

# Recurrent Segmentation for Variable Computational Budgets

Lane McIntosh\*  
Stanford University  
lmcintosh@stanford.edu

Niru Maheswaranathan  
Google Brain  
nirum@google.com

David Sussillo  
Google Brain  
sussillo@google.com

Jonathon Shlens  
Google Brain  
shlens@google.com

State-of-the-art systems for semantic image segmentation use feed-forward pipelines with fixed computational costs. Building an image segmentation system that works across a range of computational budgets is challenging and time-intensive as new architectures must be designed and trained for every computational setting. To address this problem we develop a recurrent neural network that successively improves prediction quality with each iteration. Importantly, the RNN may be deployed across a range of computational budgets by merely running the model for a variable number of iterations. We find that this architecture is uniquely suited for efficiently segmenting videos. By exploiting the segmentation of past frames, the RNN can perform video segmentation at similar quality but reduced computational cost compared to state-of-the-art image segmentation methods. When applied to static images in the PASCAL VOC 2012 and Cityscapes segmentation datasets, the RNN traces out a speed-accuracy curve that saturates near the performance of state-of-the-art segmentation methods.

## 1. Introduction

Advances in the design of neural networks have driven the field of computer vision in the last few years [40, 16, 44]. State-of-the-art artificial vision systems in image recognition [40, 61, 59], object detection [55, 22, 23, 19, 62], depth estimation [18], semantic segmentation [8, 72, 23, 50, 21, 56], instance segmentation [54, 4] and many other image processing tasks [64, 45] are built on top of novel network architectures.

Designing neural networks is challenging and time-intensive (but see [74]). Additionally, deploying a trained vision system based on a neural network is computationally expensive often requiring dedicated, highly parallelized numerical linear algebra implementations [1, 11, 10] on cus-

tom SIMD hardware architectures [39, 40, 11, 65]. In addition to these challenges, applications of vision systems span a wide range of computational budgets—from severe time or energy budgets [7, 37, 66, 38] to unconstrained domains [56].

Semantic image segmentation [8, 72, 23, 50, 21, 56] is one such task that is commonly employed across problems with vastly different computational demands—from mobile applications [46, 5, 6] to medical image diagnosis [56].

As larger networks often lead to improved performance [67, 74], a critical aspect of network design is the ability to build systems that operate well across a range of computational budgets—including and especially, extremely small computational budgets [37, 71].

One may design and train a distinct network architecture for each computational budget [74], however this approach is time- and labor-intensive and may require the calibration of a cascade-based system [2]. Another approach is to build a fully convolutional network that operates on images of arbitrary size [8]. This approach leads to state-of-the-art models, but at the expense of building models that operate on images downsampled in spatial resolution.

We instead draw inspiration from the image compression literature, where a single network is often trained and deployed across a large range of computational budgets [64]. A natural way to achieve this goal is through a recurrent network architecture [26], which may be iterated a number of times in proportion to a desired accuracy. Although convolutional networks are far more common in architectures for vision problems, recurrent networks have been applied across tasks such as multiple object recognition [3], fine-grain image recognition [58], instance segmentation [54], and image generation [28].

In this work we propose a recurrent network architecture with an intermediate human-interpretable canvas for semantic image segmentation (Figure 2 and 3). The RNN output at each iteration is added to this canvas, which is then fed

\*Work done while an intern at Google Brain.

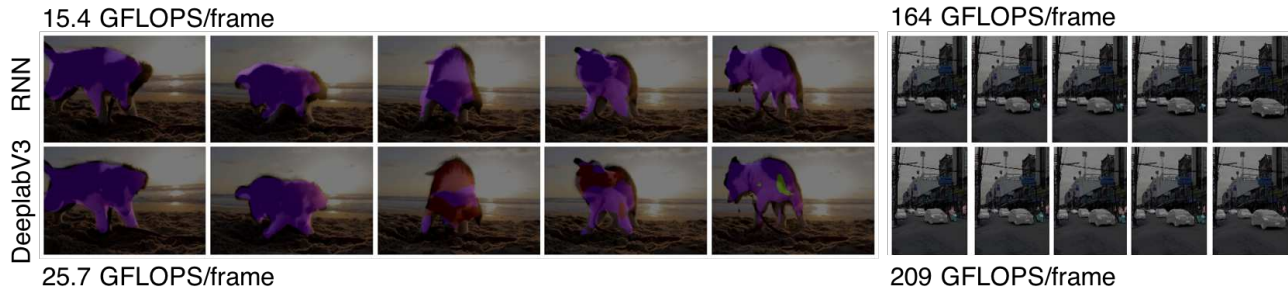


Figure 1. Leveraging a recurrent network to efficiently segment videos. Seeding a recurrent network with the segmentation from a previous frame notably reduces the computational budget required to segment a video. Segmentations from a recurrent network across  $180 \times 270$  and  $1080 \times 608$  resolution videos (top row) require about 25% less computational cost than a state-of-the-art feed-forward architecture [9] (bottom row), while maintaining higher semantic consistency across frames. Both networks are evaluated at a stride of 16 using a single inference pass.

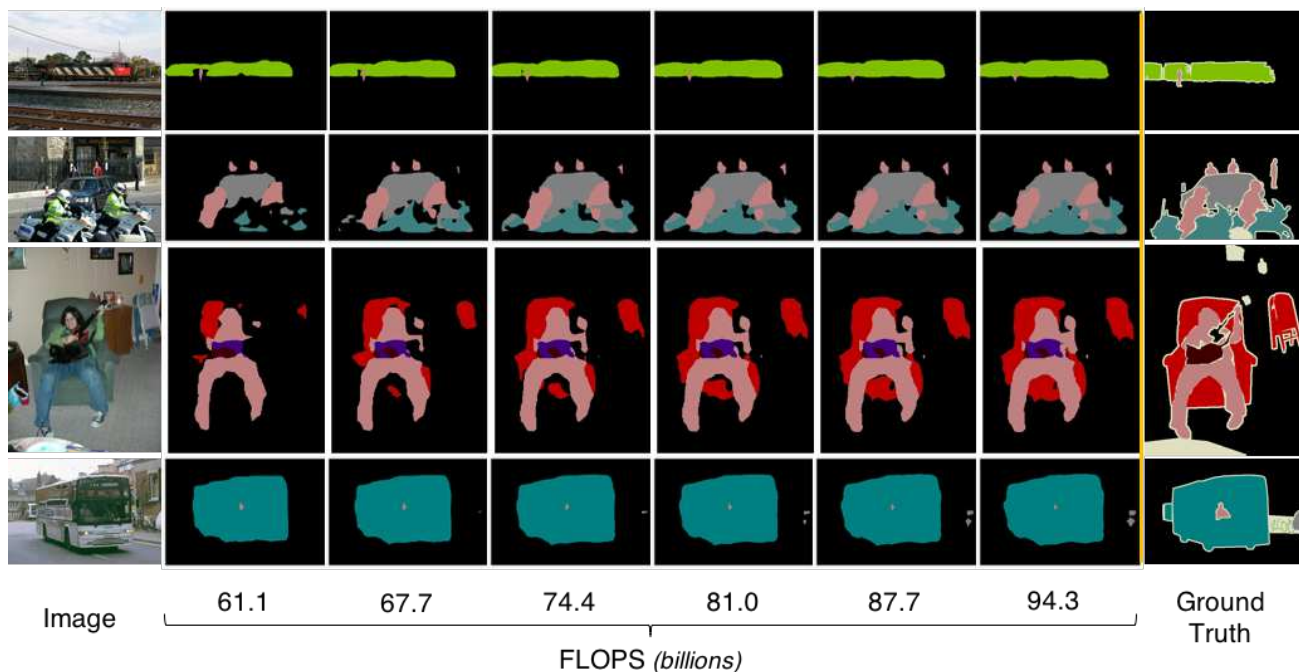


Figure 2. Recurrent segmentation iteratively improves segmentation predictions across a range of computational budgets. Images from the PASCAL VOC 2012 [20] validation dataset at  $513 \times 513$  resolution (left column) are segmented with a recurrent network for iterations one through six (columns two through seven), using single-scale evaluation at stride 16 (see Section 3). Note that most state-of-the-art segmentation networks use multi-scale inference at significantly higher computational costs. Ground truth segmentations (right column) are shown for comparison. Floating point operations per second (FLOPS) measures the cumulative computational cost of each RNN iteration. The cost at the first iteration includes both a feed-forward ResNet feature extractor and a single pass through the RNN layers, while each subsequent iteration only involves an additional pass through the RNN layers at a marginal cost of less than 7 GFLOPs.

in as input to the next iteration. The network thus updates past predictions stored in the canvas to iteratively refine the segmentation. We demonstrate that our RNN formulation may be applied to video with minimal computational cost by seeding the canvas with the segmentation from previous video frames (Figure 1). We further demonstrate that our

network corrects the errors of a segmentation on subsequent iterations of the recurrent network (Figure 6). We also visualize the spatial properties of the RNN through time (Figures 7 and 8).

On static images, we demonstrate the performance of the network on the PASCAL VOC 2012 [20] and Cityscapes

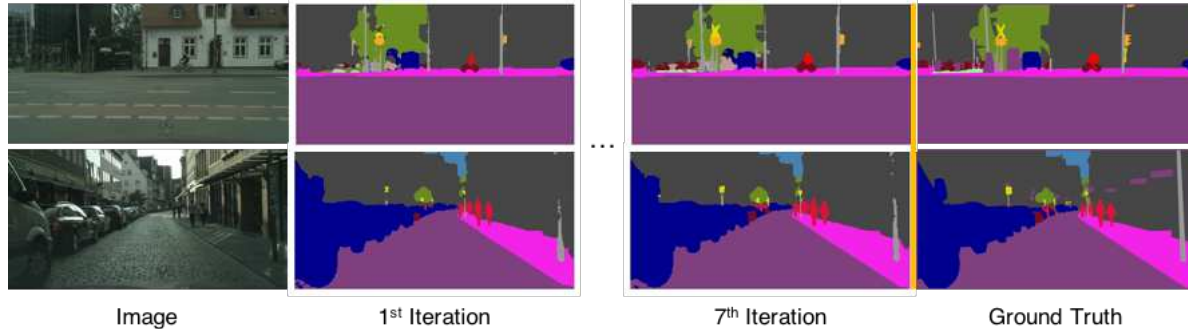


Figure 3. Recurrent segmentation iteratively improves Cityscapes segmentation masks [13]. Subsequent iterations add finer spatial details. High resolution images ( $1028 \times 2048$ ) from the Cityscapes semantic segmentation dataset (left column) are segmented by one or seven iterations of a recurrent network (columns two and three). By the seventh iteration, the RNN has added finer spatial segmentations (e.g. pedestrians, bicyclists, and sign posts) to the previous segmentation of large-scale structures. Ground truth segmentation masks (right column) are provided for comparison.

[13] semantic segmentation datasets. In this domain we find that the RNN architecture allows for more fine-grained control over computational cost than changing the output stride of a convolutional neural network [9], but at similar accuracy. One may select the number of iterations in the recurrent network, tracing out a curve of segmentation accuracy versus computational cost (Figure 5).

## 2. Related Work

Early work on semantic image segmentation learned single layer representations built on top of hand-crafted features (see [8] for review). The resurgence of neural networks enabled learning rich image features for classifying individual pixels, taking into account context and multi-scale information [30, 14, 17]. Broadly speaking, many of these systems use convolutional, feed-forward networks to extract deep features at multiple levels (and spatial resolutions) in order to densely predict the label [49, 8, 51]. Two key features of the feed-forward architectures are the use of *skip connections* [32] to bring lower level features into later processing stages [56] and *atrous* or *dilated convolutions* [36] to efficiently increase the receptive field size of activation maps without sacrificing spatial resolution [57, 8, 52, 24].

Instead of relying on a strictly feed-forward architecture, our work focuses on the application of a recurrent network architecture to semantic segmentation. This work draws inspiration from early work applying semantic segmentation to images in small recurrent neural networks [53]. Pinheiro and Collobert [53] built convolutional recurrent networks and trained the resulting system on small datasets labeled for pixel-wise segmentation (e.g. Stanford Background, SIFT Flow). In this work, we scale up these networks substantially in depth and overall size. Notably, we use the idea of a human-interpretable canvas that the net-

work continuously (additively) updates during inference as previously used in generative models of images [28, 15, 60]. A human-interpretable canvas allows the network to make a viable prediction at each successive step of the recurrence.

Segmentation systems often use Conditional Random Fields (CRFs) [42] as a post-processing step that uses image priors to improve the spatial or temporal consistency of segmentation masks [8, 41]. Zheng et al. 2015 [73] reformulated CRFs as a constrained recurrent neural network (CRF-RNN) that performs approximate variational inference, and found that end-to-end learning of a CRF-RNN on top of a feed-forward convolutional network improved upon previous state-of-the-art results. Here we do not use any CRF post-processing, and provide a generalized recurrent architecture that allows for trading off speed and accuracy.

## 3. Methods

The goal of our architecture design is to build an end-to-end fully convolutional recurrent network that may operate on arbitrarily sized input images, where a feed-forward network provides features to a recurrent network. The output of the recurrent network iteratively updates a canvas of semantic segmentation, which is then fed back to the first recurrent layer at the next time step (Figure 4). We provide a brief summary of the major features below but save the details for the Appendix.

We compute image features as the output of a conv4\_x (block4) from a ResNet-V1-101 network [33]. For a  $513 \times 513$  image, the image features are of size  $65 \times 65$  across 2048 feature channels (red box in Figure 4). This first portion of the network comprises 42.5M learned parameters and requires 54.4 gigaFLOPS (GFLOPS) to process a single  $513 \times 513$  image. Although we experimented with training the convolutional features from scratch, we found that pre-trained features suffice and permit us to accelerate ex-

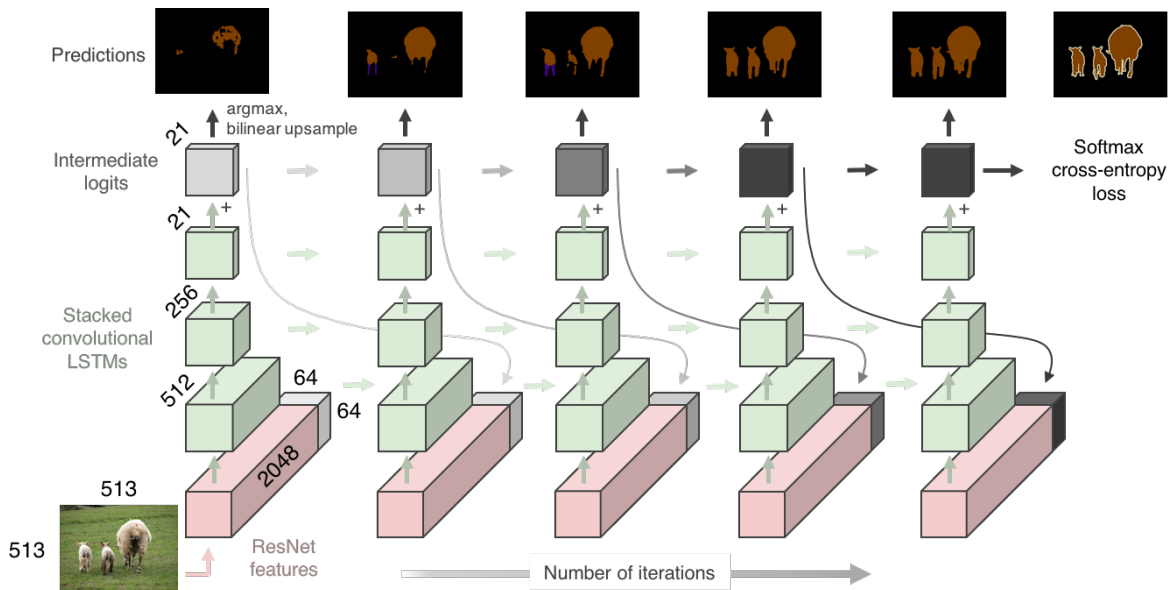


Figure 4. Diagram of recurrent network architecture for image segmentation. The model employs a human-interpretable canvas that is additively updated in subsequent iterations [28]. High-level image features from a pre-trained, fully convolutional network (e.g. [32, 37]) are extracted and concatenated with an all-zero “canvas” which stores the RNN’s intermediate logits. On each iteration, a stack of convolutional LSTMs additively updates intermediate logits from the previous iteration by adding their output to the previous iteration’s logits. The updated logits are then concatenated with the high-level image features and fed to the RNN at the next iteration.

perimentation, hence all experiments presented use ResNet features pre-trained on the COCO semantic segmentation dataset [47].

We use LSTM layers [35] for ease of training [12]. In particular, we use convolutional LSTMs [70] (for definition see Appendix 5). All of the network weights are shared across spatial locations within an activation map. We used a  $1 \times 1$  convolutional kernel size, which combines information only across activation maps. Preliminary experiments found that RNNs with larger than  $1 \times 1$  kernels—either dense or dilated—underperformed RNNs with  $1 \times 1$  kernels. In addition to learning the weights and biases during training, we also experimented with learning the initial state of the network. While these experiments produced lower error on the training dataset, they demonstrated poorer generalization; in the results presented here, we initialize the network state as zeros.

Previous work with sequence models [68] indicate that stacking multiple LSTM layers is a powerful method for capturing dependencies at multiple scales. In preliminary experiments we found that 3 layers provided a reasonable trade off between model complexity and computational cost. Each convolutional LSTM layer maintained the same spatial resolution while decreasing the output depth from the number of image feature channels to the number of semantic classes.

The final output of the network are the logits of an ad-

ditive canvas [28]. Each iteration of the recurrent network provides additive adjustments to the logits for each class at each pixel location by simply adding the output of the last RNN layer to the canvas. The logits from each iteration are concatenated to the ResNet features to provide input on subsequent iterations of the recurrent network (Figure 4). The prediction of the network at any iteration is generated by taking the *arg max* of the bilinearly upsampled canvas.

We used a softmax cross-entropy loss across the labels for each pixel. The loss was applied to the final canvas after  $N$  iterations, where  $N$  is a free parameter discussed below. In preliminary experiments we determined that weighted versions of the loss at earlier iterations of the canvas resulted in similar training performance, but changed the shape of the speed/accuracy curve. In particular, applying the loss to iteration three or earlier improved the prediction quality at earlier iterations, but decreased the quality of the best segmentation and produced a speed-accuracy trade-off curve that did not monotonically increase. All results shown here are for networks trained with a single loss applied at iteration six.

We trained the resulting architecture on the PASCAL VOC 2012 dataset [20] which contains 11,530 images of variable size, 6,929 of which have segmentation masks assigning each pixel into one of 20 classes (excluding the background class). We also trained models on the Cityscapes semantic segmentation dataset [13], which con-

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mIOU
CRF-RNN [73]	90.4	55.3	88.7	68.4	69.8	88.3	82.4	85.1	32.6	78.5	64.4	79.6	81.9	86.4	81.8	58.6	82.4	53.5	77.4	70.1	74.7
DeepLab [8]	92.6	60.4	91.6	63.4	76.3	95.0	88.4	92.6	32.7	88.5	67.6	89.6	92.1	87.0	87.4	63.3	88.3	60.0	86.8	74.5	79.7
PSPNet [72]	95.8	72.7	95.0	78.9	84.4	94.7	92.0	95.7	43.1	91.0	<b>80.3</b>	91.3	96.3	92.3	90.1	71.5	<b>94.4</b>	66.9	88.8	82.0	85.4
DeepLabv3 [9]	96.4	76.6	92.7	77.8	<b>87.6</b>	96.7	90.2	95.4	47.5	93.4	76.3	91.4	<b>97.2</b>	91.0	<b>92.1</b>	71.3	90.9	<b>68.9</b>	<b>90.8</b>	79.3	85.7
SDN+ [21]	<b>96.9</b>	<b>78.6</b>	<b>96.0</b>	<b>79.6</b>	84.1	<b>97.1</b>	<b>91.9</b>	<b>96.6</b>	<b>48.5</b>	<b>94.3</b>	78.9	<b>93.6</b>	95.5	<b>92.1</b>	91.1	<b>75.0</b>	93.8	64.8	89.0	<b>84.6</b>	<b>86.6</b>
RNN <sub>iters=6</sub>	91.3	70.1	91.5	63.7	77.2	92.5	90.6	93.6	36.8	88.2	62.2	89.7	92.4	85.8	87.6	67.2	89.8	61.8	84.4	73.8	80.3

Table 1. Comparison of the RNN to state-of-the-art image segmentation systems on PASCAL VOC 2012 validation dataset. Performance is measured as intersection-over-union (bold indicates state-of-the-art). RNN segmentations are competitive with best feed-forward architectures. Note the significant variability across segmentation labels (see Results for discussion). All results are with multi-scale evaluation, and use pre-training on the COCO dataset [47].

sists of 20,000 coarse training images, 3,475 fine validation images and 1,525 test images.

During training, optimization was performed using stochastic gradient descent with momentum, where the learning rate was reduced monotonically with a polynomial schedule (see Appendix for details).

We test the RNN model and competing architectures across *single-scale* and *multi-scale* evaluation schemes. In the single-scale scheme, we perform a single inference pass at a coarse ResNet output stride of 16. This evaluation scheme is generally steered towards mobile platforms with limited computational budgets, and we focus on this in Figure 5. Additionally, we test the model in a less computationally restricted setting by performing multiple inference passes per image: resampling the image at six different spatial scales, performing left-right flips, and averaging the logits, as described in [9]. The multi-scale setting achieves the best quantitative performance but at a notable computational expense (Table 3).

## 4. Results

We present results measuring the performance of the RNN architecture on image segmentation tasks, and highlight experiments indicating how the architecture may be applied to video segmentation in a computationally efficient manner. Finally, we examine the errors of the recurrent network to illustrate how the system operates and suggest opportunities for improvement.

### 4.1. Image segmentation with variable computational budgets

We trained the recurrent architecture on the PASCAL VOC 2012 dataset [20] to segment images into twenty semantic classes. This dataset is a popular test bed for advancements in image segmentation [8, 72, 21]. The proposed RNN was unrolled for six iterations and a cross entropy loss was applied to the final canvas state (Figure 4). The final network is fully convolutional and may operate on arbitrarily-sized images, however we focus our analysis at two operation regimes.

Method	IOU classes	iIOU classes	IOU categ.	iIOU categ.
FCN [50]	65.3	41.7	85.7	70.1
CRF-RNN [73]	62.5	34.4	82.7	66.0
DeepLab [8]	70.4	42.6	86.4	67.7
PSPNet [72]	<b>78.4</b>	<b>56.7</b>	<b>90.6</b>	<b>78.6</b>
RNN <sub>iters=6</sub>	72.7	46.0	87.6	73.9

Table 2. Comparison of the RNN to state-of-the-art image segmentation systems on the Cityscapes high-resolution validation dataset. Performance is measured as the pixel-wise intersection-over-union (IOU) and instance-level intersection-over-union (iIOU) averaged per category or per class. All models are trained only on the Cityscapes coarse dataset. The RNN consistently achieves second best segmentation despite a reduced computational budget (see Table 3).

First, we analyze the network at a  $513 \times 513$  spatial resolution in our single-scale evaluation regime. A single inference pass of one image at a single iteration results in 61.1 GFLOPs of computation (Figure 2 second column, Figure 5 single-scale, Table 3). In contrast, state-of-the-art architectures for image segmentation, such as [72] and [9], require 40%-100% more computational cost at the same output stride, indicating that our model is less computationally demanding (Table 3). Interestingly, because of the structure of the RNN, the model may produce intermediate predictions (i.e. by calculating the argmax of the canvas, see Figure 4). Specifically, at inference time, one may continue to iterate the RNN, producing an improved segmentation as the computational budget grows (Figure 2, third - seventh columns). The computational demand roughly grows linearly with the number of iterations, where each RNN iteration is 10x and 5x cheaper than the comparable post-processing performed in PSPNet [72] and DeeplabV3 [9], respectively.

We note that fully convolutional networks also have the ability to exchange accuracy for compute, e.g. by decreasing the image resolution or increasing the output stride of the feature extraction step [9]. When we compare the speed-accuracy trade-offs of changing output stride versus running recurrent segmentation for a variable number of iterations on a  $513 \times 513$  image, we find that the RNN achieves a

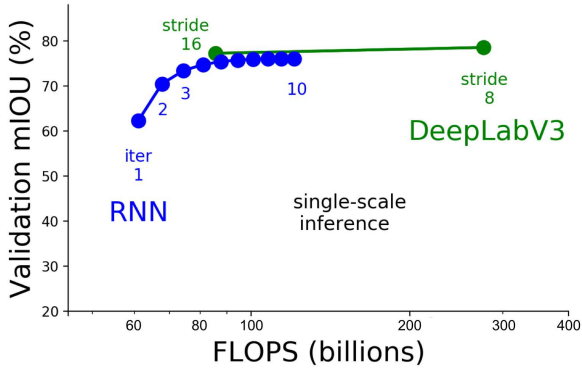


Figure 5. Trade-off between computational cost and segmentation accuracy. We plot the computational cost (FLOPS) versus model accuracy (mIOU) on the PASCAL VOC 2012 validation dataset for single-scale evaluation (see Methods). Increasing the number of iterations (“iter”) in the RNN improves model accuracy (blue curves). Additionally, a state-of-the-art fully convolutional network [9] is evaluated across two output strides (8 and 16). Importantly, recurrent segmentation achieves a finer control over the speed-accuracy trade-off with a similar best-quality segmentation as [9]. Both networks use the same pre-trained ResNet-101 feature extractor.

similar quality segmentation when controlling for computational cost, but with more fine grained control over the speed-accuracy trade-off (Figure 5).

We quantified performance on the PASCAL VOC 2012 test dataset in terms of the mean intersection-over-union (mIOU) between the ground truth labels and network predictions in Table 1. When run for many iterations, the RNN performs comparably to state-of-the-art networks [21, 69, 8, 72, 9], but performs notably worse on specific classes that have high spatial frequency features, such as bicycle and chair. We discuss reasons for this behavior in Section 4.4 and the Discussion.

We additionally examined how the RNN performs with a very small computational budget. In this case, we computed mIOU versus the number of RNN steps (Figure 5). The first iteration of the RNN leads to a large performance gain and these gains accrue as one iterates the network. Interestingly, if one runs the network for *more steps* than it was originally trained, the RNN predictions continue to improve (Figure 5, blue curves) as the canvas saturates.

We tested the same network on the Cityscapes [13] image segmentation dataset (Table 2). The performance of the network, trained only on the Cityscapes coarse dataset, was assessed using the mIOU as before – in addition, we calculated the instance-level intersection-over-union (iIOU), where the contribution of each pixel to the metric is weighted by the ratio of the class’ average instance size to the size of the ground truth instance. The RNN

Method		single-scale (GFLOPS)	multi-scale (GFLOPS)
Deeplab V3 [9]	stride 16	85.6	1,540
	stride 8	276	4,870
PSPNet [72]	stride 16	122	1,740
	stride 8	260	4,090
RNN	1 iters	61.1	1,040
	2 iters	67.7	1,160
	3 iters	74.4	1,269
	4 iters	81.0	1,380
	5 iters	87.7	1,500
	6 iters	94.3	1,610

Table 3. Computational costs for RNN and state-of-the-art segmentation methods for single-pass and multi-pass evaluation for 513x513 images. Fully convolutional models [72, 9] may be evaluated across multiple output strides to reduce computational cost (but with reduced predictive precision). The RNN may be evaluated with an increasing number of iterations to improve predictive performance. Note that the RNN requires a consistently smaller computational budget.

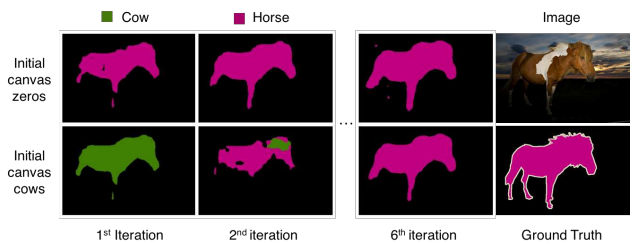


Figure 6. Error correction through RNN dynamics. When the canvas is initialized as zeros, the RNN correctly segments a horse after a single iteration (top row). When the canvas is *incorrectly* initialized as a cow, the RNN nonetheless corrects its prediction after two iterations (bottom row).

provided comparative results with respect to state-of-the-art methods in spite of the reduced computational budget.

## 4.2. Observations of error correction with recurrent networks

Given that an RNN performs identical operations on every iteration, it is natural to characterize the computational properties of this generic operation. We addressed this question by artificially perturbing the segmentation mask stored in the RNN canvas and examining how the RNN responded to these perturbations. When we seed the initial canvas to the wrong class—segmenting a horse as a cow, for instance—we find that the RNN is nonetheless able to correct the semantic segmentation in just two iterations (Figure 6, bottom row). After a few more iterations, the segmentation mask from the perturbed RNN is of similar quality to the RNN initialized with a canvas of all zeros, highlighting the robustness of the method to an incorrect or poor initial segmentation (Figure 6, top row).

The RNN’s ability to iteratively improve partial segmen-

tations (Figure 2) and correct segmentation errors (Figure 6) suggests a natural extension to segmenting video. In this setting, consecutive frames are highly correlated and segmentations from previous frames provide a good starting point for subsequent frames.

### 4.3. Leveraging recurrent networks for efficient video segmentation

Temporal correlations between video frames have been exploited in video compression [43] and the propagation of weak labels across video frames [63, 31]. In the latter case, optical flow and motion tracking algorithms may learn signals for propagating label information across frames [48]. In the case of semantic video segmentation, we ask whether the structure of the RNN provides a computationally efficient method for label propagation across subsequent frames of a video.

Section 4.1 indicates that subsequent iterations of the RNN refine previous segmentations. A natural question is if the RNN may improve the segmentation provided from a previous frame with only a few iterations, thereby saving significant computational resources. Figure 1 provides an example pipeline for segmenting recorded video from public domain videos and a mobile device recording, respectively. The first video frame is segmented with the RNN as described above. In subsequent video frames, we seed the RNN canvas with the segmentation from a previous video frame and merely run the RNN for two iterations. The resulting segmentation looks comparable to having run the RNN for many iterations, although the computational demand is significantly reduced. In Figure 1 we demonstrate that this video segmentation technique is qualitatively comparable to [8] across an array of public domain videos in which ground truth is unknown, at just over half the cost (15.4 GFLOPS/frame for  $RNN_{iter=2}$  vs. 25.7 GFLOPS/frame for DeepLab V3 [9], on a  $180 \times 270$  video).

### 4.4. Diagnosing failures of the recurrent network

We found that despite achieving near state-of-the-art in several categories of the PASCAL VOC 2012 dataset (Table 1), the performance on several classes—in particular, classes with high spatial frequency features—was poor. To investigate the capacity of the recurrent neural network to learn fine-grained features, we initialized the RNN canvas to the ground truth label of a bicycle with the frame, rims, and wheel spokes all individually segmented (Figure 7, bottom row). When the RNN is allowed to successively iterate on the already correct canvas, the RNN nonetheless progressively fills in the empty areas of the segmentation mask with each iteration, covering a convex hull of the object. The final segmentation mask closely resembles the low spatial frequency mask discovered when the RNN evolves from a blank canvas (Figure 7, top row). We found that even

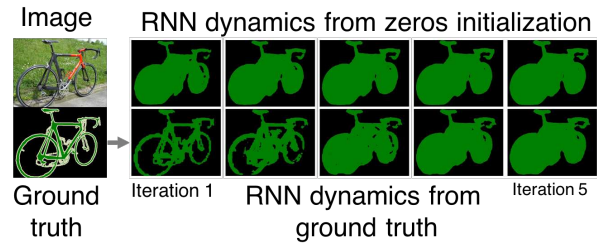


Figure 7. Visualizing RNN dynamics after initializing the canvas with zeros (top row) the ground-truth label (bottom row). Even when the canvas is initialized to the higher spatial frequency ground truth labels, the network fills in pixels in the object’s interior. Since the PASCAL VOC dataset segments individual bike spokes and wheels, this represents a failure mode of the RNN.

for classes with higher spatial frequencies, the RNN predictions had a spatial frequency distribution that matched the average spatial power spectral density across all classes (Supplemental Figure 8).

## 5. Discussion

We propose a recurrent neural network architecture for semantic segmentation that enables variable quality segmentation across computational budgets. This yields a natural framework for segmenting videos, resulting in a high quality segmentation of video frames at a fraction of the computational cost previously feasible.

When applied to images, the RNN model can finely trade-off between accuracy and computational demand. We note however that RNN dynamics tend to fill in the interior of objects with a single semantic label. To address this, it would be possible to target the fine-grained features in these difficult images by bootstrapping [9] or by explicitly modifying the loss function to penalize object boundaries [4].

Due to memory constraints during training, the RNN processes a relatively coarse spatial resolution due to the output stride applied to the image features. While it is remarkable that the RNN achieves good performance with such coarse-grained features, we anticipate gains in GPU hardware memory will permit increased spatial resolution in convolutional RNN architectures and thus allow for more capacity in the network. In the interim, switching from an LSTM to other RNN architectures with fewer parameters may alleviate some of these memory limitations during training [26]. Preliminary experiments also indicate that pre-training on larger internal datasets [34] disproportionately improves the performance of the RNN compared to baseline models, resulting in RNNs that achieve state-of-the-art performance on some categories in Pascal VOC. These results suggest that the RNN architecture has a large capacity and may scale to larger problems. Another di-

rection may be to relax the strict requirement that every iteration of the RNN must perform an identical computation. One method for approaching this is to build a hypernetwork [29] whereby a second network predicts the network weights for each iteration. Another approach is to instead build an RNN architecture that adaptively updates the number of iterations [27], perhaps on a per-pixel level, in order to achieve a given accuracy level.

The intersection of computational constraints, predictive power, and memory limitations necessitate a diversity of architectural approaches for image and video segmentation. Here we provide one such method, a recurrent architecture that trades off speed and accuracy for semantic segmentation. Developing novel architectures for computer vision problems remains fertile ground for future research.

## Acknowledgements

We thank Kevin Murphy and George Papandreou for comments on the manuscript, Liang-Chieh Chen for helpful comments about running Deeplab V3, and George Toderici and Ben Poole for useful discussions.

## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. J. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Józefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. G. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. A. Tucker, V. Vanhoucke, V. Vasudevan, F. B. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2015. **1**
- [2] A. Angelova, A. Krizhevsky, and V. Vanhoucke. Pedestrian detection with a large-field-of-view deep network. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 704–711. IEEE, 2015. **1**
- [3] J. Ba, V. Mnih, and K. Kavukcuoglu. Multiple object recognition with visual attention. In *International Conference of Learned Representations (ICLR)*, 2015. **1**
- [4] M. Bai and R. Urtasun. Deep watershed transform for instance segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **1, 7**
- [5] J. T. Barron, A. Adams, Y. Shih, and C. Hernández. Fast bilateral-space stereo for synthetic defocus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4466–4474, 2015. **1**
- [6] J. T. Barron and B. Poole. The fast bilateral solver. In *European Conference on Computer Vision*, pages 617–632. Springer, 2016. **1**
- [7] M. Bojarski, D. D. Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba. End to end learning for self-driving cars. *CoRR*, abs/1604.07316, 2016. **1**
- [8] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. **1, 3, 5, 6, 7**
- [9] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. **2, 3, 5, 6, 7, 11, 12**
- [10] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015. **1**
- [11] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759*, 2014. **1**
- [12] J. Collins, J. Sohl-Dickstein, and D. Sussillo. Capacity and trainability in recurrent neural networks. In *International Conference of Learned Representations (ICLR)*, 2016. **4**
- [13] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. **3, 4, 6**
- [14] J. Dai, K. He, and J. Sun. Convolutional feature masking for joint object and stuff segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3992–4000, 2015. **3**
- [15] E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015. **3**
- [16] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International Conference on Machine Learning*, volume 32, pages 647–655, 2014. **1**
- [17] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2650–2658, 2015. **3**
- [18] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2366–2374. Curran Associates, Inc., 2014. **1**
- [19] D. Erhan, C. Szegedy, A. Toshev, and D. Anguelov. Scalable object detection using deep neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. **1**
- [20] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010. **2, 4, 5**



- [21] J. Fu, J. Liu, Y. Wang, and H. Lu. Stacked deconvolutional network for semantic segmentation. *arXiv preprint arXiv:1708.04943*, 2017. 1, 5, 6
- [22] R. Girshick. Fast r-cnn. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015. 1
- [23] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 1
- [24] A. Giusti, D. C. Ciresan, J. Masci, L. M. Gambardella, and J. Schmidhuber. Fast image scanning with deep max-pooling convolutional neural networks. In *Image Processing (ICIP), 2013 20th IEEE International Conference on*, pages 4034–4038. IEEE, 2013. 3
- [25] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010. 11
- [26] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016. 1, 7
- [27] A. Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016. 8
- [28] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, 2015. 1, 3, 4
- [29] D. Ha, A. Dai, and Q. V. Le. Hypernetworks. In *International Conference on Learning Representations*, 2017. 8
- [30] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 447–456, 2015. 3
- [31] G. Hartmann, M. Grundmann, J. Hoffman, D. Tsai, V. Kwatra, O. Madani, S. Vijayanarasimhan, I. Essa, J. Rehg, and R. Sukthankar. Weakly supervised learning of object segmentations from web-scale video. In *Computer Vision—ECCV 2012. Workshops and Demonstrations*, pages 198–208. Springer, 2012. 7
- [32] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 4
- [33] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 3
- [34] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 7
- [35] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 4
- [36] M. Holschneider, R. Kronland-Martinet, J. Morlet, and P. Tchamitchian. A real-time algorithm for signal analysis with the help of the wavelet transform. In *Wavelets*, pages 286–297. Springer, 1990. 3
- [37] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1, 4
- [38] M.-C. Hu, Y.-N. Tseng, J.-L. Wu, C.-H. Kuo, Y.-M. Hsiao, and K.-L. Hua. Real-time depth of field rendering with bokeh effect. In *Consumer Electronics (ISCE), 2013 IEEE 17th International Symposium on*, pages 99–100. IEEE, 2013. 1
- [39] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-dataloader performance analysis of a tensor processing unit. *arXiv preprint arXiv:1704.04760*, 2017. 1
- [40] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing System*, 2012. 1
- [41] A. Kundu, V. Vineet, and V. Koltun. Feature space optimization for semantic video segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 3168–3175. IEEE, 2016. 3
- [42] J. Lafferty, A. McCallum, and F. C. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001. 3
- [43] D. Le Gall. Mpeg: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58, 1991. 7
- [44] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. 1
- [45] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [46] M. Levoy and Y. Pritch. Portrait mode on the pixel 2 and pixel 2 xl smartphones, 2017. 1
- [47] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 4, 5
- [48] C. Liu, J. Yuen, and A. Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):978–994, 2011. 7
- [49] Z. Liu, X. Li, P. Luo, C.-C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1377–1385, 2015. 3
- [50] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 1, 5
- [51] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1520–1528, 2015. 3
- [52] G. Papandreou, I. Kokkinos, and P.-A. Savalle. Modeling local and global deformations in deep learning: Epitomic convolution, multiple instance learning, and sliding window detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 390–399, 2015. 3

- [53] P. Pinheiro and R. Collobert. Recurrent convolutional neural networks for scene labeling. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 82–90, Beijing, China, 22–24 Jun 2014. PMLR. 3
- [54] M. Ren and R. S. Zemel. End-to-end instance segmentation and counting with recurrent attention. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [55] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015. 1
- [56] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015. 1, 3
- [57] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013. 3
- [58] P. Sermanet, A. Frome, and E. Real. Attention for fine-grained categorization. In *International Conference of Learned Representations (ICLR)*, 2015. 1
- [59] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 1
- [60] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448, 2015. 3
- [61] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 1
- [62] C. Szegedy, A. Toshev, and D. Erhan. Deep neural networks for object detection. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2553–2561. Curran Associates, Inc., 2013. 1
- [63] K. Tang, R. Sukthankar, J. Yagnik, and L. Fei-Fei. Discriminative segment annotation in weakly labeled video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2483–2490, 2013. 7
- [64] G. Toderici, S. M. O’Malley, S. J. Hwang, D. Vincent, D. Minnen, S. Baluja, M. Covell, and R. Sukthankar. Variable rate image compression with recurrent neural networks. In *International Conference of Learned Representations (ICLR)*, 2016. 1
- [65] A. Viebke and S. Pillana. The potential of the intel (r) xeon phi for supervised deep learning. In *High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on Cyberspace Safety and Security (CSS), 2015 IEEE 12th International Conference on Embedded Software and Systems (ICESSE), 2015 IEEE 17th International Conference on*, pages 758–765. IEEE, 2015. 1
- [66] P. Warden. Running tensorflow graphs on microcontrollers, 2017. 1
- [67] W. Williams, N. Prasad, D. Mrva, T. Ash, and T. Robinson. Scaling recurrent neural network language models. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 5391–5395. IEEE, 2015. 1
- [68] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016. 4
- [69] Z. Wu, C. Shen, and A. v. d. Hengel. Bridging category-level and instance-level semantic image segmentation. *arXiv preprint arXiv:1605.06885*, 2016. 6
- [70] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015. 4, 12
- [71] X. Zhang, X. Zhou, L. Mengxiao, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *arXiv preprint arXiv:1707.01083*, 2017. 1
- [72] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. *arXiv preprint arXiv:1612.01105*, 2016. 1, 5, 6, 12
- [73] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1529–1537, 2015. 3, 5
- [74] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. Learning transferable architectures for scalable image recognition. *CoRR*, abs/1707.07012, 2017. 1