

Convolutional Random Walk Networks for Semantic Image Segmentation

Gedas Bertasius¹, Lorenzo Torresani², Stella X. Yu³, Jianbo Shi¹

¹University of Pennsylvania, ²Dartmouth College, ³UC Berkeley ICSI

{gberta, jshi}@seas.upenn.edu lt@dartmouth.edu stealla.yu@berkeley.edu

Abstract

Most current semantic segmentation methods rely on fully convolutional networks (FCNs). However, their use of large receptive fields and many pooling layers cause low spatial resolution inside the deep layers. This leads to predictions with poor localization around the boundaries. Prior work has attempted to address this issue by post-processing predictions with CRFs or MRFs. But such models often fail to capture semantic relationships between objects, which causes spatially disjoint predictions. To overcome these problems, recent methods integrated CRFs or MRFs into an FCN framework. The downside of these new models is that they have much higher complexity than traditional FCNs, which renders training and testing more challenging.

In this work we introduce a simple, yet effective Convolutional Random Walk Network (RWN) that addresses the issues of poor boundary localization and spatially fragmented predictions with very little increase in model complexity. Our proposed RWN jointly optimizes the objectives of pixelwise affinity and semantic segmentation. It combines these two objectives via a novel random walk layer that enforces consistent spatial grouping in the deep layers of the network. Our RWN is implemented using standard convolution and matrix multiplication. This allows an easy integration into existing FCN frameworks and it enables end-to-end training of the whole network via standard back-propagation. Our implementation of RWN requires just 131 additional parameters compared to the traditional FCNs, and yet it consistently produces an improvement over the FCNs on semantic segmentation and scene labeling.

1. Introduction

Fully convolutional networks (FCNs) were first introduced in [20] where they were shown to yield significant improvements in semantic image segmentation. Adopting the FCN approach, many subsequent methods have achieved even better performance [6, 27, 8, 12, 17, 6, 27, 8, 19, 23]. However, traditional FCN-based methods tend to suffer from several limitations. Large receptive fields in the

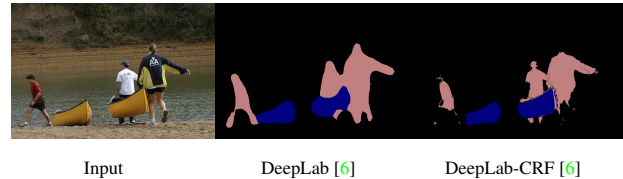


Figure 1: Examples illustrating shortcomings of prior semantic segmentation methods. Segments produced by FCNs are poorly localized around object boundaries, while Dense-CRF produce spatially-disjoint object segments.

convolutional layers and the presence of pooling layers lead to low spatial resolution in the deepest FCN layers. As a result, their predicted segments tend to be blobby and lack fine object boundary details. We report in Fig. 1 some examples illustrating typical poor localization of objects in the outputs of FCNs. Recently, Chen et al. [6] addressed this issue by applying a Dense-CRF post-processing step [15] on top of coarse FCN segmentations. However, such approaches often fail to accurately capture semantic relationships between objects and lead to spatially fragmented segmentations (see an example in the last column of Fig. 1).

To address these problems several recent methods integrated CRFs or MRFs directly into the FCN framework [27, 19, 17, 5]. However, these new models typically involve (1) a large number of parameters, (2) complex loss functions requiring specialized model training or (3) recurrent layers, which make training and testing more complicated. We summarize the most prominent of these approaches and their model complexities in Table 1.

We note that we do not claim that using complex loss functions always makes the model overly complex and too difficult to use. If a complex loss is integrated into an FCN framework such that the FCN can still be trained in a standard fashion, and produce better results than using standard losses, then such a model is beneficial. However, in the context of prior segmentation methods [27, 19, 17, 5], such complex losses often require: 1) modifying the network structure (casting CNN into an RNN) [27, 5], or 2) using a complicated multi-stage learning scheme, where different layers are optimized during a different training stage [19, 17]. Due to such complex training procedures,

	[6]	[5]	[23]	[27]	[19]	RWN
requires post-processing?	✓	✗	✗	✗	✗	✗
uses complex loss?	✗	✗	✗	✓	✓	✗
requires recurrent layers?	✗	✓	✗	✓	✗	✗
model size (in MB)	79	79	961	514	>1000	79

Table 1: Summary of recent semantic segmentation models. For each model, we report whether it requires: (1) CRF post-processing, (2) complex loss functions, or (3) recurrent layers. We also list the size of the model (using the size of Caffe [13] models in MB). We note that unlike prior methods, our RWN does not need post-processing, it is implemented using standard layers and loss functions, and it also has a compact model.

which are adapted for specific tasks and datasets, these models can be quite difficult to adapt for new tasks and datasets, which is disadvantageous.

Inspired by random walk methods [21, 3, 24], in this work we introduce a simple, yet effective alternative to traditional FCNs: a Convolutional Random Walk Network (RWN) that combines the strengths of FCNs and random walk methods. Our model addresses the issues of (1) poor localization around the boundaries suffered by FCNs and (2) spatially disjoint segments produced by dense CRFs. Additionally, unlike recent semantic segmentation approaches [23, 27, 19, 17], our RWN does so without significantly increasing the complexity of the model.

Our proposed RWN jointly optimizes (1) pixelwise affinity and (2) semantic segmentation learning objectives that are linked via a novel random walk layer, which enforces spatial consistency in the deepest layers of the network. The random walk layer is implemented via matrix multiplication. As a result, RWN seamlessly integrates both affinity and segmentation branches, and can be jointly trained end-to-end via standard back-propagation with minimal modifications to the existing FCN framework. Additionally, our implementation of RWN requires only 131 additional parameters. Thus, the effective complexity of our model is the same as the complexity of traditional FCNs (see Table 1). We compare our approach to several variants of the DeepLab semantic segmentation system [6, 7], and show that our proposed RWN consistently produces better performance over these baselines for the tasks of semantic segmentation and scene labeling.

2. Related Work

The recent introduction of fully convolutional networks (FCNs) [20] has led to remarkable advances in semantic segmentation. However, due to the large receptive fields and many pooling layers, segments predicted by FCNs tend to be blobby and lack fine object boundary details. Recently there have been several attempts to address these problems. These approaches can be divided into several groups.

The work in [4, 6, 25, 14, 4] used FCN predictions as unary potentials in a separate globalization model that refines the segments using similarity cues based on regions or boundaries. One disadvantage of these methods is that the learning of the unary potentials and the training of the globalization model are completely disjoint. As a result, these methods often fail to capture semantic relationships between objects, which produces segmentation results that are spatially disjoint (see the right side of Fig. 1).

To address these issues, several recent methods [27, 19, 17] have proposed to integrate a CRF or a MRF into the network, thus enabling end-to-end training of the joint model. However, the merging of these two models leads to a dramatic increase in complexity and number of parameters. For instance, the method in [27], requires to cast the original FCN into a Recurrent Neural Network (RNN), which renders the model much bigger in size (see Table 1). A recent method [5] jointly predicts boundaries and segmentations, and then combines them using a recurrent layer, which also requires complex modifications to the existing FCN framework.

The work in [19] proposes to use **local** convolutional layers, which leads to a significantly larger number of parameters. Similarly, the method in [17] proposes to model unary and pairwise potentials by separate multi-scale branches. This leads to a network with at least twice as many parameters as the traditional FCN and a much more complex multi-stage training procedure.

In addition to the above methods, it is worth mentioning deconvolutional networks [23, 12], which use deconvolution and unpooling layers to recover fine object details from the coarse FCN predictions. However, in order to effectively recover fine details one must employ nearly as many deconvolutional layers as the number of convolutional layers, which yields a large growth in number of parameters (see Table 1).

Unlike these prior methods, our implementation of RWN needs only 131 additional parameters over the base FCN. These additional parameters represent only 0.0008% of the total number of parameters in the network. In addition, our RWN uses standard convolution and matrix multiplication. Thus, it does not need to incorporate complicated loss functions or new complex layers [27, 5, 19, 17]. Finally, unlike the methods in [4, 6, 25, 14] that predict and refine the segmentations disjointly, our RWN model jointly optimizes pixel affinity and semantic segmentation in an end-to-end fashion. Our experiments show that this leads to spatially smoother segmentations.

3. Background

Random Graph Walks. Random walks are one of the most widely known and used methods in graph theory [21]. Most notably, the concept of random walks led to

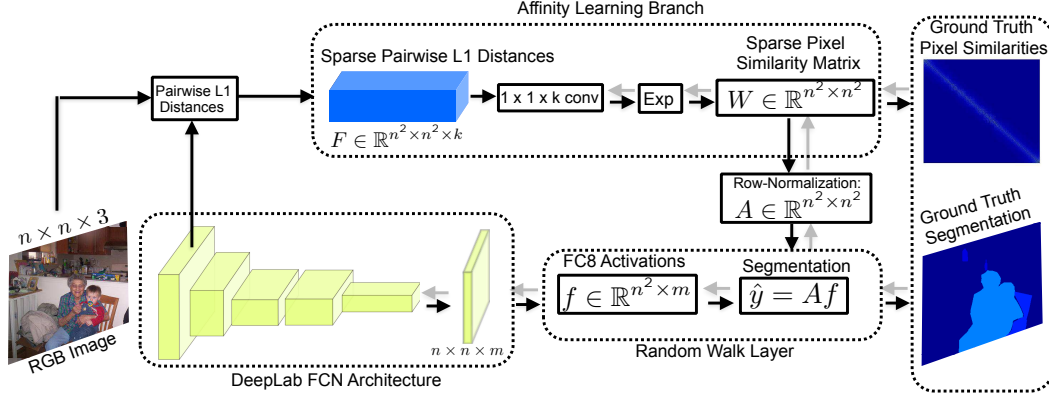


Figure 2: The architecture of our proposed Random Walk Network (RWN) (best viewed in color). Our RWN consists of two branches: (1) one branch devoted to the segmentation predictions, and (2) another branch predicting pixel-level affinities. These two branches are then merged via a novel random walk layer that encourages spatially smooth segmentation predictions. The entire RWN is jointly optimized end-to-end via a standard back-propagation algorithm.

the development of PageRank [24] and Personalized PageRank [3], which are widely used for many applications. Let $G = (V, E)$ denote an undirected graph with a set of vertices V and a set of edges E . Then a random walk in such graph can be characterized by the transition probabilities between its vertices. Let W be a symmetric $n \times n$ affinity matrix, where n denotes the number of nodes in the graph and where $W_{ij} \in [0, 1]$ denotes how similar the nodes i and j are. In the context of a semantic segmentation problem, each pixel in the image can be viewed as a separate node in the graph, where the similarity between two nodes can be evaluated according to some metric (e.g. color or texture similarity etc). Then let D indicate a diagonal $n \times n$ matrix, which stores the degree values for each node: $D_{ii} = \sum_{j=1}^n W_{ij}$ for all j except $i = j$. Then, we can express our random walk transition matrix as $A = D^{-1}W$.

Given this setup, we want to model how the information in the graph spreads if we start at a particular node, and perform a random walk in this graph. Let y_t be a $n \times 1$ vector denoting a node distribution at time t . In the context of the PageRank algorithm, y_t may indicate the rank estimates associated with each of the n Web pages at time t . Then, according to the random walk theory, we can spread the rank information in the graph by performing a one-step random walk. This process can be expressed as $y_{t+1} = Ay_t$, where y_{t+1} denotes a newly obtained rank distribution after one random walk step, the matrix A contains the random walk transition probabilities, and y_t is the rank distribution at time step t . Thus, we can observe that the information among the nodes can be diffused, by simply multiplying the random walk transition probability matrix A , with the rank distribution y_t at a particular time t . This process can be repeated multiple times, until convergence is reached. For a more detailed survey please see [21, 24].

Difference from MRF/CRF Approaches. CRFs and

MRFs have been widely used in structured prediction problems [16]. Recently, CRFs and MRFs have also been integrated into the fully convolutional network framework for semantic segmentation [27, 19, 17]. We want to stress that while the goals of CRF/MRF and random walk methods are the same (i.e. to globally propagate information in the graph structures), the mechanism to achieve this objective is very different in these two approaches. While MRFs and CRFs typically employ graphs with a fixed grid structure (e.g., one where each node is connected to its four closest neighbors), random walk methods are much more flexible and can implement any arbitrary graph structure via the affinity matrix specification. Thus, since our proposed RWN is based on random walks, it can employ any arbitrary graph structure, which can be beneficial as different problems may require different graph structures.

Additionally, to globally propagate information among the nodes, MRFs and CRFs need to employ approximate inference techniques, because exact inference tends to be intractable in graphs with a grid structure. Integrating such approximate inference techniques into the steps of FCN training and prediction can be challenging and may require lots of domain-specific modifications. In comparison, random walk methods globally propagate information among the nodes via a simple matrix multiplication. Not only is the matrix multiplication efficient and exact, but it is also easy to integrate into the traditional FCN framework for both training and prediction schemes. Additionally, due to the use of standard convolution and matrix multiplication operations, our RWN can be trivially trained via standard back-propagation in an end-to-end fashion.

4. Convolutional Random Walk Networks

In this work, our goal is to integrate a random walk process into the FCN architecture to encourage coherent se-

mantic segmentation among pixels that are similar to each other. Such a process introduces an explicit grouping mechanism, which should be beneficial in addressing the issues of (1) poor localization around the boundaries, and (2) spatially fragmented segmentations.

A schematic illustration of our proposed RWN architecture is presented in Fig. 2. Our RWN is a network composed of two branches: (1) one branch that predicts semantic segmentation potentials, and (2) another branch devoted to predicting pixel-level affinities. These two branches are merged via a novel random walk layer that encourages spatially coherent semantic segmentation. The entire RWN can be jointly optimized end-to-end. We now describe each of the components of the RWN architecture in more detail.

4.1. Semantic Segmentation Branch

For the semantic segmentation branch, we present results for several variants of the DeepLab segmentation systems, including DeepLab-LargeFOV [6], DeepLab-attention [7], and DeepLab-v2, which is one of the top performing segmentation systems. DeepLab-largeFOV is a fully convolutional adaptation of the VGG [26] architecture, which contains 16 convolutional layers. DeepLab-attention [7], is a multi-scale VGG based network, for which each multi-scale branch focuses on a specific part of the image. Finally, DeepLab-v2 is a multi-scale network based on the residual network [11] implementation. We note that even though we use a DeepLab architecture in our experiments, other architectures such as [2] and many others could be integrated into our framework.

4.2. Pixel-Level Affinity Branch

To learn the pairwise pixel-level affinities, we employ a separate affinity learning branch with its own learning objective (See Fig. 2). The affinity branch is connected with the input $n \times n \times 3$ RGB image, and low-level *conv1_1* and *conv1_2* layers. The feature maps corresponding to these layers are $n \times n$ in width and height but they have a different number of channels (3, 64, and 64 respectively). Let k be the total number of affinity learning parameters (in our case $k = 3 + 64 + 64 = 131$). Then, let F be a **sparse** $n^2 \times n^2 \times k$ matrix that stores $L1$ distances between each pixel and all of its neighbors within a radius R , according to each channel. Note that the distances are **not** summed up across the k channels, but instead they are computed and stored separately for each channel. The resulting matrix F is then used as an input to the affinity branch, as shown in Figure 2.

The affinity branch consists of a $1 \times 1 \times k$ convolutional layer and an exponential layer. The output of the exponential layer is then attached to the Euclidean loss layer and is optimized to predict the ground truth pixel affinities, which are obtained from the original semantic segmentation an-

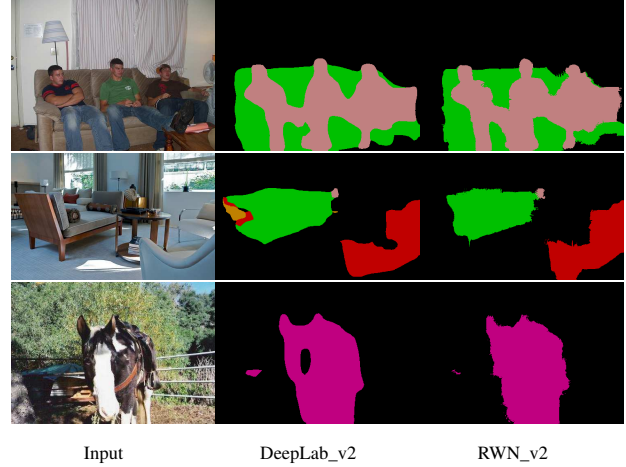


Figure 3: A figure illustrating the segmentation results of our RWN and the DeepLab-v2 network. Note that RWN produced segmentations are spatially smoother and produce less false positive predictions than the DeepLab-v2 system.

notations. Specifically, we set the ground truth affinity between two pixels to 1 if the pixels share the same semantic label and have distance less than R from each other. Note that F , which is used as an input to the affinity branch, is a **sparse** matrix, as only a small fraction of all the entries in F are populated with non-zero values. The rest of the entries are ignored during the computation.

Also note that we only use features from RGB, *conv1_1* and *conv1_2* layers, because they are not affected by pooling, and thus, preserve the original spatial resolution. We also experimented with using features from deeper FCN layers such as *fc6*, and *fc7*. However, we observed that features from deeper layers are highly correlated to the predicted semantic segmentation unary potentials, which causes redundancy and little improvement in the segmentation performance. We also experimented with using more than one convolutional layer in the affinity learning branch, but observed that additional layers provide negligible improvements in accuracy.

4.3. Random Walk Layer

To integrate the semantic segmentation potentials and our learned pixel-level affinities, we introduce a novel random walk layer, which propagates the semantic segmentation information based on the learned affinities. The random walk layer is connected to the two bottom layers: (1) the *fc8* layer containing the semantic segmentation potentials, and (2) the affinity layer that outputs a sparse $n^2 \times n^2$ random walk transition matrix A . Then, let f denote the activation values from the *fc8* layer, reshaped to the dimensions of $n^2 \times m$, where n^2 refers to the number of pixels, and m is the number of object classes in the dataset. A single random walk layer implements one step of the random walk

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean	overall
DeepLab-largeFOV	79.8	71.5	78.9	70.9	72.1	87.9	81.2	85.7	46.9	80.9	56.5	82.6	77.9	79.3	80.1	64.4	77.6	52.7	80.3	70.0	73.8	76.0
RWN-largeFOV	81.6	72.1	82.3	72.0	75.4	89.1	82.5	87.4	49.1	83.6	57.9	84.8	80.7	80.2	81.2	65.7	79.7	55.5	81.5	74.0	75.8	77.9
DeepLab-attention	83.4	76.0	83.0	74.2	77.6	91.6	85.2	89.1	54.4	86.1	62.9	86.7	83.8	84.2	82.4	70.2	84.7	61.0	84.8	77.9	79.0	80.5
RWN-attention	84.7	76.6	85.5	74.0	79.0	92.4	85.6	90.0	55.6	87.4	63.5	88.2	85.0	84.8	83.4	70.1	85.9	62.6	85.1	79.3	79.9	81.5
DeepLab-v2	85.5	50.6	86.9	74.4	82.7	93.1	88.4	91.9	62.1	89.7	71.5	90.3	86.2	86.3	84.6	75.1	87.6	72.2	87.8	81.3	81.4	83.4
RWN-v2	86.0	50.0	88.4	73.5	83.9	93.4	88.6	92.5	63.9	90.9	72.6	90.9	87.3	86.9	85.7	75.0	89.0	74.0	88.1	82.3	82.1	84.3

Table 2: Semantic segmentation results on the SBD dataset according to the per-pixel Intersection over Union evaluation metric. From the results, we observe that our proposed RWN consistently outperforms DeepLab-LargeFOV, DeepLab-attention, and DeepLab-v2 baselines.

process, which can be performed as $\hat{y} = Af$, where \hat{y} indicates the diffused segmentation predictions, and A denotes the random walk transition matrix.

The random walk layer is then attached to the softmax loss layer, and is optimized to predict ground truth semantic segmentations. One of the advantages of our proposed random walk layer is that it is implemented as a matrix multiplication, which makes it possible to back-propagate the gradients to both (1) the affinity branch and (2) the segmentation branch. Let the softmax-loss gradient be an $n^2 \times m$ matrix $\frac{\partial L}{\partial \hat{y}}$, where n^2 is the number of pixels in the *fc8* layer, and m is the number of predicted object classes. Then the gradients, which are back-propagated to the semantic segmentation branch are computed as $\frac{\partial L}{\partial f} = A^T \frac{\partial L}{\partial \hat{y}}$, where A^T is the transposed random walk transition matrix. Also, the gradients, that are back-propagated to the affinity branch are computed as $\frac{\partial L}{\partial A} = \frac{\partial L}{\partial \hat{y}} f^T$, where f^T is a $m \times n^2$ matrix that contains transposed activation values from the *fc8* layer of the segmentation branch. We note that $\frac{\partial L}{\partial A}$ is a **sparse** $n^2 \times n^2$ matrix, which means that the above matrix multiplication only considers the pixel pairs that correspond to the non-zero entries in the random walk transition matrix A .

4.4. Random Walk Prediction at Testing

In the previous subsection, we mentioned that the prediction in the random walk layer can be done via a simple matrix multiplication operation $\hat{y} = Af$, where A denotes the random walk transition matrix, and f depicts the activation values from the *fc8* layer. Typically, we want to apply multiple random walk steps until convergence is reached. However, we also do not want to deviate too much from our initial segmentation predictions, in case the random walk transition matrix is not perfectly accurate, which is a reasonable expectation. Thus, our prediction scheme needs to balance two effects: (1) propagating the segmentation information across the nodes using a random walk transition matrix, and (2) not deviating too much from the initial segmentation.

This tradeoff between the two quantities is very similar to the idea behind MRF and CRF models, which try to minimize an energy formed by unary and pairwise terms. However, as discussed earlier, MRF and CRF methods tend to use 1) grid structure graphs and 2) various approximate

inference techniques to propagate segmentation information globally. In comparison, our random walk approach is advantageous because it can use 1) any arbitrary graph structure and 2) an exact matrix multiplication operation to achieve the same goal.

Let us first denote our segmentation prediction after $t+1$ random walk steps as \hat{y}^{t+1} . Then our general prediction scheme can be written as:

$$\hat{y}^{t+1} = \alpha A \hat{y}^t + (1 - \alpha) f \quad (1)$$

where α denotes a parameter $[0, 1]$ that controls the tradeoff between (1) diffusing segmentation information along the connections of a random walk transition matrix and (2) not deviating too much from initial segmentation values (i.e. the outputs of the last layer of the FCN). Let us now initialize \hat{y}^0 to contain the output values from the *fc8* layer, which we denoted with f . Then we can write our prediction equation by substituting the recurrent expressions:

$$\hat{y}^{t+1} = (\alpha A)^{t+1} f + (1 - \alpha) \sum_{i=0}^t (\alpha A)^i f \quad (2)$$

Now, because we want to apply our random walk procedure until convergence we set $t = \infty$. Then, because our random walk transition matrix is stochastic we know that $\lim_{t \rightarrow \infty} (\alpha A)^{t+1} = 0$. Furthermore, we can write $S_t = \sum_{i=0}^t A^i = I + A + A^2 + \dots + A^t$, where I is an identity matrix, and where S_t denotes a partial sum of random walk transitions until iteration t . We can then write $S_t - AS_t = I - A^{t+1}$, which implies:

$$(I - A)S_t = I - A^{t+1} \quad (3)$$

From our previous derivation, we already know that $\lim_{t \rightarrow \infty} (A)^{t+1} = 0$, which implies that

$$S_\infty = (I - A)^{-1} \quad (4)$$

Thus, our final prediction equation, which corresponds to applying repeated random walk steps until convergence, can be written as

$$\hat{y}^\infty = (I - \alpha A)^{-1} f \quad (5)$$

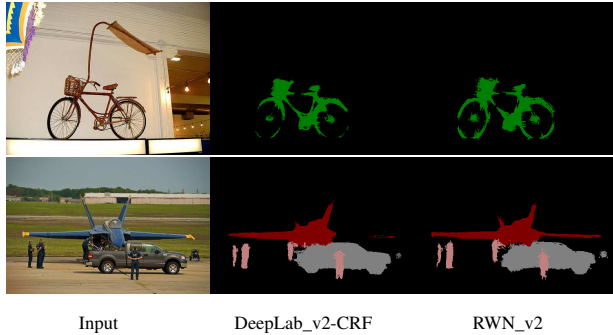


Figure 4: Comparison of segmentation results produced by our RWN versus the DeepLab-v2-CRF system. It can be noticed that, despite not using any post-processing steps, our RWN predicts fine object details (e.g., bike wheels or plane wings) more accurately than DeepLab-v2-CRF, which fails to capture some these object parts.

In practice, the random walk transition matrix A is pretty large, and inverting it is impractical. To deal with this problem, we shrink the matrix $(I - \alpha A)$ using a simple and efficient technique presented in [1], and then invert it to compute the final segmentation. In the experimental section, we show that such a prediction scheme produces solid results and is still pretty efficient (≈ 1 second per image). We also note that we use this prediction scheme only during testing. During training we employ a scheme that uses a single random walk step (but with a much larger radius), which is faster. We explain this procedure in the next subsection.

4.5. Implementation Details

We jointly train our RWN in an end-to-end fashion for 2000 iterations, with a learning rate of 10^{-5} , 0.9 momentum, the weight decay of $5 \cdot 10^{-5}$, and 15 samples per batch. For the RWN model, we set the tradeoff parameter α to 0.01. During testing we set the random walk connectivity radius $R = 5$ and apply the random walk procedure until convergence. However, during training we set $R = 40$, and apply a single random walk step. This training strategy works well because increasing the radius size eliminates the need for multiple random walk steps, which speeds up the training. However, using $R = 5$ and applying an infinite number of random walk steps until convergence still yields slightly better results (see study in 5.4), so we use it during testing. For all of our experiments, we use a Caffe library [13]. During training, we also employ data augmentation techniques such as cropping, and mirroring.

5. Experimental Results

In this section, we present our results for semantic segmentation on the SBD [10] dataset, which contains objects and their per-pixel labels for 20 Pascal VOC classes (excluding the background class). We also include scene label-

Method	mean IOU	overall IOU
DeepLab-largeFOV-CRF	75.7	77.7
RWN-largeFOV	75.8	77.9
DeepLab-attention-CRF	79.9	81.6
RWN-attention	79.9	81.5
DeepLab-v2-CRF	81.9	84.2
RWN-v2	82.1	84.3
DeepLab-DT	76.6	78.7
RWN	76.7	78.8

Table 3: Quantitative comparison between our RWN model and several variants of the DeepLab system that use a dense CRF or a domain-transfer (DT) filter for post-processing. These results suggest that our RWN acts as an effective globalization scheme, since it produces results that are similar or even better than the results achieved by post-processing the DeepLab outputs with a CRF or DT.

ing results on the commonly used Stanford background [9] and Sift Flow [18] datasets. We evaluate our segmentation results on these tasks using the standard metric of the intersection-over-union (IOU) averaged per pixels across all the classes from each dataset. We also include the class-agnostic overall pixel intersection-over-union score, which measures the per-pixel IOU across all classes.

We experiment with several variants of the DeepLab system [6, 7] as our main baselines throughout our experiments: DeepLab-LargeFOV [6], DeepLab-attention [7], and DeepLab-v2.

Our evaluations provide evidence for four conclusions:

- In subsections 5.1, 5.2, we demonstrate that our proposed RWN outperforms DeepLab baselines for both semantic segmentation, and scene labeling tasks.
- In subsection 5.1, we demonstrate that, compared to the dense CRF approaches, RWN predicts segmentations that are spatially smoother.
- In Subsection 5.3, we show that our approach is more efficient than the denseCRF inference.
- Finally, in Subsection, 5.4, we demonstrate that our random walk layer is beneficial and that our model is flexible to use different graph structures.

5.1. Semantic Segmentation Task

Standard Evaluation. In Table 2, we present semantic segmentation results on the Pascal SBD dataset [10], which contains 8055 training and 2857 testing images. These results indicate that RWN consistently outperforms all three of the DeepLab baselines. In Figure 3, we also compare qualitative segmentation results of a DeepLab-v2 network and our RWN model. We note that the RWN segmentations contain fewer false positive predictions and are also spatially smoother across the object regions.

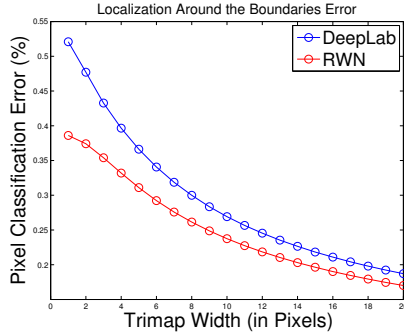


Figure 5: Localization error around the object boundaries within a trimap. Compared to the DeepLab system (blue), our RWN (red) achieves lower segmentation error around object boundaries for all trimap widths.

Furthermore, in Table 3, we present experiments where we compare RWN with models using dense CRFs [15] to post-process the predictions of DeepLab systems. We also include DeepLab-DT [5], which uses domain-transfer filtering to refine the segmentations inside an FCN. Based on these results, we observe that, despite not using any post-processing, our RWN produces results similar to or even better than the DeepLab models employing post-processing. These results indicate that RWN can be used as a globalization mechanism to ensure spatial coherence in semantic segmentation predictions. In Figure 4 we present qualitative results where we compare the final segmentation predictions of RWN and the DeepLab-v2-CRF system. Based on these qualitative results, we observe that RWN captures more accurately the fine details of the objects, such as the bike wheels, or plane wings. The DeepLab-v2-CRF system misses some of these object parts.

Localization Around the Boundaries. Earlier we claimed that due to the use of large receptive fields and many pooling layers, FCNs tend to produce blobby segmentations that lack fine object boundary details. We want to show that our RWN produces more accurate segmentations around object boundaries than the traditional FCNs. Thus, adopting the practice from [15], we evaluate segmentation accuracy around object boundaries. We do so by counting the relative number of misclassified pixels within a narrow band (“trimap”) surrounding the ground truth object boundaries. We present these results in Figure 5. The results show that RWN achieves higher segmentation accuracy than the DeepLab (DL) system for all trimap widths considered in this test.

Spatial Smoothness. We also argued that applying the dense CRF [15] as a post-processing technique often leads to spatially fragmented segmentations (see the right side of Fig. 1). How can we evaluate whether a given method produces spatially smooth or spatially fragmented segmenta-

Method	MF	AP
DeepLab-largeFOV-CRF	0.676	0.457
RWN-largeFOV	0.703	0.494
DeepLab-attention-CRF	0.722	0.521
RWN-attention	0.747	0.556
DeepLab-v2-CRF	0.763	0.584
RWN-v2	0.773	0.595

Table 4: Quantitative comparison of spatial segmentation smoothness. We extract the boundaries from the predicted segmentations and evaluate them against ground truth object boundaries using max F-score (MF) and average precision (AP) metrics. These results suggest that RWN segmentations are spatially smoother than the DeepLab-CRF segmentations across all baselines.

tions? Intuitively, spatially fragmented segmentations will produce many false boundaries that do not correspond to actual object boundaries. Thus, to test the spatial smoothness of a given segmentation, we extract the boundaries from the segmentation and then compare these boundaries against ground truth object boundaries using the standard maximum F-score (MF) and average precision (AP) metrics, as done in the popular BSDS benchmark [22]. We perform this experiment on the Pascal SBD dataset and present these results in Table 4. We can see that the boundaries extracted from the RWN segmentations yield better MF and AP results compared to the boundaries extracted from the different variants of the DeepLab-CRF system. Thus, these results suggest that RWN produces spatially smoother segmentations than DeepLab-CRF.

5.2. Scene Labeling

We also tested our RWN on the task of scene labeling using two popular datasets: Stanford Background [9] and Sift Flow [18]. Stanford Background is a relatively small dataset for scene labeling. It contains 715 images, which we randomly split into 600 training images and 115 testing images. In contrast, the Sift Flow dataset contains 2489 training examples and 201 testing images. For all of our experiments, we use the DeepLab-largeFOV [6] architecture since it is smaller and more efficient to train and test. To evaluate scene labeling results, we use the overall IOU evaluation metric which is a commonly used metric for this task. In Table 5, we present our scene labeling results on both of these datasets. Our results indicate that our RWN method outperforms the DeepLab baseline by 2.57%, and 2.54% on these two datasets, respectively.

5.3. Runtime Comparisons

We also include the runtime comparison of our RWN approach versus the denseCRF inference. We note that using a single core of a 2.7GHz Intel Core i7 processor, the denseCRF inference requires 3.301 seconds per image on

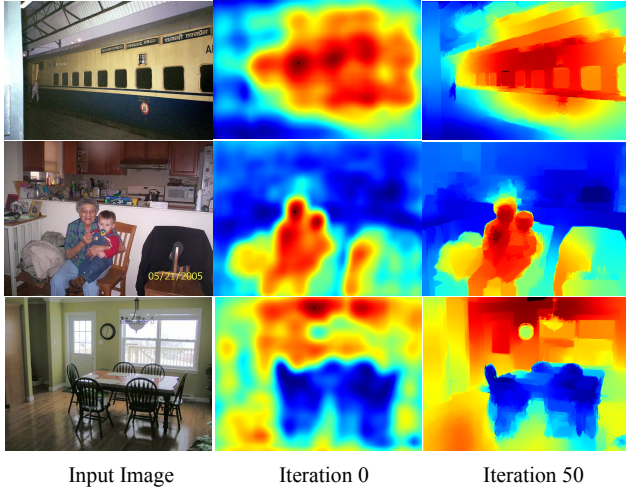


Figure 6: A figure illustrating how the probability predictions change as we apply more random walk steps. Note that the RWN predictions become more refined and better localized around the object boundaries as more random walk steps are applied.

average on a Pascal SBD dataset. In comparison, a single iteration of a random walk, which is simply a sparse matrix multiplication, takes 0.032 seconds on average on the same Pascal SBD dataset. A DeepLab_v2 post-processed with denseCRF achieves 81.9% IOU score on this same Pascal SBD dataset. In comparison, RWN_v2 with a single random walk iteration and with $R=40$ (radius) achieves 82.2% IOU, which is both more accurate and more than 100 times more efficient than the denseCRF inference.

5.4. Ablation Experiments

Optimal Number of Random Walk Steps. In Figure 7, we illustrate how the IOU accuracy changes when we use a different number of random walk steps. We observe that the segmentation accuracy keeps increasing as we apply more random walk steps, and that it reaches its peak performance when the random walk process converges, which indicates the effectiveness of our random walk step procedure. In Figure 6, we also illustrate how the predicted object segmentation probabilities change as we apply more random walk steps. We observe that the object boundaries become much better localized as more iterations of random walk are applied.

Radius Size. To analyze the effect of a radius size in the RWN architecture, we test alternative versions of our model with different radii sizes. Our results indicate, that the RWN model produces similar results with different radii in the interval of $R > 3$ and $R < 20$ if the random walk step process is applied until convergence. We also note that, if we select $R = 40$, and apply a random walk step only **once**, we can achieve the segmentation accuracy of 75.5% and 77.6% according to the two evaluation metrics,

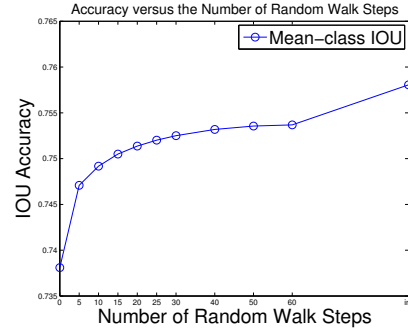


Figure 7: IOU accuracy as a function of the number of random walk steps. From this plot we observe that the segmentation accuracy keeps improving as we apply more random walk steps and that it reaches its peak when the random walk process converges.

	DeepLab-largeFOV	RWN-largeFOV
Stanford-BG	65.74	68.31
Sift-Flow	67.31	69.85

Table 5: Scene labeling results on the Stanford Background and Sift-Flow datasets measured according to the overall IOU evaluation metric. We use a DeepLab-largeFOV network as base model, and show that our RWN yields better results on both of these scene labeling datasets.

respectively. In comparison, choosing $R = 5$ and applying random walk until convergence yields the accuracies of 75.8% and 77.9%, which is slightly better. However, note that selecting $R = 40$, and applying multiple random walk steps does not yield any improvement in segmentation accuracy. These experiments show the flexibility of our model compared to the MRF or CRF models, which typically use graphs with a fixed grid structure. Our model has the ability to use different graph structures depending on the problem.

6. Conclusion

In this work, we introduced Random Walk Networks (RWNs), and showed that, compared to traditional fully convolutional networks (FCNs), they produce improved accuracy for the same model complexity. Our RWN addresses the issues of 1) poor localization around the segmentation boundaries and 2) spatially disjoint segmentations. Additionally, our implementation of RWN uses only 131 additional learnable parameters (0.0008% of the original number of the parameters in the network) and it can be easily integrated into the standard FCN learning framework for a joint end-to-end training. Finally, RWN provides a more efficient alternative to the denseCRF approaches.

Our future work includes experimenting with alternative RWN architectures and applying RWN to new domains such as language processing or speech recognition.

References

- [1] Pablo Arbelaez, J. Pont-Tuset, Jon Barron, F. Marqués, and Jitendra Malik. Multiscale combinatorial grouping. In *Computer Vision and Pattern Recognition (CVPR)*, 2014. 6
- [2] Vijay Badrinarayanan, Ankur Handa, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *arXiv preprint arXiv:1505.07293*, 2015. 4
- [3] Bahman Bahmani, Abdur Chowdhury, and Ashish Goel. Fast incremental and personalized pagerank. *Proc. VLDB Endow.*, 4(3):173–184, December 2010. 2, 3
- [4] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Semantic segmentation with boundary neural fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [5] Liang-Chieh Chen, Jonathan T. Barron, George Papandreou, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. *CVPR*, 2016. 1, 2, 7
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully. In *ICLR*, 2015. 1, 2, 4, 6, 7
- [7] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L. Yuille. Attention to scale: Scale-aware semantic image segmentation. *CVPR*, 2016. 2, 4, 6
- [8] Jifeng Dai, Kaiming He, and Jian Sun. Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015. 1
- [9] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2009. 6, 7
- [10] Bharath Hariharan, Pablo Arbelaez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *International Conference on Computer Vision (ICCV)*, 2011. 6
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 4
- [12] Seunghoon Hong, Hyeonwoo Noh, and Bohyung Han. Decoupled deep neural network for semi-supervised semantic segmentation. In *NIPS*, December 2015. 1, 2
- [13] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 2, 6
- [14] Iasonas Kokkinos. Surpassing humans in boundary detection using deep learning. *CoRR*, abs/1511.07386, 2015. 2
- [15] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In J. Shawe-Taylor, R.S. Zemel, P.L. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 109–117. Curran Associates, Inc., 2011. 1, 7
- [16] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. 3
- [17] 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. 1, 2, 3
- [18] Ce Liu, Jenny Yuen, and Antonio Torralba. *Nonparametric Scene Parsing via Label Transfer*, pages 207–236. Springer International Publishing, Cham, 2016. 6, 7
- [19] Ziwei Liu, Xiao Xiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015. 1, 2, 3
- [20] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 1, 2
- [21] Laszlo Lovasz. Random walks on graphs: A survey, 1993. 2, 3
- [22] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001. 7
- [23] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, 2015. 1, 2
- [24] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999. Previous number = SIDL-WP-1999-0120. 2, 3
- [25] Xiaojuan Qi, Jianping Shi, Shu Liu, Renjie Liao, and Jia-ya Jia. Semantic segmentation with object clique potential. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015. 2
- [26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 4
- [27] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip Torr. Conditional random fields as recurrent neural networks. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 2, 3