

# Learning Short-Cut Connections for Object Counting

Daniel Oñoro-Rubio<sup>1</sup>  
daniel.onoro@neclab.eu

Mathias Niepert<sup>1</sup>  
mathias.niepert@neclab.eu

Roberto J. López-Sastre<sup>2</sup>  
robertoj.lopez@uah.es

<sup>1</sup> SysML,  
NEC Lab Europe,  
Heidelberg, Germany

<sup>2</sup> GRAM,  
University of Alcalá,  
Alcalá de Henares, Spain

---

## Abstract

Object counting is an important task in computer vision due to its growing demand in applications such as traffic monitoring or surveillance. In this paper, we consider object counting as a learning problem of a joint feature extraction and pixel-wise object density estimation with Convolutional-Deconvolutional networks. We introduce a novel counting model, named Gated U-Net (GU-Net). Specifically, we propose to enrich the U-Net architecture with the concept of *learnable* short-cut connections. Standard short-cut connections are connections between layers in deep neural networks which skip at least one intermediate layer. Instead of simply setting short-cut connections, we propose to learn these connections from data. Therefore, our short-cuts can work as gating units, which optimize the flow of information between convolutional and deconvolutional layers in the U-Net architecture. We evaluate the introduced GU-Net architecture on three commonly used benchmark data sets for object counting. GU-Nets consistently outperform the base U-Net architecture, and achieve state-of-the-art performance.

## 1 Introduction

Counting objects in images is an important problem with numerous applications. Counting people in crowds [17], cells in a microscopy image [62], or vehicles in a traffic jam [13] are typical instances of this problem. Computer vision systems that automate these tasks have the potential to reduce costs and processing time of otherwise labor intensive tasks.

An intuitive approach to object counting is through object detection [8, 9, 42]. The accuracy of object localization approaches, however, rapidly deteriorates as the density of the objects increases. Thus, methods that approach the counting problem as one of object density estimation, following the pioneering work of Lempitsky *et al.* [20], have been shown to achieve state-of-the-art results [3, 27, 45, 47].

With this paper, we introduce a deep learning model that generates object density maps for counting problems. Figure 1 illustrates the architecture which we refer to as the Gated U-Net (GU-Net). We propose the use of *learnable* short-cut connections in the context of neural architectures with deconvolutional layers, such as the U-Net [33] and V-net [26]. Short-cut connections are connections between layers that skip at least one intermediate layer. There

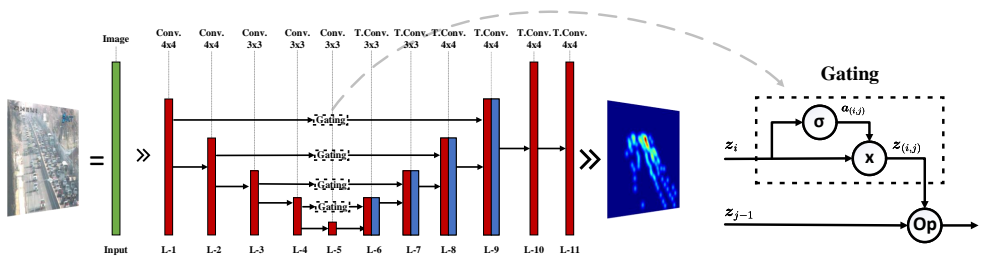


Figure 1: Proposed Gated U-Net Network architecture, where each gating unit is a *learnable* short-cut connection between convolutional and deconvolutional layers.

is an increasing body of work on short-cut connections (sometimes also referred to as skip-connections) in deep neural networks. Examples of architectures with short-cut connections are: highway [67] and residual networks [15], DenseNets [16], and U-Nets [63]. While a clear benefit of these connections has been shown repeatedly, the exact reasons for their success are still not fully understood. Short-cut connections allow one to train very deep networks by facilitating a more direct flow of information between far apart layers during backpropagation [15]. A more recent and complementary explanation is that networks with short-cut connections behave similar to ensembles of shallower networks [69].

Instead of simply setting short-cut connections, in this paper, we propose to *learn* these connections from data. Therefore, our learnable short-cut connection can work as a gating unit, which optimizes the flow of information between convolutional and deconvolutional layers (see again Figure 1).

The main contributions of our work are as follows:

- We first introduce, in Section 3, our learnable short-cut connections, *i.e.* the Gating units, as a general bypass mechanism applicable to any neural network.
- In Section 4, we explore the U-Net architecture [63], a popular model for image segmentation problems, as a reference approach for object counting. We extend the U-Net model by integrating the proposed Gating units. The resulting GU-Net architecture is able to learn which feature maps are beneficial for deconvolutional layers, and also the extent to which feature maps should be passed to later layers.
- We conduct an extensive set of experiments (Section 5) that explore the benefits of the proposed model for the problems of crowd and vehicles counting.

## 2 Related Work

In recent years, the benefits of short-cut connections in deep neural networks have been demonstrated in various empirical studies [6, 14, 15, 16, 65, 64]. Initial research focuses on augmenting top layers with feature maps from bottom layers, where the RCNN framework [17] is modified by connecting the last convolutional layers of the VGG16 architecture using a concatenation operation [6]. In later works, the importance of short-cut connections for back-propagating the gradients during training, and its mitigating effect on the vanishing gradient problem, have been emphasized. [67] propose “high way” networks which, inspired by LSTMs [10], have an architecture where the input of a layer is the sum of the gated input

of the previous layer and the gated output of the previous layer. Following a similar strategy, residual networks [15] and DenseNets [16] take advantage of multiple shortcut connections.

More recent work introduces a gating mechanism that operates directly on the short-cut connections [9]. The proposed gated refinement network predicts the ground truth at multiple scales, where a superior scale is predicted based on the concatenation of the prediction of the previous scale, and the output of its corresponding gating unit. On their gating units, the feature maps of two consecutive scales are combined creating a totally new feature map that is used to refine the prediction.

In contrast to previous work, we present a novel self-gating mechanism that performs a spatial attention over the information forwarded through the short-cut connections. The presented method improves the results on models with a single output, therefore further refinement is still allowed afterward.

We refer the reader to the survey of prior work on object counting of Loy *et al.* [24]. In summary, different type of counting techniques have been applied such as counting by detection [7, 8, 9, 19, 21, 28, 40, 41], counting by clustering [51, 58], and counting by regression [3, 6, 10, 20, 29, 52, 45]. At this point in time, the best-performing methods belong to the group of counting by regression which was introduced by Lempitsky *et al.* [20]. There a method for learning a linear mapping from local image features to object density maps is introduced. Fiaschi *et al.* [10] use a random forest trained with the aid of a collection of multi-spectral filters. In [45] they propose 5 layers convolutional neural network trained with a switching loss function. The work [47] proposes a neural that combines in parallel multiple filter sizes. In [9] they present a framework where a neural network analyzes the image by region and selects an existing model to estimate the output. In the work of [46] they divide the image, and they train a deep neural network that contains a regressor for each region. In [56] they present a system that combines local and global features by aggregating multiple neural networks trained with a reconstruction loss and an adversarial loss. In Yuhong *et al.* [22] they extend the VGG16 with the dilated convolution [24] achieving the top performance of the state of the art. With the presented work we proposed a general gating mechanism that can be easily implemented on different models. In our experiments we show how it improves the performance of our U-Net baseline model. Moreover, the proposed models achieve the state of the art performance without requiring a pretraining on large datasets such as ImageNet and with just a fraction of the parameters of the top performing methods.

### 3 Learnable Short-Cut Connections

Short-cut connections are links between layers in deep neural networks (DNN) which skip at least one intermediate layer. We here introduce a framework for fine-grained learning of the strength of those connections. More specifically, we propose a gating mechanism that can be applied to any deep neural network, to learn short-cuts from the training data, hence optimizing the flow of information between earlier and later layers of the architecture.

We define a DNN as a function approximator that maps feature tensors to output tensors,

$$\mathbf{y} = F_W(\mathbf{x}), \quad (1)$$

where  $W$  represents the parameters of the model.

A DNN  $F_W(\mathbf{x})$  can be written as a composition of functions  $F_W(\mathbf{x}) = f_{W_n}^n \circ \dots \circ f_{W_2}^2 \circ f_{W_1}^1(\mathbf{x})$ , where each  $f_{W_i}^i$  represents the transformation function for layer  $i \in \{1, \dots, n\}$  with its parameters  $W_i$ . Each such layer maps its input to an output tensor  $\mathbf{z}_i$ . Typically, the output  $\mathbf{z}_i$  of

layer  $i$  is used as input to the proceeding layer  $f_{W_{i+1}}^{i+1}(\mathbf{z}_i)$  to generate output  $\mathbf{z}_{i+1}$ , and so on. Each layer, therefore, is only connected locally to its preceding and proceeding layers.

However, we can also connect the output of a layer  $i$  to the input of a layer  $j$ , being  $j > i + 1$ , that is, we can create *short-cut connections* that skip a number of layers in the network. The first class of network architectures popularizing these skip-connections have been introduced with the ResNet model class [15]. Here, we do not only create certain short-cut connections, but *learn* their connection strength using a *gated short-cut unit*, see Figure 1. The short-cut units determine the amount of information which is passed to other layers, and also the ways in which this information is combined with the input of these later layers.

Whenever a layer  $i$  with transformation function  $f_{W_i}^i$  is connected to a layer  $j$ , with  $j > i + 1$ , we introduce a new convolutional layer  $g^{(i,j)}$ , whose hyper parameters are identical to that of layer  $i$ . With  $\mathbf{z}_i$  being the output of layer  $i$ , we compute

$$\mathbf{a}_{(i,j)} = \sigma(g^{(i,j)}(\mathbf{z}_i)), \quad (2)$$

where  $\sigma$  is the sigmoid function. Tensor  $\mathbf{a}_{(i,j)}$  consists of scalars between 0 and 1 that determine the amount of information passed through to layer  $j$ , for each of the local convolutions applied to  $\mathbf{z}_i$ . We then compute the element-wise product of the output of layer  $i$  and the tensor  $\mathbf{a}_{(i,j)}$ , as follows,

$$\mathbf{z}_{(i,j)} = \mathbf{z}_i \odot \mathbf{a}_{(i,j)}, \quad (3)$$

where  $\mathbf{a}_{(i,j)}$  are again the values of the gating function in charge of deciding how much information is let through.

Figure 1 illustrates the gating unit.  $\mathbf{z}_i$  is used as input to a convolutional layer to compute the gating scores  $\mathbf{a}_{(i,j)}$ . A point-wise multiplication is performed to obtain  $\mathbf{z}_{(i,j)}$ . These gated features are then combined with the input  $\mathbf{z}_{j-1}$  of layer  $j$  by using an operation  $\circ$  (sum, point-wise multiplication, or concatenation) between  $\mathbf{z}_{(i,j)}$  and  $\mathbf{z}_{j-1}$ . The resulting tensor  $\mathbf{z}_{(i,j)} \circ \mathbf{z}_{j-1}$  is then used as the input for layer  $j$ .

## 4 GU-Net: Gated Short-Cuts for Object Counting

A class of network architecture with standard short-cut connections is the U-Net [13]. U-Nets are fully convolutional neural networks which can be divided into a compression (convolution) and a construction (deconvolution) part. Both parts are connected with short-cut connections leading from a compression layer to a construction layer with the same size. We chose U-Nets since they have shown good performance on several data sets for image segmentation and since they contain short-cut connections which we aim to learn.

Figure 1 illustrates the DNN architecture we propose for the object counting problem. With the exception of the gating units, which we introduce with this paper, we have a standard U-Net architecture [13]. The U-Net model consists of 11 layers. The first five layers are convolutional layers that comprise the encoding of the input images. The following five layers are transpose-convolutional layers [23] that perform the reconstruction (or decoding) of the bottleneck representation. In all convolutional and deconvolutional layers we use a stride of [2, 2] and we make use of the Leaky ReLu [25] as activation function. Finally, the eleventh layer is a convolution operation with a single filter that produces the output.

For our experiments, we have setup on the receptive field of our models to cover an area of [96, 96] pixels. Since we test our algorithm for the counting problem, we find out that an area of [96, 96] pixels can hold a big object surrounded by background. The size of the filter

is chosen to be  $[4, 4]$  for those features maps that are greater or equal to  $[24, 24]$ , while for those that are smaller we use a slightly reduced size of  $[3, 3]$ .

In this work, we tackle the problem of object counting with a deep network that is trained to learn how to map the appearance of the images to their corresponding object density maps. Therefore, the goal of our network is to generate image-wise density maps, that can be later integrated to get the total object count which is presented in the image. Therefore, we propose to solve a multivariate regression task, where the network is trained to minimize the squared difference between the prediction and the ground truth. Therefore, we define our loss function as:  $\mathcal{L}(\mathbf{y}, \mathbf{x}) = \|\mathbf{y} - F_W(\mathbf{x})\|_2$ . In contrast to most previous works, we can achieve state-of-the-art performance with just a single loss term.

In the GU-Net, the gating units act as bit-wise soft masks that determine the amount of information to pass to the respective layers, so as to improve the quality of the final feature maps. In a deep neural network, the first layers are specialized in detecting low-level features such as edges, textures, colors, etc. These low-level features are needed in a normal feed-forward neural network, but when they are combined with deeper features, they might or might not contribute to the overall performance. For this reason, the gating units are specially useful to automate the feature selection mechanism for improving the *short-cuts* connections, while strongly back propagating the updates to the early layers during learning.

The main limitation of the presented idea is that the gating units do not add more freedom to the general network but add additional parameters (one additional convolutional layer per short-cut connection). Therefore, the fitting capability of GU-Net is the same as the original U-Net. However, the gating strategy leads to more robust models that produce better results.

## 5 Experiments

We have conducted experiments on three publicly available object counting data sets: TRANCOS [13], SHANGHAITECH [14], and UCSD [6]. We perform a detailed comparison with state-of-the-art object counting methods. Moreover, we provide empirical insights into the advantages of the proposed learnable short-cut connections.

### 5.1 General Setup

We use Tensorflow [15] to implement the proposed models. To train our models, we initialize all the weights with samples from a normal distribution with mean zero and standard deviation 0.02. We make use of the  $L2$  regularizer with a scale value of  $2.5 \times 10^{-5}$  in all layers. To perform gradient decent we employ Adam [16] with a learning rate of  $10^{-4}$ ,  $\beta_1$  of 0.9, and  $\beta_2$  of 0.999. We train our solutions for  $2 \times 10^5$  iterations, with mini-batches consisting of 128 image patches, which are extracted from a random image of the training set, such that 50% of the patches contain a centered object and the remaining patches are randomly sampled from the image. We perform data augmentation by flipping images horizontally with a probability of 0.5. All the pixel values from all channels are scaled to the range  $[0, 1]$ . The ground truth of each dataset is generated by placing a normal distribution on top of each of the annotated objects in the image. The standard deviation  $\sigma$  of the normal distribution varies depending on the data set under consideration.

To perform the object count we feed entire images to our models. The proposed models are fully convolutional neural networks, therefore they are not constrained to a fixed input size. Finally, note that the presented models have significantly fewer parameters than most

Table 1: Fusion operation comparison for the GU-Net model in the TRANCOS dataset.

Model	GAME 0	GAME 1	GAME 2	GAME 3
Ew. Multiplication	6.19	7.24	8.64	10.51
Summation	4.81	6.09	7.63	9.60
Concatenation	<b>4.44</b>	<b>5.84</b>	<b>7.34</b>	<b>9.29</b>

of the state-of-the-art methods. To analyze an image of size  $[640, 480, 3]$  with a GU-Net architecture, takes on average only 23ms when executed in a NVIDIA GeForce 1080 Ti.

We use mean absolute error (*MAE*) and mean squared error (*MSE*) for the evaluation of the results on SHANGHAITECH [47] and UCSD [6]:

$$MAE = \frac{1}{N} \sum_{n=1}^N |c(F_W(\mathbf{x})) - c(\mathbf{y})|, \quad (4) \quad MSE = \sqrt{\frac{1}{N} \sum_{n=1}^N (c(F_W(\mathbf{x})) - c(\mathbf{y}))^2}. \quad (5)$$

$N$  is the total number of images,  $c(F_W(\mathbf{x}))$  represents the object count based on the density map output of the neural network, and  $c(\mathbf{y})$  is the count based on the ground truth density map.

The Grid Average Mean absolute Error (GAME) metric is used to perform the evaluation with TRANCOS [13]. According to this metric, each image has to be divided into a set of  $4^s$  non-overlapping rectangles. When  $s = 0$  it is the entire image, when  $s = 1$  it is the four quadrants, and so on. For a specific value of  $s$ , the GAME( $s$ ) is computed as the average of the MAE in each of the corresponding  $4^s$  subregions. Let  $\mathbf{R}_s$  be the set of image regions resulting from dividing the image into  $4^s$  non-overlapping rectangles. Then, the GAME( $s$ ) is obtained as follows,

$$GAME(s) = \frac{1}{N} \sum_{n=1}^N \left( \frac{1}{|\mathbf{R}_s|} \sum_{\mathbf{x} \in \mathbf{R}_s} |c(F_W(\mathbf{x})) - c(\mathbf{y})| \right). \quad (6)$$

This way, this metric provides a spatial measure of the error for  $s > 0$ . Note that GAME(0) is simply the Mean Absolute Error (MAE).

## 5.2 Counting vehicles

TRANCOS is a publicly available dataset consisting of images depicting traffic jams in various road scenarios, and under multiple lighting conditions and different perspectives. It provides 1.244 images obtained from video cameras where a total of 46.796 vehicles have been annotated. It also includes a region of interest (ROI) per image.

For the experiments using this dataset, we follow the experimental setup of [13] to generate the ground truth density maps by setting the standard deviation of the normal distributions to  $\sigma = 10$ . We also perform a random gamma transformation to the images.

We first use this dataset to evaluate the performance of the different possible operations (element-wise multiplication or summation, as well as a simple concatenation) for the operator  $\circ$  of our learnable short-cut connections. Table 1 shows the results for the GU-Net architecture and all these operations. We observe that the element-wise multiplication, even though it produces accurate density maps, does not offer the best results. The summation is a linear operation that shows satisfying results. It also has the advantage that it does not add channels to the resulting features, making this type of operation suitable for situations in

Table 2: Results comparison for TRANCOS dataset.

TRANCOS				
Model	GAME 0	GAME 1	GAME 2	GAME 3
Reg. Forest [14]	17.77	20.14	23.65	25.99
MESA [20]	13.76	16.72	20.72	24.36
Hydra 3s [27]	10.99	13.75	16.69	19.32
FCN-ST [8]	5.47	-	-	-
CSRNet [22]	<b>3.56</b>	<b>5.49</b>	8.57	15.04
U-Net	4.58	6.69	8.69	10.83
GU-Net	4.44	5.84	<b>7.34</b>	<b>9.29</b>

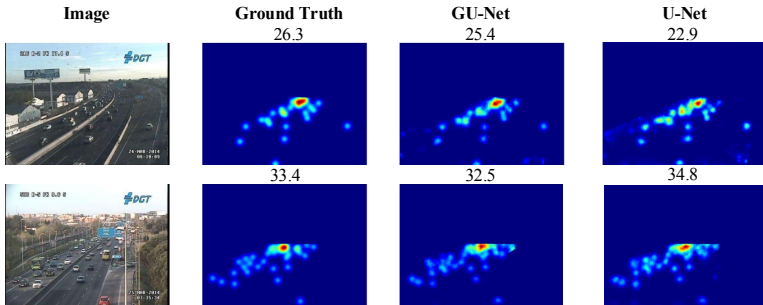


Figure 2: TRANCOS qualitative results.

which memory consumption and processing capabilities are constrained. The best performing operation is the concatenation, hence, for the remainder of our experiments, we use it as the merging operation of the short-cut connections.

We now compare in Table 2 our solutions with the state-of-the-art in the TRANCOS dataset. Shallow methods are listed in the upper part of the table. In the second part we find the deep learning methods. The last two rows show the results obtained by the U-Net and the GU-Net models. First, our GU-Net achieves better results than the U-Net base architecture. The improvements are consistent across the different GAME settings used. This validates our hypothesis about the proposed learnable gating units. Second, in our analysis one can appreciate that the shallow methods are outperformed by deep learning based approaches. Among the deep models, the recent work of [22] gets the best performance for GAME(0) (or MAE). However, it reports high errors for GAME(2) and GAME(3). This indicates that the model is probably under-counting and over-counting in different regions of the image, but the count over the whole image is compensated. Overall, the best vehicle counting performance for both GAME(2) and GAME(3) is obtained by our GU-Net. Figure 2 shows a few qualitative results of the proposed models.

Adding learnable short-cut units helps to determine the strength of the information that is passed to the later construction layers. It actually increases the robustness of the final model by blocking information that is not relevant for these layers. Intuitively, in a deep convolutional neural network, the first layers learn low-level features such as edges, textures, or colors. The later layers are able to capture more complex feature patterns such as eyes, legs, etc. Therefore, when low-level features are combined with high-level features, the resulting feature map might contain irrelevant information, potentially adding noise. As an example, in Figure 3(a) we present a situation that clearly shows how a low level texture is adding



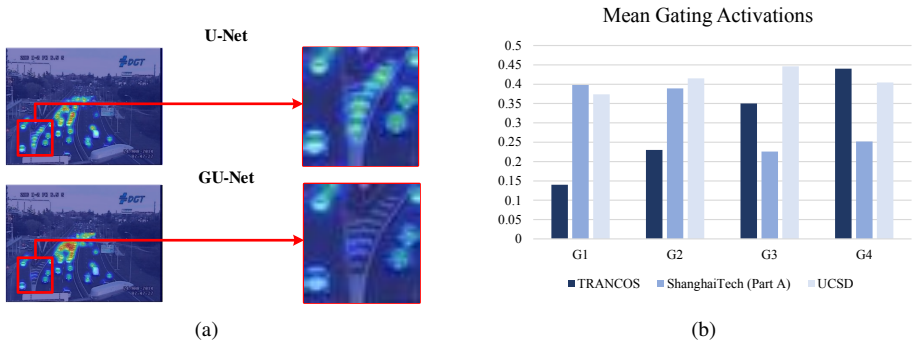


Figure 3: a) Qualitative error analysis between U-Net and GU-Net. (b) Mean activation scores of the gating units of the analyzed datasets.

noise. The U-Net is confusing the road lines as cars while the GU-Net can efficiently handle the problem. Our learnable short-cut units learn to detect these situation and effectively block the forwarded information through the short-cut connections. To explore this hypothesis more thoroughly, we measure the mean values of the activation scalars (the output of the sigmoids) of the four learnable short-cut connections in the GU-Net for several data sets. Figure 3(b) depicts the results of this experiment. It clearly shows that the effect of the short-cut units is data dependent, that is, the learnable short-cut units automatically adapt to the data set under consideration.

### 5.3 Crowd counting

In addition to the previous experiments, we also include a detailed comparison of the U-Net and GU-Net architectures using SHANGHAITECH [47] and UCSD [9], two publicly available data sets for the problem of crowd counting.

SHANGHAITECH [47] contains 1.198 annotated images with a total of 330.165 persons. The dataset consists of two parts: Part A contains images randomly crawled on the Internet, and Part B is made of images taken from the metropolitan areas of Shanghai. Both sets are divided into training and testing subsets, where: i) Part A contains 300 training images and 182 test images; ii) Part B consists of 400 training images and 316 test images. In addition to the standard sets, we have created our own validation set for each part, by randomly taking 50 images out of the training sets. We resize the images to have a maximum height or width of 380 pixels, and we generate the ground truth density maps by placing a Gaussian kernel on each annotated position with a standard derivation of  $\sigma = 4$ . We follow the training procedure described in section 5.1.

Table 3(a) shows the results obtained for this dataset. First, one can again observe how the GU-Net improves the results of the standard U-Net baseline approach. In the Part A of the dataset, the CSRNet [22] approach gets the best performance. It is followed by the CP-CNN [36] model. However, in Part B, CSRNet reports the lowest error, but it is followed by our models. Note that both CP-CNN and CSRNet approaches are based on or include the VGG16 architecture. While VGG16 has around 16.8M of learnable parameters and is pretrained on Imagenet, our most complex model, *i.e.* GU-Net, has only 2.8M parameters and is not pretrained on any external dataset.



Table 3: Crowd counting results for SHANGHAI TECH and UCSD datasets.

(a) SHANGHAI TECH state of the art.					(b) UCSD state of the art.		
SHANGHAI TECH					UCSD		
Model	Part A		Part B		Model	MAE	MSE
	MAE	MSE	MAE	MSE			
Cross-scene [15]	181.8	277.7	32.0	49.8	FCN-MT [16]	1.67	3.41
MCNN [17]	110.2	173.2	26.4	41.3	Count-forest [18]	1.61	4.4
CP-CNN [19]	73.6	<b>106.4</b>	20.1	30.1	Cross-scene [15]	1.6	5.97
CSRNet [20]	<b>68.2</b>	115.0	<b>10.6</b>	<b>16.0</b>	CSRNet [20]	1.16	1.47
U-Net	104.9	173.3	17.1	25.8	MCNN [17]	<b>1.07</b>	<b>1.35</b>
GU-Net	101.4	167.8	14.7	23.3	U-Net	1.28	1.57
					GU-Net	1.25	1.59

UCSD is a standard benchmark in the crowd counting community. It provides 2000 video frames recorded by a camera capturing a single scene. The images have been annotated with a dot on each visible pedestrian. As in previous work [6], we train the models on frames 601 to 1400. The remaining frames are used for testing. For our experiments we sample 100 frames uniformly at random from the set of training, and use them as our validation subset. To generate the ground truth density maps, we set the standard deviation of the normal distributions placed on the objects to  $\sigma = 5$ . To train our models, we follow the procedure described in section 5.1.

Table 3(b) shows the results for the UCSD data set. Without hyper-parameter tuning, we are able to report state of the art results. Moreover, we observe how the GU-Net model consistently improves the U-Net. Our approaches are outperformed by CSRNet [20] and MCNN [17] models. However, note that our models are not pretrained, and that they are lighter in terms of parameters.

## 6 Conclusion

In this paper, we have introduced the concept of learnable gating unit for convolutional-deconvolutional deep networks, like U-Net. Technically, our gating units work as short-cut connections, which can be directly learned from the training data, in order to optimize the flow of information between the layers they connect.

We demonstrate the effectiveness of this novel gating units for the problem of object counting. For doing so, we introduce a novel counting model named GU-Net, which is a U-Net architecture enriched with the learnable short-cut connections proposed. We have evaluated our solutions with three commonly used data sets for object counting. According to our results, GU-Net consistently outperforms the base U-Net architecture, and achieves a state-of-the-art performance.

## Acknowledgments

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 761508, and the project PREPEATE, with reference TEC2016-80326-R, of the Spanish Ministry of Economy, Industry and Competitiveness. We acknowledge the support of NVIDIA Corporation with the donation of a GPU. Cloud computing resources were provided through a Microsoft Azure for Research Award.

## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [2] Md Amirul Islam, Mrigank Rochan, Neil D. B. Bruce, and Yang Wang. Gated feedback refinement network for dense image labeling. In *CVPR*, July 2017.
- [3] C. Arteta, V. Lempitsky, J. A. Noble, and A. Zisserman. Interactive object counting. In *ECCV*, 2014.
- [4] Deepak Babu Sam, Shiv Surya, and R. Venkatesh Babu. Switching convolutional neural network for crowd counting. In *CVPR*, July 2017.
- [5] Sean Bell, C. Lawrence Zitnick, Kavita Bala, and Ross Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [6] A.-B. Chan, Z.-S.-J. Liang, and N. Vasconcelos. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *CVPR*, 2008.
- [7] Sheng Chen, Alan Fern, and Sinisa Todorovic. Person count localization in videos from noisy foreground and detections. In *CVPR*, 2015.
- [8] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
- [9] Pedro F. Felzenszwalb, Ross B. Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2010.
- [10] Luca Fiaschi, Ullrich Köthe, Rahul Nair, and Fred A. Hamprecht. Learning to count with regression forest and structured labels. In *ICPR*, 2012.
- [11] Felix A. Gers, Jürgen A. Schmidhuber, and Fred A. Cummins. Learning to forget: Continual prediction with LSTM. *Neural Comput.*, 12(10):2451–2471, October 2000.
- [12] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [13] R. Guerrero-Gómez-Olmedo, B. Torre-Jiménez, R. J. López-Sastre, S. Maldonado-Bascón, and D. Oñoro-Rubio. Extremely overlapping vehicle counting. In *Iberian Conference on Pattern Recognition and Image Analysis (IbPRIA)*, 2015.

- [14] Bharath Hariharan, Pablo Arbelaez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, June 2015.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [16] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [17] Haroon Idrees, Imran Saleemi, Cody Seibert, and Mubarak Shah. Multi-source multi-scale counting in extremely dense crowd images. In *CVPR*, 2013.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 2015.
- [19] Bastian Leibe, Edgar Seemann, and Bernt Schiele. Pedestrian detection in crowded scenes. In *CVPR*, 2005.
- [20] V. Lempitsky and A. Zisserman. Learning to count objects in images. In *NIPS*, 2010.
- [21] Min Li, Zhaoxiang Zhang, Kaiqi Huang, and Tieniu Tan. Estimating the number of people in crowded scenes by mid based foreground segmentation and head-shoulder detection. In *ICPR*, 2008.
- [22] Yuhong Li, Xiaofan Zhang, and Deming Chen. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In *CVPR*, 2018.
- [23] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, June 2015.
- [24] ChenChange Loy, Ke Chen, Shaogang Gong, and Tao Xiang. Crowd counting and profiling: Methodology and evaluation. In *Modeling, Simulation and Visual Analysis of Crowds*. 2013.
- [25] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *in ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- [26] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *CoRR*, abs/1606.04797, 2016.
- [27] Daniel Oñoro-Rubio and Roberto J. López-Sastre. Towards perspective-free object counting with deep learning. In *ECCV*, 2016.
- [28] Michael Patzold, Ruben Heras Evangelio, and Thomas Sikora. Counting people in crowded environments by fusion of shape and motion information. In *AVSS*, 2010.
- [29] V-Q Pham, T. Kozakaya, O. Yamaguchi, and R. Okada. COUNT forest: CO-voting uncertain number of targets using random forest for crowd density estimation. In *ICCV*, 2015.

- [30] Viet-Quoc Pham, Tatsuo Kozakaya, Osamu Yamaguchi, and Ryuzo Okada. Count forest: Co-voting uncertain number of targets using random forest for crowd density estimation. In *ICCV*, December 2015.
- [31] Vincent Rabaud and Serge Belongie. Counting crowded moving objects. In *CVPR*, 2006.
- [32] Mikel Rodriguez, Ivan Laptev, Josef Sivic, and Jean-Yves Audibert. Density-aware person detection and tracking in crowds. In *ICCV*, 2011.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015.
- [34] J. Selinummi, J. Seppälä, O. Yli-Harja, and J.A. Puhakka. Software for quantification of labeled bacteria from digital microscope images by automated image analysis. *BioTechniques*, 39(6):859–862, 2005. ISSN 0736-6205. doi: 10.2144/000112018.
- [35] Pierre Sermanet, Koray Kavukcuoglu, Soumith Chintala, and Yann Lecun. Pedestrian detection with unsupervised multi-stage feature learning. In *CVPR*, June 2013.
- [36] Vishwanath A. Sindagi and Vishal M. Patel. Generating high-quality crowd density maps using contextual pyramid cnns. In *ICCV*, 2017.
- [37] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Training very deep networks. In *NIPS*, 2015.
- [38] Peter H. Tu, Thomas Sebastian, Gianfranco Doretto, Nils Krahnstoeber, Jens Rittscher, and Ting Yu. Unified crowd segmentation. In *ECCV*. 2008.
- [39] Andreas Veit, Michael Wilber, and Serge Belongie. Residual networks behave like ensembles of relatively shallow networks. *NIPS*, 2016.
- [40] Paul Viola and Michael J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 2004.
- [41] Meng Wang and Xiaogang Wang. Automatic adaptation of a generic pedestrian detector to a specific traffic scene. In *CVPR*, 2011.
- [42] B. Wu and R. Nevatia. Detection of multiple, partially occluded humans in a single image by bayesian combination of edgelet part detectors. In *ICCV*, 2005.
- [43] Songfan Yang and Deva Ramanan. Multi-scale recognition with DAG-CNNs. In *ICCV*, December 2015.
- [44] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2016.
- [45] Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. Cross-scene crowd counting via deep convolutional neural networks. In *CVPR*, June 2015.
- [46] Shanghang Zhang, Guanhang Wu, Joao P. Costeira, and Jose M. F. Moura. Understanding traffic density from large-scale web camera data. In *CVPR*, July 2017.
- [47] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*, June 2016.