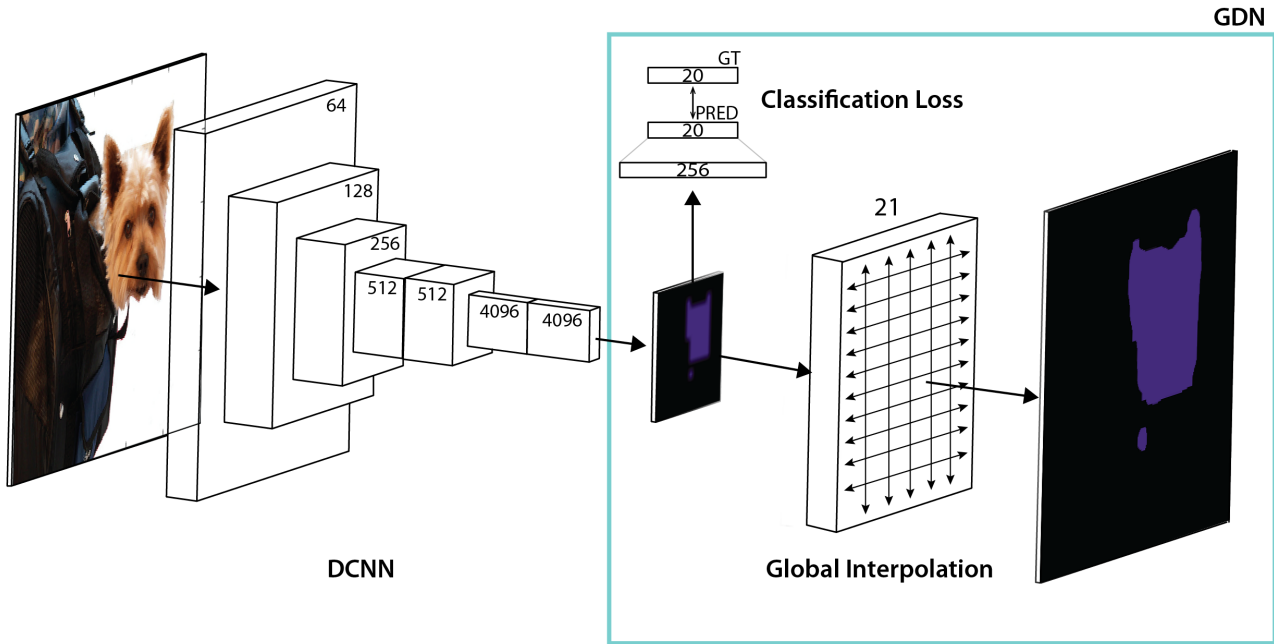


Figure 1: **Global Deconvolutional Network.** Our adaptation of FCN-32s [31]. After hierarchical blocks of convolutional-subsampling-nonlinearity layers, we are upsampling the reduced signal with the help of the global interpolation block. In addition to the pixel-wise softmax loss (not shown here), we also use the multi-label classification loss to increase the recognition accuracy.

[4, 28, 52] and Markov Random Fields (MRFs) [30]. These approaches usually refine per-pixel features extracted by CNNs (so-called ‘unary potentials’) with the help of pairwise similarities between the pixels based on location and colour features, followed by an approximate inference of the obtained fully connected graphical model [23].

In this work, we address two main challenges of current CNN-based semantic segmentation methods [4, 31]: an effective deconvolution, or upsampling, of low-resolution output from a neural network; and an inclusion of global information, or context, into existing models without relying on graphical models. Our contribution is twofold: i) we propose a novel approach for performing the deconvolution operation of the encoded signal and ii) demonstrate that this new architecture, called ‘*Global Deconvolutional Network*’, achieves close to the state-of-the-art performance on semantic segmentation with a simpler architecture and significantly lower number of parameters in comparison to the existing models [4, 30, 34].

The rest of the paper is structured as follows. We briefly explore recent common practices of semantic segmentation models in Section 2. Section 3 presents our approach designed to overcome the issues outlined above. Section 4 describes the experimental part, including the evaluation results of the proposed method on the popular PASCAL VOC dataset. Finally, Section 5 contains conclusions.

2 Related Work

Exploitation of fully convolutional neural networks has become ubiquitous in semantic image segmentation ever since the publication of the paper by Long *et al.* [31]. Further re-

search has been concerned with the combination of CNNs and probabilistic graphical models [4, 28, 30, 52], training in the presence of weakly-labelled data [18, 35, 40], learning an additional (deconvolutional) network [34].

The problem of incorporation of contextual information has been an active research topic in computer vision [7, 8, 16, 33, 36]. To some extent, probabilistic graphical models address this issue in semantic segmentation and can be either a) used as a separate post-processing step [4] or b) trained end-to-end with CNNs [28, 30, 52]. In setting a) graphical models are unable to refine the parameters of the CNN and thus the errors from the CNN will essentially be presented during post-processing. On the other hand, in b) one need to carefully design the inference part in terms of usual neural networks operations, and it still relies on computing high-dimensional Gaussian kernels [1]. Besides that, Yu and Koltun [49] have recently shown that dilated convolution filters are generally applicable and allow to increase the contextual capacity of the network as well.

In terms of improving the deconvolutional part of the network for dense predictions, there has been a prevalence of using information from lower layers: so-called ‘Skip Architecture’ [31, 38] and Multi-scale [4] are two notable examples. Noh *et al.* [34] proposed to train a separate deconvolutional network to effectively decode information from the original fully convolutional model. While these methods have given better results, all of them contain significantly more parameters than the corresponding baseline models.

In turn, we propose another approach, called ‘*Global Deconvolutional Network*’, which includes a global interpolation block with an additional recognition loss, and gives better results than multi-scale and ‘skip’ variants. The depiction of our architecture is presented in Figure 1.

3 Global Deconvolutional Network

In this section, we describe our approach intended to boost the performance of deep learning models on the semantic segmentation task.

3.1 Baseline Models

As baseline models, we choose two publicly available deep CNN models: FCN-32s¹ [31] and DeepLab² [4]. Both of them are based on the VGG 16-layer net [44] from the ILSVRC-2014 competition [41]. This network contains 16 weight layers, including two fully-connected ones, and can be represented as hierarchical stacks of convolutional layers with rectified linear unit non-linearity [13] followed by pooling operations after each stack. The output of the fully-connected layers is fed into a softmax classifier.

For semantic segmentation, where one needs to acquire dense predictions, the fully-connected layers have been replaced by convolution filters followed by a learnable deconvolution or fixed bilinear interpolation to match the original spatial dimensions of the image. The pixel-wise softmax loss represents the objective function.

¹<https://github.com/BVLC/caffe/wiki/Model-Zoo>

²<https://bitbucket.org/deeplab/deeplab-public/>

3.2 Global Interpolation

The output of multiple blocks of convolutional and pooling layers is an encoded image with severely reduced dimensions. To predict the segmentation mask of the same resolution as the original image, one needs to simultaneously decode and upsample this coarse output. A natural approach is to perform an interpolation. In this work, instead of applying conventional local methods, we devise a learnable global interpolation.

We denote the decoded information of the RGB-image $\mathbf{I} : \mathbf{I} \in \mathbb{R}^{3 \times H \times W}$, as $\mathbf{x} : \mathbf{x} \in \mathbb{R}^{C \times h \times w}$, where C represents the number of channels, h and w define the reduced height H and width W , respectively. To acquire $\mathbf{y} : \mathbf{y} \in \mathbb{R}^{C \times H \times W}$, an upsampled signal, we apply the following formula:

$$\mathbf{y}_c = \mathbf{K}_h \mathbf{x}_c \mathbf{K}_w^T, \forall c \in \mathbf{C} \quad (1)$$

where the matrices $\mathbf{K}_h \in \mathbb{R}^{H \times h}$ and $\mathbf{K}_w \in \mathbb{R}^{W \times w}$ are interpolating each feature map of \mathbf{x} through the corresponding spatial dimensions. Opposite to a simple bilinear interpolation, which operates only on the closest four points, the equation above allows to include much more information on the rectangular grid. An illustrative example can be seen in Figure 2.

$$\begin{bmatrix} & \mathbf{K}_h & \\ & & \mathbf{x}_c \\ & & & \mathbf{K}_w^T & \\ & & & & \mathbf{y}_c \end{bmatrix}
 \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}
 \begin{bmatrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix}
 \begin{bmatrix} 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}
 =
 \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

Figure 2: Depiction of Equation (1) for synthetic data.

Note that this operation is differentiable, and during the backpropagation algorithm [39] the derivatives of the pixelwise cross-entropy loss function \mathcal{L}_s with respect to the input \mathbf{x} and parameters $\mathbf{K}_h, \mathbf{K}_w$ can be found as follows:

$$\begin{aligned} \frac{\partial \mathcal{L}_s}{\partial \mathbf{x}_c} &= \mathbf{K}_h^T \frac{\partial \mathcal{L}_s}{\partial \mathbf{y}_c} \mathbf{K}_w^T, \\ \frac{\partial \mathcal{L}_s}{\partial \mathbf{K}_w} &= \left(\frac{\partial \mathcal{L}_s}{\partial \mathbf{y}_c} \right)^T \mathbf{K}_h \mathbf{x}_c, \quad \frac{\partial \mathcal{L}_s}{\partial \mathbf{K}_h} = \left(\frac{\partial \mathcal{L}_s}{\partial \mathbf{y}_c} \right) \mathbf{K}_w \mathbf{x}_c^T. \end{aligned} \quad (2)$$

We call the operation performed by Equation (1) ‘*global deconvolution*’. We only use this term to underline the fact that we mimic the behaviour of standard deconvolution using a global function; note that our method is not the inverse of the convolution operation and therefore does not represent deconvolution in the strictest sense as, for example, in [51].

3.3 Multi-task loss

It is not uncommon to force intermediate layers of deep learning networks to preserve meaningful and discriminative representations. For example, Szegedy *et al.* [46] appended several auxiliary classifiers to the middle blocks of their model.

As semantic image segmentation essentially comprises image classification as one of its sub-tasks, we append an additional objective function on the top of the coarse output to

further improve the model performance on the particular task of classification (Figure 1). This supplementary block consists of 3 fully-connected layers, with the length of the last one being equal to the pre-defined number of possible labels (excluding the background). As there are usually multiple instances of the same label presented in the image, we do not explicitly encode the quantitative components and only denote the presence of a particular class or its absence. The scores from the last layer are transformed with the sigmoid function followed by the multinomial cross entropy loss.

The loss functions are defined as follows:

$$\begin{aligned}\mathcal{L}_s(I, G) &= \frac{-1}{|I|} \sum_{i \in I} \log(\hat{p}_{iG_i}), \\ \mathcal{L}_c(I, L) &= \frac{-1}{|C|} \sum_{c \in C} [p_c \log(\hat{p}_c) + (1 - p_c) \log(1 - \hat{p}_c)],\end{aligned}\tag{3}$$

where \mathcal{L}_c is the multi-label classification loss; \mathcal{L}_s is the pixelwise cross-entropy loss; I is the set of pixels; G is a ground truth map; C is the number of possible labels; L is a ground truth binary vector of length $|C|$; $\hat{p}_{ic} \in [0, 1]$ is the softmax probability of pixel i being assigned to class c ; $p_c \in \{0, 1\}$ indicates the presence of class c or its absence; $\hat{p}_c \in [0, 1]$ corresponds to the predicted probability of class c being presented in the image. Note that it is possible to use a weighted sum of the two losses depending on which task's performance we want to optimise.

Overall, each component of the proposed approach aims to capture global information and incorporate it into the network, hence the name *global deconvolutional network*. Besides that, the proposed interpolation also effectively upsamples the coarse output and a nonlinear upsampling can be achieved with the addition of an activation function on the top of the block. The complete architecture of our approach is presented in Figure 1.

4 Experiments

4.1 Implementation details

We have implemented the proposed methods using Caffe [19], the popular deep learning framework. Our training procedure follows the practice of the corresponding baseline models: DeepLab [4] and FCN [31]. Both of them employ the VGG-16 net pre-trained on ImageNet [6].

We use Stochastic Gradient Descent (SGD) with momentum and train with a minibatch size of 20 images. We start the training process with the learning rate equal to 10^{-8} and divide it by 10 when the validation accuracy stops improving. We use momentum of 0.9 and weight decay of 0.0005. We initialise all additional layers randomly as in [12] and fine-tune them by backpropagation with a lower learning rate before finally training the whole network.

4.2 Dataset

We evaluate performance of the proposed approach on the PASCAL VOC 2012 segmentation benchmark [9], which consists of 20 semantic categories and one background category.

Following [14], we augment the training data to 8498 images and to 10582 images for the FCN and DeepLab models, respectively.

The segmentation performance is evaluated by the mean pixel-wise intersection-over-union (mIoU) score [9], defined as follows:

$$mIoU = \frac{1}{|C|} \sum_{c \in C} n_{cc} / \left(\sum_{c' \in C} (n_{cc'} + n_{c'c}) - n_{cc} \right), \quad (4)$$

where $n_{cc'}$ represents the number of pixels of class c predicted to belong to class c' .

First, we conduct all our experiments on the PASCAL VOC *val* set, and then compare the best performing models with their corresponding baseline models on the PASCAL VOC *test* set. As the annotations for the test data are not available, we send our predictions to the PASCAL VOC Evaluation Server.³

4.3 Experiments with FCN-32s

We conduct several experiments with FCN-32s as a baseline model. During the training stage the images are resized to 500×500 .⁴ We evaluate all the models on the holdout dataset of 736 images as in [31], and send the test results of the best performing ones to the evaluation server.

The original FCN-32s model employs the standard deconvolution operation (also known as *backwards convolution*) to upsample the coarse output. We replace it with the proposed global deconvolution and randomly initialise the new parameters as in [12]. We fix the rest of the network to pre-train the added block, and after that train the whole network. Global interpolation already improves its baseline model on the validation dataset, as can be seen in Table 1.


The baseline model deals with inputs of different sizes via cropping the predicted mask to the same resolution as the corresponding input. Other popular options include either 1) padding or 2) resizing to the fixed input dimensions. In case of global deconvolution, we propose a more elegant solution. Recall that the parameters of this block can be represented as matrices $\mathbf{K}_h \in \mathbb{R}^{H \times h}$ and $\mathbf{K}_w \in \mathbb{R}^{W \times w}$, where $H = W = 500$, $h = w = 16$ during the training stage. Then, given a test image $\{\mathbf{I} \in \mathbb{R}^{3 \times \hat{H} \times \hat{W}} | \hat{W}, \hat{H} \leq 500\}$, we subset the learned matrices to acquire $\hat{\mathbf{K}}_h \in \mathbb{R}^{\hat{H} \times \hat{h}}$ and $\hat{\mathbf{K}}_w \in \mathbb{R}^{\hat{W} \times \hat{w}}$ ($\hat{w}, \hat{h} \leq 16$) and proceed with the same operation. To subset, we only leave first \hat{H}, \hat{W} rows and \hat{h}, \hat{w} columns of the corresponding matrices, and discard all the rest. We have found that this do not affect the final performance of our model.

Next, to increase the recognition accuracy we also append the multi-label classification loss. This slightly improves the validation score in comparison to the baseline model, while the combination with global interpolation gives a further boost in performance (FCN-32s+GDN).


Besides that, we have also conducted additional experiments with FCN-32s, where we insert a fully-connected layer directly after the coarse output (FCN-32s+FC). The idea behind this trick is to allow the network to refine the local predictions based on the information from all the pixels. One drawback in such an approach is the requirement of the fully-connected layers to have the fixed-size input, although the solutions discussed above are also

³<http://host.robots.ox.ac.uk/>

⁴This is the maximum value for both the height and width in the PASCAL VOC dataset.



(a) FCN-32s as the baseline model



(b) DeepLab-LargeFOV as the baseline model

Figure 3: **Qualitative results on the validation set.** (a) Last column represents our approach, which includes the replacement of standard deconvolution with global deconvolution and the addition of the multi-label classification loss. (b) Fourth and sixth columns demonstrate our model, where bilinear interpolation is replaced with global deconvolution. Last two columns also incorporate a conditional random field (CRF) as a post-processing step. *Best viewed in colour.*

Method	mean IoU
FCN-32s [31]	59.4
FCN-32s + Label Loss	59.8
FCN-32s + Global Interp.	60.9
FCN-32s + GDN	61.2
FCN-32s + GDN + FC	62.5
DL-LargeFOV [4]	73.3
DL-LargeFOV + Label Loss	73.9
DL-LargeFOV + Global Interp.	74.2
DL-LargeFOV + GDN	75.1

Table 1: Mean intersection over union accuracy of our approach (GDN), which includes the addition of multi-label classification loss and global interpolation, compared with the baseline model on the reduced validation dataset (for FCN-32s) and on the PASCAL VOC 2012 validation dataset (for DL-LargeFOV).

applicable here. Nevertheless, neither of these methods gives satisfactory results during the empirical evaluations. Therefore, we proceed with a slightly different architecture: before appending the fully-connected layer, we first add a spatial pyramid pooling layer [15], which produces the output of the same length given an arbitrarily sized input. In particular, we are using a 5-level pyramid with max-pooling. Though during evaluation this approach alone does not give any improvement in the validation score over the baseline model, its ensemble with the global deconvolution model (FCN-32s+GDN+FC) improves previous results, which indicates that these models may be complementary to each other.

We continue with the evaluation of the best performing models on the test set (Table 2). Both of them improve their baseline model, FCN-32s, and even outperform FCN-8s, another model by Long *et al.* [31] with the skip-architecture, which combines information from lower layers with the final prediction layer.

Some examples of our approach can be seen in Figure 3.

4.4 Experiments with DeepLab

As the next baseline model we consider DeepLab-LargeFOV [4]. With the help of the *algorithme à trous* [17, 43], the model has a larger Field-of-View (FOV), which results in the finer predictions from the network. Besides that, this model is significantly faster and contains fewer parameters, than the plain modification of the VGG-16 net due to the reduced number of filters of the last two convolutional layers. The model employs simple bilinear interpolation to acquire the output of the same resolution as the input.

We proceed with the same experiments as for the FCN-32s model, except for the ones involving the fully-connected layer. As DeepLab-LargeFOV has a higher resolution coarse output, the inclusion of the fully-connected layer would result in the weight matrix of several billions parameters. Therefore, we omit these experiments.

We separately replace the bilinear interpolation with global deconvolution, append the label loss and estimate the joint GDN model. We carry out the same strategy outlined above during the testing stage to deal with variable-size inputs. All the experiments lead to improvements over the baseline model, with GDN showing a significantly higher score on the PASCAL VOC 2012 val set (Table 1).




Figure 4: Failure cases of our approach, Global Deconvolutional Network (GDN), on the PASCAL VOC 2012 val set. *Best viewed in colour.*

The DeepLab-LargeFOV model also incorporates a fully-connected CRF [22, 23, 25] as a post-processing step. To set the parameters of the fully connected CRF, we employ the same method of cross-validation as in [4] on a subset of the validation data. Then we send our best performing model enriched by CRF to the evaluation server. On the PASCAL VOC 2012 test set our single model (DL-LargeFOV-CRF+GDN) achieves 73.2% mIoU, a significant improvement over the baseline model (around 2.9%), and even excels the multiscale DeepLab-MSc-LargeFOV model by 1.6% (Table 2); the predictions averaged across our several models (DL-L_FOV-CRF+GDN_ENS) give a further improvement of 0.8%, showing a competitive score to the models that do not exploit the Microsoft COCO dataset [29]. As is the case with the FCN-32s model, we obtain performance on par with the multi-resolution variant using a much simpler architecture. Moreover, our single CRF-equipped global deconvolutional network (DL-LargeFOV-CRF+GDN) even surpasses the results of the competing approach (DeconvNet+CRF [34]) by 2.7%, where the deconvolutional part of the network contains significantly more parameters: almost 126M compared to less than 70K of global deconvolution; in case of ensembles, the improvement is over 1.5%.

The illustrative examples are presented in Figures 3 and 4.

5 Conclusion

In this paper we addressed two important problems of semantic image segmentation: an upsampling of the low-resolution output from the network and refinement of this coarse output, incorporating global information and the additional classification loss. We proposed

METHOD	AERO	BIKE	BIRD	BOAT	BOTTLE	BUS	CAR	CAT	CHAIR	COW	TABLE	DOG	HORSE	MBIKE	PERSON	PLANT	SHEEP	SOFA	TRAIN	TV	mIoU
FCN-8s [31]	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	62.20
FCN-32s + GDN	74.5	31.8	66.6	49.7	60.5	76.9	75.8	76.0	22.8	57.5	54.5	72.9	59.4	74.9	73.6	50.9	67.5	43.2	70.0	56.4	62.22
FCN-32s + GDN + FC	75.6	31.5	69.2	51.6	62.9	78.8	76.7	78.6	24.6	61.6	60.3	74.5	62.6	76.0	74.3	51.4	70.6	47.3	73.9	58.3	64.37
DL-LARGEFOV-CRF [4]	83.4	36.5	82.5	62.2	66.4	85.3	78.4	83.7	30.4	72.9	60.4	78.4	75.4	82.1	79.6	58.2	81.9	48.8	73.6	63.2	70.34
DECONVNET+CRF_VOC[34]	87.8	41.9	80.6	63.9	67.3	88.1	78.4	81.3	25.9	73.7	61.2	72.0	77.0	79.9	78.7	59.5	78.3	55.0	75.2	61.5	70.50
DL-MSC-LARGEFOV-CRF [4]	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	71.60
EDECONVNET+CRF_VOC[34]	89.9	39.3	79.7	63.9	68.2	87.4	81.2	86.1	28.5	77.0	62.0	79.0	80.3	83.6	80.2	58.8	83.4	54.3	80.7	65.0	72.50
DL-LARGEFOV-CRF + GDN	87.9	37.8	88.8	64.5	70.7	87.7	81.3	87.1	32.5	76.7	66.7	80.4	76.6	82.2	82.3	57.9	84.5	55.9	78.5	64.2	73.21
DL-L_FOV-CRF + GDN_ENS	88.6	48.6	88.8	64.7	70.4	87.2	81.8	86.4	32.0	77.1	64.1	80.5	78.0	84.0	83.3	59.2	85.9	56.8	77.9	65.0	74.02
ADELAIDE_CONT_CRF_VOC [30]	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	75.30

Table 2: Mean intersection over union accuracy of our approach (GDN), compared with the competing models on the PASCAL VOC 2012 test set.⁵ Lacking the results of FCN-32s on the test set, we thus compare it directly with a more powerful model, FCN-8s [31]. All the methods presented here do not use MS COCO dataset.

a novel approach, *global deconvolution*, to acquire the output of the same size as the input for images of variable resolutions. We showed that *global deconvolution* effectively replaces standard approaches, and can easily be trained in a straightforward manner.

On the benchmark competition, PASCAL VOC 2012, we showed that the proposed approach outperforms the results of the baseline models. Furthermore, our method even surpasses the performance of more powerful multi-resolution models, which combine information from several blocks of the deep neural network.

Acknowledgements The authors would like to thank the anonymous reviewers for their helpful and constructive comments, and Gae Kim for making Fig. 1. This work is supported by the Ministry of Science, ICT & Future Planning (MSIP), Korea, under Basic Science Research Program through the National Research Foundation of Korea (NRF) grant (NRF-2014R1A1A1002662), under the ITRC (Information Technology Research Center) support program (IITP-2016-R2720-16-0007) supervised by the IITP (Institute for Information & communications Technology Promotion) and under NIPA (National IT Industry Promotion Agency) program (NIPA-S0415-15-1004).

References

- [1] Andrew Adams, Jongmin Baek, and Myers Abraham Davis. Fast high-dimensional filtering using the permutohedral lattice. *Comput. Graph. Forum*, 29(2):753–762, 2010.
- [2] José Manuel Álvarez, Theo Gevers, Yann LeCun, and Antonio M. López. Road scene segmentation from a single image. In *ECCV*, 2012.
- [3] Vijay Badrinarayanan, Ankur Handa, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *CoRR*, abs/1505.07293, 2015. URL <http://arxiv.org/abs/1505.07293>.
- [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. *CoRR*, abs/1412.7062, 2014.

⁵<http://host.robots.ox.ac.uk:8080/leaderboard/displaylb.php?challengeid=11&compid=6>

- [5] Dan C. Ciresan, Alessandro Giusti, Luca Maria Gambardella, and Jürgen Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *NIPS*, 2012.
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [7] Santosh Kumar Divvala, Derek Hoiem, James Hays, Alexei A. Efros, and Martial Hebert. An empirical study of context in object detection. In *CVPR*, 2009.
- [8] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Context as supervisory signal: Discovering objects with predictable context. In *ECCV*, 2014.
- [9] Mark Everingham, Luc J. Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [10] Ross B. Girshick. Fast R-CNN. *CoRR*, abs/1504.08083, 2015.
- [11] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*.
- [12] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *AISTATS*, 2010.
- [13] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *AISTATS*, 2011.
- [14] Bharath Hariharan, Pablo Arbelaez, Lubomir D. Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1904–1916, 2015.
- [16] Jeremy Heitz and Daphne Koller. Learning spatial context: Using stuff to find things. In *ECCV*, 2008.
- [17] Matthias Holschneider, Richard Kronland-Martinet, Jean Morlet, and Ph Tchamitchian. A real-time algorithm for signal analysis with the help of the wavelet transform. In *Wavelets*, pages 286–297. 1989.
- [18] Seunghoon Hong, Hyeonwoo Noh, and Bohyung Han. Decoupled deep neural network for semi-supervised semantic segmentation. *CoRR*, abs/1506.04924, 2015.
- [19] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross B. Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *CoRR*, abs/1408.5093, 2014.
- [20] Andrej Karpathy and Fei-Fei Li. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.

- [21] Stefan Kluckner, Thomas Mauthner, Peter M. Roth, and Horst Bischof. Semantic classification in aerial imagery by integrating appearance and height information. In *ACCV*, 2009.
- [22] Pushmeet Kohli, Lubor Ladicky, and Philip H. S. Torr. Robust higher order potentials for enforcing label consistency. *International Journal of Computer Vision*, 82(3):302–324, 2009.
- [23] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [25] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 2001.
- [26] Yann LeCun, Bernhard E. Boser, John S. Denker, Donnie Henderson, R. E. Howard, Wayne E. Hubbard, and Lawrence D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [27] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [28] Guosheng Lin, Chunhua Shen, Ian D. Reid, and Anton van den Hengel. Efficient piecewise training of deep structured models for semantic segmentation. *CoRR*, abs/1504.01013, 2015.
- [29] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, 2014.
- [30] Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Semantic image segmentation via deep parsing network. *CoRR*, abs/1509.02634, 2015.
- [31] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [32] Volodymyr Mnih and Geoffrey E. Hinton. Learning to detect roads in high-resolution aerial images. In *ECCV*, 2010.
- [33] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan L. Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, 2014.
- [34] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. *CoRR*, abs/1505.04366, 2015.
- [35] George Papandreou, Liang-Chieh Chen, Kevin Murphy, and Alan L. Yuille. Weakly- and semi-supervised learning of a DCNN for semantic image segmentation. *CoRR*, abs/1502.02734, 2015.

- [36] Andrew Rabinovich, Andrea Vedaldi, Carolina Galleguillos, Eric Wiewiora, and Serge J. Belongie. Objects in context. In *ICCV*, 2007.
- [37] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015.
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [39] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [40] Olga Russakovsky, Amy L. Bearman, Vittorio Ferrari, and Fei-Fei Li. What’s the point: Semantic segmentation with point supervision. *CoRR*, abs/1506.02106, 2015.
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [42] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR*, abs/1312.6229, 2013.
- [43] Mark J. Shensa. The discrete wavelet transform: wedding the a trous and mallat algorithms. *IEEE Transactions on Signal Processing*, 40(10):2464–2482, 1992.
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [45] Paul Sturgess, Karteek Alahari, Lubor Ladicky, and Philip H. S. Torr. Combining appearance and structure from motion features for road scene understanding. In *BMVC*, 2009.
- [46] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014.
- [47] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.
- [48] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [49] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *CoRR*, abs/1511.07122, 2015.
- [50] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014.

- [51] Matthew D. Zeiler, Dilip Krishnan, Graham W. Taylor, and Robert Fergus. Deconvolutional networks. In *CVPR*, 2010.
- [52] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip H. S. Torr. Conditional random fields as recurrent neural networks. *CoRR*, abs/1502.03240, 2015.