

Multi-View Deep Learning for Consistent Semantic Mapping with RGB-D Cameras

Lingni Ma¹, Jörg Stückler², Christian Kerl¹ and Daniel Cremers¹

Abstract—Visual scene understanding is an important capability that enables robots to purposefully act in their environment. In this paper, we propose a novel deep neural network approach to predict semantic segmentation from RGB-D sequences. The key innovation is to train our network to predict multi-view consistent semantics in a self-supervised way. At test time, its semantics predictions can be fused more consistently in semantic keyframe maps than predictions of a network trained on individual views. We base our network architecture on a recent single-view deep learning approach to RGB and depth fusion for semantic object-class segmentation and enhance it with multi-scale loss minimization. We obtain the camera trajectory using RGB-D SLAM and warp the predictions of RGB-D images into ground-truth annotated frames in order to enforce multi-view consistency during training. At test time, predictions from multiple views are fused into keyframes. We propose and analyze several methods for enforcing multi-view consistency during training and testing. We evaluate the benefit of multi-view consistency training and demonstrate that pooling of deep features and fusion over multiple views outperforms single-view baselines on the NYUDv2 benchmark for semantic segmentation. Our end-to-end trained network achieves state-of-the-art performance on the NYUDv2 dataset in single-view segmentation as well as multi-view semantic fusion.

I. INTRODUCTION

Intelligent robots require the ability to understand their environment through parsing and segmenting the 3D scene into meaningful objects. The rich appearance-based information contained in images renders vision a primary sensory modality for this task.

In recent years, large progress has been achieved in semantic segmentation of images. Most current state-of-the-art approaches apply deep learning for this task. With RGB-D cameras, appearance as well as shape modalities can be combined to improve the semantic segmentation performance. Less explored, however, is the usage and fusion of multiple views onto the same scene which appears naturally in the domains of 3D reconstruction and robotics. Here, the camera is moving through the environment and captures the scene from multiple view points. Semantic SLAM aims at aggregating several views in a consistent 3D geometric and semantic reconstruction of the environment.

In this paper, we propose a novel deep learning approach for semantic segmentation of RGB-D images with multi-view context. We base our network on a recently proposed deep

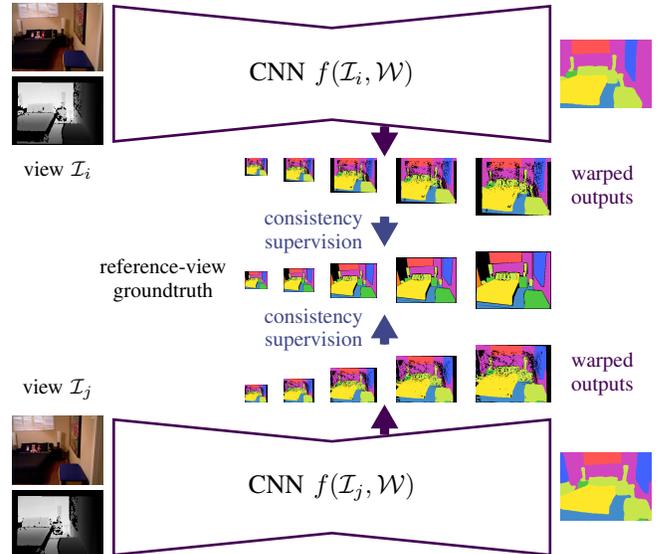


Fig. 1: We train our CNN to predict multi-view consistent semantic segmentations for RGB-D images. The key innovation is to enforce consistency by warping CNN feature maps from multiple views into a common reference view using the SLAM trajectory and to supervise training at multiple scales. Our approach improves performance for single-view segmentation and is specifically beneficial for multi-view fused segmentation.

convolutional neural network (CNN) for RGB and depth fusion [1] and enhance the approach with multi-scale deep supervision. Based on the trajectory obtained through RGB-D simultaneous localization and mapping (SLAM), we further regularize the CNN training with multi-view consistency constraints as shown in Fig. 1. We propose and evaluate several variants to enforce multi-view consistency during training. A shared principle is using the SLAM trajectory estimate to warp network outputs of multiple frames into the reference view with ground-truth annotation. By this, the network not only learns features that are invariant under view-point change. Our semi-supervised training approach also makes better use of the annotated ground-truth data than single-view learning. This alleviates the need for large amounts of annotated training data which is expensive to obtain. Complementary to our training approach, we aggregate the predictions of our trained network in keyframes to increase segmentation accuracy at testing. The predictions of neighboring images are fused into the keyframe based on the SLAM estimate in a probabilistic way.

In experiments, we evaluate the performance gain achieved through multi-view training and fusion at testing over single-view approaches. Our results demonstrate that multi-view

This work is funded by ERC Consolidator Grant 3D Reloaded (649323).

¹ Lingni Ma, Christian Kerl and Daniel Cremers are with Computer Vision Group, Department of Computer Science, Technical University of Munich, ([lingni](mailto:lingni@in.tum.de), [kerl](mailto:kerl@in.tum.de), [cremers](mailto:cremers@in.tum.de))@in.tum.de

² Jörg Stückler (stueckler@vision.rwth-aachen.de) is with Computer Vision Group, Visual Computing Institute, RWTH Aachen University. Jörg Stückler has been supported by ERC Starting Grant CV-SUPER (ERC-StG-2012-307432)

max-pooling of feature maps during training best supports multi-view fusion at testing. Overall we find that enforcing multi-view consistency during training significantly improves fusion at test time versus fusing predictions from networks trained on single views. Our end-to-end training achieves state-of-the-art performance on the NYUDv2 dataset in single-view segmentation as well as multi-view semantic fusion. While the fused keyframe segmentation can be directly used in robotic perception, our approach can also be useful as a building block for semantic SLAM using RGB-D cameras.

II. RELATED WORK

Recently, remarkable progress has been achieved in semantic image segmentation using deep neural networks and, in particular, CNNs. On many benchmarks, these approaches exceed previous techniques by a great margin.

Image-based Semantic Segmentation. As one early attempt, Couprie et al. [2] propose a multiscale CNN architecture to combine information at different receptive field resolutions and achieved reasonable segmentation results. Gupta et al. [3] integrate depth into the R-CNN approach by Girshick et al. [4] to detect objects in RGB-D images. They convert depth into 3-channel HHA, *i.e.*, disparity, height and angle encoding and achieve semantic segmentation by training a classifier for superpixels based on the CNN features. Long et al. [5] propose a fully convolutional network (FCN) which enables end-to-end training for semantic segmentation. Since CNNs reduce the input spatial resolution by a great factor through layers pooling, FCN presents an upsample stage to output high-resolution segmentation by fusing low-resolution predictions. Inspired by FCN and auto-encoders [6], encoder-decoder architectures have been proposed to learn upsampling with unpooling and deconvolution [7]. For RGB-D images, Eigen et al. [8] propose to train CNNs to predict depth, surface normals and semantics with a multi-task network and achieve very good performance. FuseNet [1] proposes an encoder-decoder CNN to fuse color and depth cues in an end-to-end training for semantic segmentation, which is shown to be more efficient in learning RGB-D features in comparison to direct concatenation of RGB and depth or the use of HHA. Recently, more complex CNN architectures have been proposed that include multi-resolution refinement [9], dilated convolutions [10] and residual units (e.g., [11]) to achieve state-of-the-art single image semantic segmentation. Li et al. [12] use a LSTM recurrent neural network to fuse RGB and depth cues and obtain smooth predictions. Lin et al. [13] design a CNN that corresponds to a conditional random field (CRF) and use piecewise training to learn both unary and pairwise potentials end-to-end. Our approach trains a network on multi-view consistency and fuses the results from multiple view points. It is complementary to the above single-view CNN approaches.

Semantic SLAM. In the domain of semantic SLAM, Salas-Moreno et al. [14] developed the SLAM++ algorithm to perform RGB-D tracking and mapping at the object instance level. Hermans et al. [15] proposed 3D semantic mapping for indoor RGB-D sequences based on RGB-D

visual odometry and a random forest classifier that performs semantic image segmentation. The individual frame segmentations are projected into 3D and smoothed using a CRF on the point cloud. Stückler et al. [16] perform RGB-D SLAM and probabilistically fuse the semantic segmentations of individual frames obtained with a random forest in multi-resolution voxel maps. Recently, Armeni et al. [17] propose a hierarchical parsing method for large-scale 3D point clouds of indoor environments. They first separate point clouds into disjoint spaces, *i.e.*, single rooms, and then further cluster points at the object level according to handcrafted features.

Multi-View Semantic Segmentation. In contrast to the popularity of CNNs for image-based segmentation, it is less common to apply CNNs for semantic segmentation on multi-view 3D reconstructions. Recently, Riegler et al. [18] apply 3D CNNs on sparse octree data structures to perform semantic segmentation on voxels. Nevertheless, the volumetric representations may discard details which are present at the original image resolution. McCormac et al. [19] proposed to fuse CNN semantic image segmentations on a 3D surfel map [20]. He et al. [21] propose to fuse CNN semantic segmentations from multiple views in video using superpixels and optical flow information. In contrast to our approach, these methods do not impose multi-view consistency during CNN training and cannot leverage the view-point invariant features learned by our network. Kundu et al. [22] extend dense CRFs to videos by associating pixels temporally using optical flow and optimizing their feature similarity. Closely related to our approach for enforcing multi-view consistency is the approach by Su et al. [23] who investigate the task of 3D shape recognition. They render multiple views onto 3D shape models which are fed into a CNN feature extraction stage that is shared across views. The features are max-pooled across view-points and fed into a second CNN stage that is trained for shape recognition. Our approach uses multi-view pooling for the task of semantic segmentation and is trained using realistic imagery and SLAM pose estimates. Our trained network is able to classify single views, but we demonstrate that multi-view fusion using the network trained on multi-view consistency improves segmentation performance over single-view trained networks.

III. CNN ARCHITECTURE FOR SEMANTIC SEGMENTATION

In this section, we detail the CNN architecture for semantic segmentation of each RGB-D image of a sequence. We base our encoder-decoder CNN on FuseNet [1] which learns rich features from RGB-D data. We enhance the approach with multi-scale loss minimization, which gains additional improvement in segmentation performance.

A. RGB-D Semantic Encoder-Decoder

Fig. 2 illustrates our CNN architecture. The network follows an encoder-decoder design, similar to previous work on semantic segmentation [7]. The encoder extracts a hierarchy of features through convolutional layers and aggregates spatial information by pooling layers to increase the

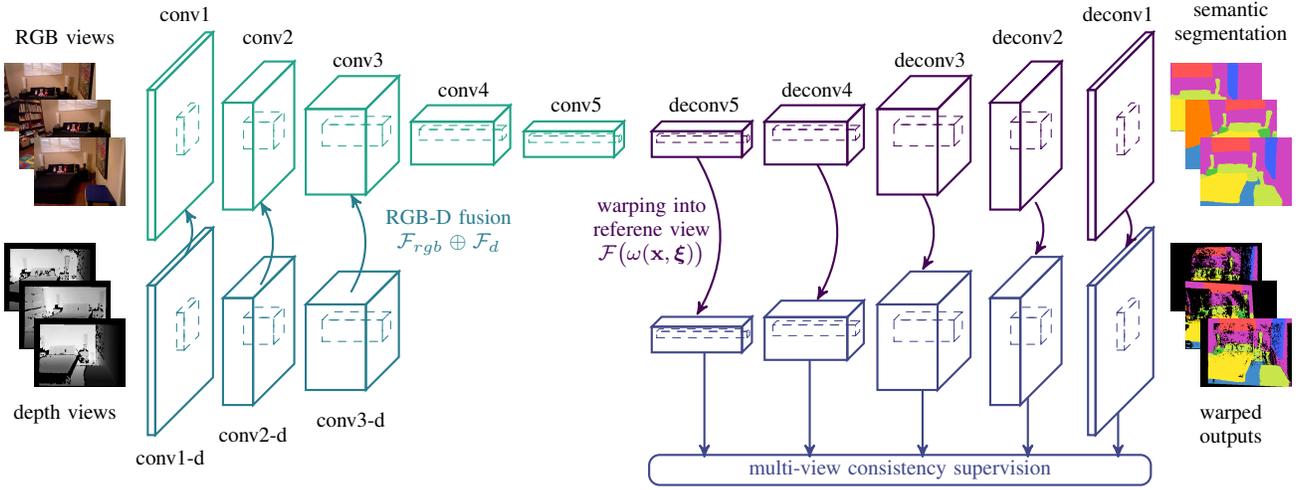


Fig. 2: The CNN encoder-decoder architecture used in our approach. Input to the network are RGB-D sequences with corresponding poses from SLAM trajectory. The encoder contains two branches to learn features from RGB-D data as inspired by FuseNet [1]. The obtained low-resolution high-dimension feature maps are successively refined through deconvolutions in the decoder. We warp feature maps into a common reference view and enforce multi-view consistency with various constraints. The network is trained in a deeply-supervised manner where loss is computed at all scales of the decoder.

receptive field. The encoder outputs low-resolution high-dimensional feature maps, which are upsampled back to the input resolution by the decoder through layers of memorized unpooling and deconvolution. Following FuseNet [1], the network contains two branches to learn features from RGB (\mathcal{F}_{rgb}) and depth (\mathcal{F}_d), respectively. The feature maps from the depth branch are consistently fused into the RGB branch at each scale. We denote the fusion by $\mathcal{F}_{rgb} \oplus \mathcal{F}_d$.

The semantic label set is denoted as $\mathcal{L} = \{1, 2, \dots, K\}$ and the category index is indicated with subscript j . Following notation convention, we compute the classification score $\mathcal{S} = (s_1, s_2, \dots, s_K)$ at location \mathbf{x} and map it to the probability distribution $\mathcal{P} = (p_1, p_2, \dots, p_K)$ with the softmax function $\sigma(\cdot)$. Network inference obtains the probability

$$p_j(\mathbf{x}, \mathcal{W} | \mathcal{I}) = \sigma(s_j(\mathbf{x}, \mathcal{W})) = \frac{\exp(s_j(\mathbf{x}, \mathcal{W}))}{\sum_k^K \exp(s_k(\mathbf{x}, \mathcal{W}))}, \quad (1)$$

of all pixels \mathbf{x} in the image for being labelled as class j , given input RGB-D image \mathcal{I} and network parameters \mathcal{W} .

We use the cross-entropy loss to learn network parameters for semantic segmentation from ground-truth annotations l_{gt} ,

$$L(\mathcal{W}) = -\frac{1}{N} \sum_i^N \sum_j^K \mathbb{1}[j = l_{gt}] \log p_j(\mathbf{x}_i, \mathcal{W} | \mathcal{I}), \quad (2)$$

where N is the number of pixels. This loss minimizes the Kullback-Leibler (KL) divergence between predicted distribution and the ground-truth, assuming the ground-truth has a one-hot distribution on the true label.

B. Multi-Scale Deep Supervision

The encoder of our network contains five 2×2 pooling layers and downsamples the input resolution by a factor of 32. The decoder learns to refine the low resolution back to the original one with five memorized unpooling followed

by deconvolution. In order to guide the decoder through the successive refinement, we adopt the deeply supervised learning method [24], [25] and compute the loss for all upsample scales. For this purpose, we append a classification layer at each deconvolution scale and compute the loss for the respective resolution of ground-truth which is obtained through stochastic pooling [26] over the full resolution annotation (see Fig. 3 for an example).

IV. MULTI-VIEW CONSISTENT LEARNING AND PREDICTION

While CNNs have been shown to obtain the state-of-the-art semantic segmentation performances for many datasets, most of these studies focus on single views. When observing a scene from a moving camera such as on a mobile robot, the system obtains multiple different views onto the same objects. The key innovation of this work is to explore the use of temporal multi-view consistency within RGB-D sequences for CNN training and prediction. For this purpose, we perform 3D data association by warping multiple frames into a common reference view. This then enables us to impose multi-view constraints during training. In this section, we describe several variants of such constraints. Notably, these methods can also be used at test time to fuse predictions from multiple views in a reference view.

A. Multi-view Data Association Through Warping

Instead of single-view training, we train our network on RGB-D sequences with poses estimated by a SLAM algorithm. We define each training sequence to contain one reference view \mathcal{I}_k with ground-truth semantic annotations and several overlapping views \mathcal{I}_i that are tracked towards \mathcal{I}_k . The relative poses ξ of the neighboring frames are estimated through tracking algorithms such as DVO SLAM [27]. In

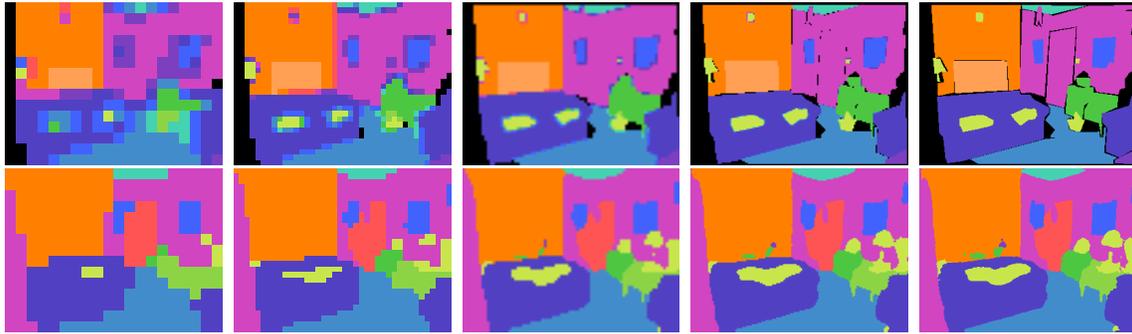


Fig. 3: Example of multi-scale ground-truth and predictions. Upper row: successive subsampled of ground-truth annotation obtained through stochastic pooling. Lower row: CNN prediction on each scale. The resolutions are coarse to fine from left to right with 20×15 , 40×30 , 80×60 , 160×120 and 320×240 .

order to impose temporal consistency, we adopt the warping concept from multi-view geometry to associate pixels between view points and introduce warping layers into our CNN. The warping layers synthesize CNN output in a reference view from a different view at any resolution by sampling given a known pose estimate and the known depth. The warping layers can be viewed as a variant of spatial transformers [28] with fixed transformation parameters.

We now formulate the warping. Given 2D image coordinate $\mathbf{x} \in \mathbb{R}^2$, the warped pixel location

$$\mathbf{x}^\omega := \omega(\mathbf{x}, \boldsymbol{\xi}) = \pi(\mathbf{T}(\boldsymbol{\xi}) \pi^{-1}(\mathbf{x}, Z_i(\mathbf{x}))), \quad (3)$$

is determined through the warping function $\omega(\mathbf{x}, \boldsymbol{\xi})$ which transforms the location from one camera view to the other using the depth $Z_i(\mathbf{x})$ at pixel \mathbf{x} in image \mathcal{I}_i and the SLAM pose estimate $\boldsymbol{\xi}$. The functions π and its inverse π^{-1} project homogeneous 3D coordinates to image coordinates and vice versa, while $\mathbf{T}(\boldsymbol{\xi})$ denotes the homogeneous transformation matrix derived from pose $\boldsymbol{\xi}$.

Using this association by warping, we synthesize the output of the reference view by sampling the feature maps of neighboring views using bilinear interpolation. Since the interpolation is differentiable, it is straight-forward to back-propagate gradients through the warping layers. With a slight abuse of notation, we denote the operation of synthesizing the layer output \mathcal{F} given the warping by $\mathcal{F}^\omega := \mathcal{F}(\omega(\mathbf{x}, \boldsymbol{\xi}))$.

We also apply deep supervision when training for multi-view consistency through warping. As shown in Fig. 2, feature maps at each resolution of the decoder are warped into the common reference view. Despite the need to perform warping at multiple scales, the warping grid is only required to be computed once at the input resolution, and is normalized to the canonical coordinates within the range of $[-1, 1]$. The lower-resolution warping grids can then be efficiently generated through average pooling layers.

B. Consistency Through Warp Augmentation

One straight-forward solution to enforce multi-view segmentation consistency is to warp the predictions of neighboring frames into the ground-truth annotated keyframe and computing a supervised loss there. This approach can be interpreted as a type of data augmentation using the

available nearby frames. We implement this consistency method by warping the keyframe into neighboring frames, and synthesize the classification score of the nearby frame from the keyframe’s view point. We then compute the cross-entropy loss on this synthesized prediction. Within RGB-D sequences, objects can appear at various scales, image locations, view perspective, color distortion given uncontrolled lighting and shape distortion given rolling shutters of RGB-D cameras. Propagating the keyframe annotation into other frames implicitly regulates the network predictions to be invariant under these transformations.

C. Consistency Through Bayesian Fusion

Given a sequence of measurements and predictions at test time, Bayesian fusion is frequently applied to aggregate the semantic segmentations of individual views. Let us denote the semantic labelling of a pixel by y and its measurement in frame i by z_i . We use the notation z^i for the set of measurements up to frame i . According to Bayes rule,

$$p(y | z^i) = \frac{p(z_i | y, z^{i-1}) p(y | z^{i-1})}{p(z_i | z^{i-1})} \quad (4)$$

$$= \eta_i p(z_i | y, z^{i-1}) p(y | z^{i-1}). \quad (5)$$

Suppose measurements satisfy the *i.i.d.* condition, i.e. $p(z_i | y, z^{i-1}) = p(z_i | y)$, and equal a-priori probability for each class, then Equation (4) simplifies to

$$p(y | z^i) = \eta_i p(z_i | y) p(y | z^{i-1}) = \prod_i \eta_i p(z_i | y). \quad (6)$$

Put simple, Bayesian fusion can be implemented by taking the product over the semantic labelling likelihoods of individual frame at a pixel and normalizing the product to yield a valid probability distribution. This process can also be implemented recursively on a sequence of frames.

When training our CNN for multi-view consistency using Bayesian fusion, we warp the predictions of neighboring frames into the keyframe using the SLAM pose estimate. We obtain the fused prediction at each keyframe pixel by summing up the unnormalized log labelling likelihoods instead of the individual frame softmax outputs. Applying softmax on the sum of log labelling likelihoods yields the

fused labelling distribution. This is equivalent to Eq. (6) since

$$\frac{\prod_i p_{i,j}^\omega}{\sum_k \prod_i p_{i,k}^\omega} = \frac{\prod_i \sigma(s_{i,j}^\omega)}{\sum_k \prod_i \sigma(s_{i,k}^\omega)} = \sigma\left(\sum_i s_{i,j}^\omega\right), \quad (7)$$

where $s_{i,j}^\omega$ and $p_{i,j}^\omega$ denote the warped classification scores and probabilities, respectively, and $\sigma(\cdot)$ is the softmax function as defined in Equation (1).

D. Consistency Through Multi-View Max-Pooling

While Bayesian fusion provides an approach to integrate several measurements in the probability space, we also explore direct fusion in the feature space using multi-view max-pooling of the warped feature maps. We warp the feature maps preceding the classification layers at each scale in our decoder into the keyframe and apply max-pooling over corresponding feature activations at the same warped location to obtain a pooled feature map in the keyframe,

$$\mathcal{F} = \text{max_pool}(\mathcal{F}_1^\omega, \mathcal{F}_2^\omega, \dots, \mathcal{F}_N^\omega). \quad (8)$$

The fused feature maps are classified and the resulting semantic segmentation is compared to the keyframe ground-truth for loss calculation.

V. EVALUATION

We evaluate our proposed approach using the NYUDv2 RGB-D dataset [29]. The dataset provides 1449 pixelwise annotated RGB-D images capturing various indoor scenes, and is split into 795 frames for training/validation (trainval) and 654 frames for testing. The original sequences that contain these 1449 images are also available with NYUDv2, whereas sequences are unfortunately not available for other large RGB-D semantic segmentation datasets. Using DVO-SLAM [27], we determine the camera poses of neighboring frames around each annotated keyframe to obtain multi-view sequences. This provides us with in total 267,675 RGB-D images, despite that tracking fails for 30 out of 1449 keyframes. Following the original trainval/test split, we use 770 sequences with 143,670 frames for training and 649 sequences with 124,005 frames for testing. For benchmarking, our method is evaluated for the 13-class [2] and 40-class [30] semantic segmentation tasks. We use the raw depth images without inpainted missing values.

A. Training Details

We implemented our approach using the Caffe framework [31]. For all experiments, the network parameters are initialized as follows. The convolutional kernels in the encoder are initialized with the pretrained 16-layer VGGNet [32] and the deconvolutional kernels in the decoder are initialized using He’s method [33]. For the first layer of the depth encoder, we average the original three-channel VGG weights to obtain a single-channel kernel. We train the network with stochastic gradient descent (SGD) [34] with 0.9 momentum and 0.0005 weight decay. The learning rate is set to 0.001 and decays by a factor of 0.9 every 30k iterations. All the images are resized to a resolution of 320×240 pixels as input to the network and the predictions are also up to

TABLE I: Single-view semantic segmentation accuracy of our network in comparison to the state-of-the-art methods for NYUDv2 13-class and 40-class segmentation tasks.

	methods	input	pixelwise	classwise	IoU
NYUDv2 13 classes	Coupric et al. [2]	RGB-D	52.4	36.2	-
	Hermans et al. [15]	RGB-D	54.2	48.0	-
	SceneNet [35]	DHA	67.2	52.5	-
	Eigen et al. [8]	RGB-D-N	75.4	66.9	52.6
	FuseNet-SF3 [1]	RGB-D	75.8	66.2	54.2
	MVCNet-Mono	RGB-D	77.6	68.7	56.9
	MVCNet-Augment	RGB-D	77.6	69.3	57.2
	MVCNet-Bayesian	RGB-D	77.8	69.4	57.3
	MVCNet-MaxPool	RGB-D	77.7	69.5	57.3
	NYUDv2 40 classes	RCNN [3]	RGB-HHA	60.3	35.1
FCN-16s [5]		RGB-HHA	65.4	46.1	34.0
Eigen et al. [8]		RGB-D-N	65.6	45.1	34.1
FuseNet-SF3 [1]		RGB-D	66.4	44.2	34.0
Context-CRF [13]		RGB	67.6	49.6	37.1
MVCNet-Mono		RGB-D	68.6	48.7	37.6
MVCNet-Augment		RGB-D	68.6	49.9	38.0
MVCNet-Bayesian		RGB-D	68.4	49.5	37.4
MVCNet-MaxPool		RGB-D	69.1	50.1	38.0

this scale. To downsample, we use cubic interpolation for RGB images and nearest-neighbor interpolation for depth and label images. During training, we use a minibatch of 6 that comprises two sequences, with one keyframe and two tracking frames for each sequence. We apply random shuffling after each epoch for both inter and intra sequences. The network is trained until convergence. We observed that multi-view CNN training does not require significant extra iterations for convergence. For multi-view training, we sample from the nearest frames first and include 10 further-away frames every 5 epochs. By this, we alleviate that typically tracking errors accumulate and image overlap decreases as the camera moves away from the keyframe.

B. Evaluation Criteria

We measure the semantic segmentation performance with three criteria: global pixelwise accuracy, average classwise accuracy and average intersection-over-union (IoU) scores. These three criteria can be calculated from the confusion matrix. With K classes, each entry of the $K \times K$ confusion matrix c_{ij} is the total amount of pixels belonging to class i that are predicted to be class j . The global pixelwise accuracy is computed by $\sum_i c_{ii} / \sum_{ij} c_{ij}$, the average classwise accuracy is computed by $\frac{1}{K} \sum_i (c_{ii} / \sum_j c_{ij})$, and the average IoU score is calculated by $\frac{1}{K} \sum_i (c_{ii} / (\sum_i c_{ij} + \sum_j c_{ij} - c_{ii}))$.

C. Single Frame Segmentation

In a first set of experiments, we evaluate the performance of several variants of our network for direct semantic segmentation of individual frames. This means we do not fuse predictions from nearby frames to obtain the final prediction in a frame. We predict semantic segmentation with our trained models on the 654 test images of the NYUDv2 dataset and compare our methods with state-of-art approaches. The results are shown in Table I. Unless otherwise stated, we take the results from the original papers

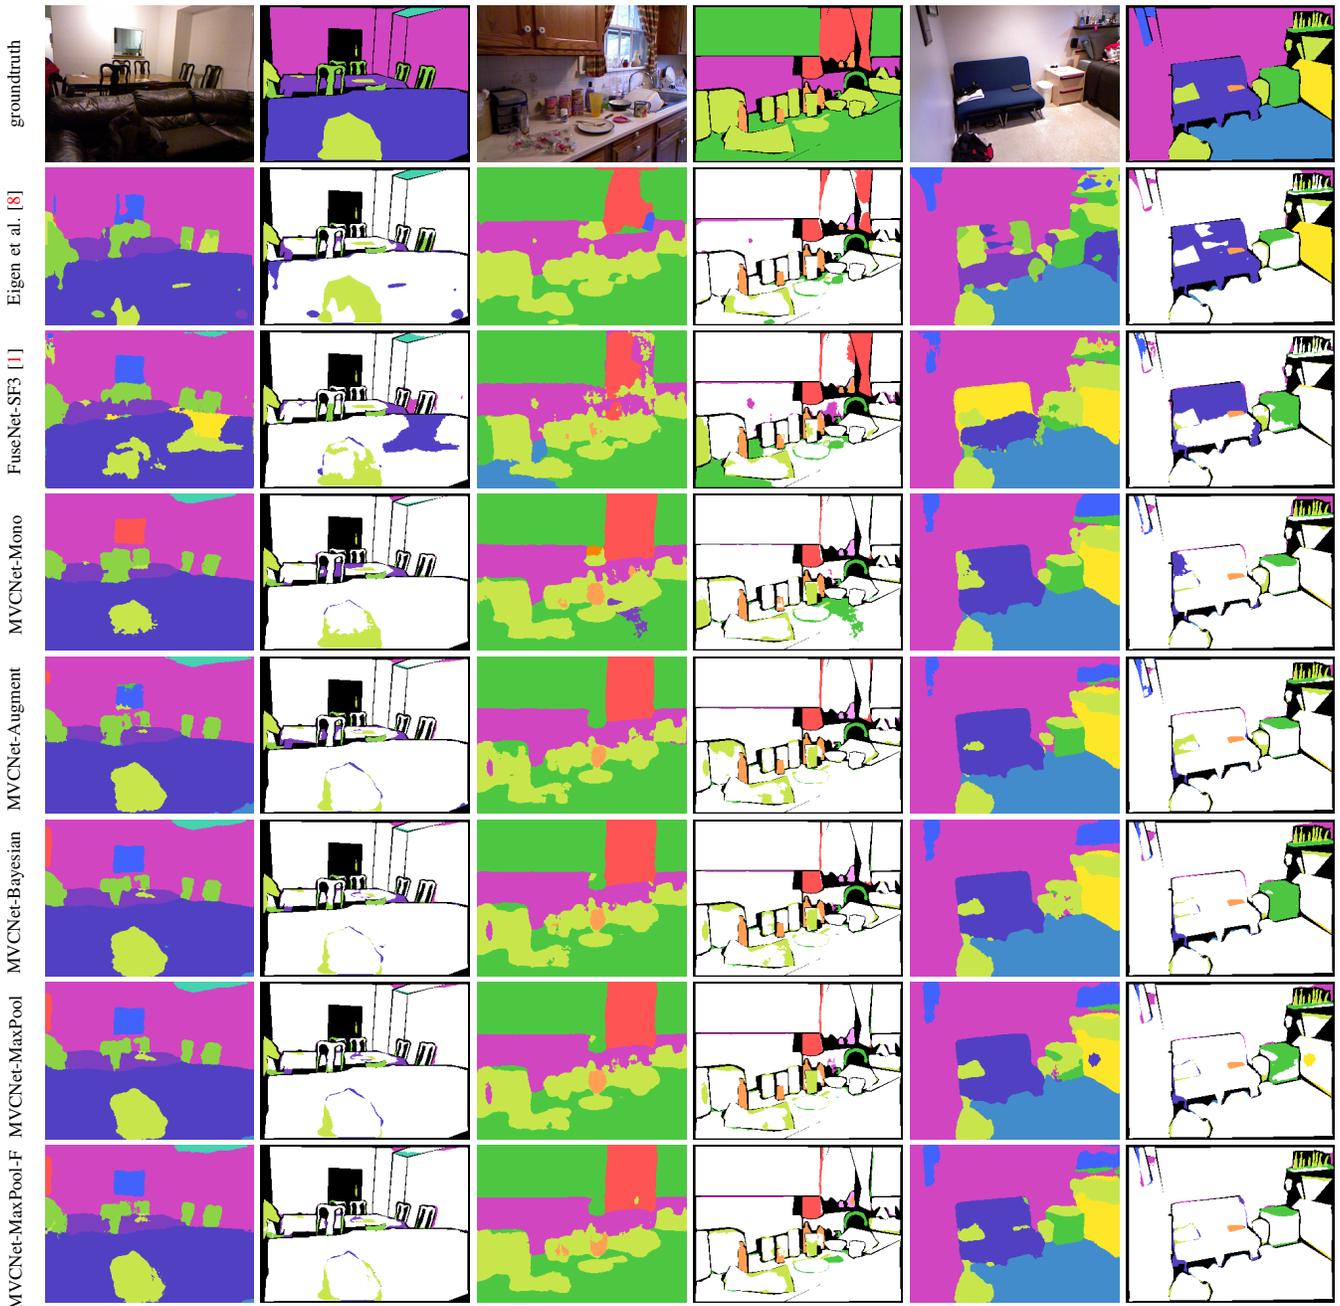


Fig. 4: Qualitative semantic segmentation results of our methods and several state-of-the-art baselines on NYUDv2 13-class segmentation (see Table III for color coding, left columns: semantic segmentation, right columns: falsely classified pixels, black is void). Our multi-view consistency trained models produce more accurate and homogeneous results than single-view methods. Bayesian fusion further improves segmentation quality (e.g. MVCNet-MaxPool-F).

for comparison and report their best results (i.e. SceneNet-FT-NYU-DO-DHA model for SceneNet [35], VGG-based model for Eigen et al. [8]). The result of Hermans et al. [15] is obtained after applying a dense CRF [36] for each image and in-between neighboring 3D points to further smoothen their results. We also remark that the results reported here for the Context-CRF model are finetuned on NYUDv2 like in our approach to facilitate comparison. Furthermore, the network output is refined using a dense CRF [36] which is claimed to increase the accuracy of the network by

approximately 2%. The results for FuseNet-SF3 are obtained by our own implementation. Our baseline model MVCNet-Mono is trained without multi-view consistency, which amounts to FuseNet with multiscale deeply supervised loss at decoder. However, we apply single image augmentation to train the FuseNet-SF3 and MVCNet-Mono with random scaling between $[0.8, 1.2]$, random crop and mirror. This data augmentation is not used for multi-view training. Nevertheless, our results show that the different variants of multi-view consistency training outperform the state-of-art

TABLE II: Multi-view segmentation accuracy of our network using Bayesian fusion for NYUDv2 13-class and 40-class segmentation.

	methods	pixelwise	classwise	IoU
NYUDv2 13 classes	FuseNet-SF3 [1]	77.19	67.46	56.01
	MVCNet-Mono	78.70	69.61	58.29
	MVCNet-Augment	78.94	70.48	58.93
	MVCNet-Bayesian	79.13	70.48	59.04
	MVCNet-MaxPool	79.13	70.59	59.07
NYUDv2 40 classes	FuseNet-SF3 [1]	67.74	44.92	35.36
	MVCNet-Mono	70.03	49.73	39.12
	MVCNet-Augment	70.34	51.73	40.19
	MVCNet-Bayesian	70.24	51.18	39.74
	MVCNet-MaxPool	70.66	51.78	40.07

methods for single image semantic segmentation. Overall, multi-view max-pooling (MVCNet-MaxPool) has a small advantage over the other multi-view consistency training approaches (MVCNet-Augment and MVCNet-Bayesian).

D. Multi-View Fused Segmentation

Since we train on sequences, in the second set of experiment, we also evaluate the fused semantic segmentation over the test sequences. The number of fused frames is fixed to 50, which are uniformly sampled over the entire sequence. Due to the lack of ground-truth for neighboring frames, we fuse the prediction of neighboring frames in the keyframes using Bayesian fusion according to Equation (7). This fusion is typically applied for semantic mapping using RGB-D SLAM. The results are shown in Table II. Bayesian multi-view fusion improves the semantic segmentation by approx. 2% on all evaluation measures towards single-view segmentation. Also, the training for multi-view consistency achieves a stronger gain over single-view training (MVCNet-Mono) when fusing segmentations compared to single-view segmentation. This performance gain is observed in the qualitative results in Fig. 4. It can be seen that our multi-view consistency training and Bayesian fusion produces more accurate and homogeneous segmentations. Fig. 5 shows typical challenging cases for our model.

We also compare classwise and average IoU scores for 13-class semantic segmentation on NYUDv2 in Table III. The results of Eigen et al. [8] are from their publicly available model tested on 320×240 resolution. The results demonstrate that our approach gives high performance gains across all occurrence frequencies of the classes in the dataset.

VI. CONCLUSION

In this paper we propose methods for enforcing multi-view consistency during the training of CNN models for semantic RGB-D image segmentation. We base our CNN design on FuseNet [1], a recently proposed CNN architecture in an encoder-decoder scheme for semantic segmentation of RGB-D images. We augment the network with multi-scale loss supervision to improve its performance. We present and evaluate three different approaches for multi-view consistency training. Our methods use an RGB-D SLAM trajectory estimate to warp semantic segmentations or feature maps

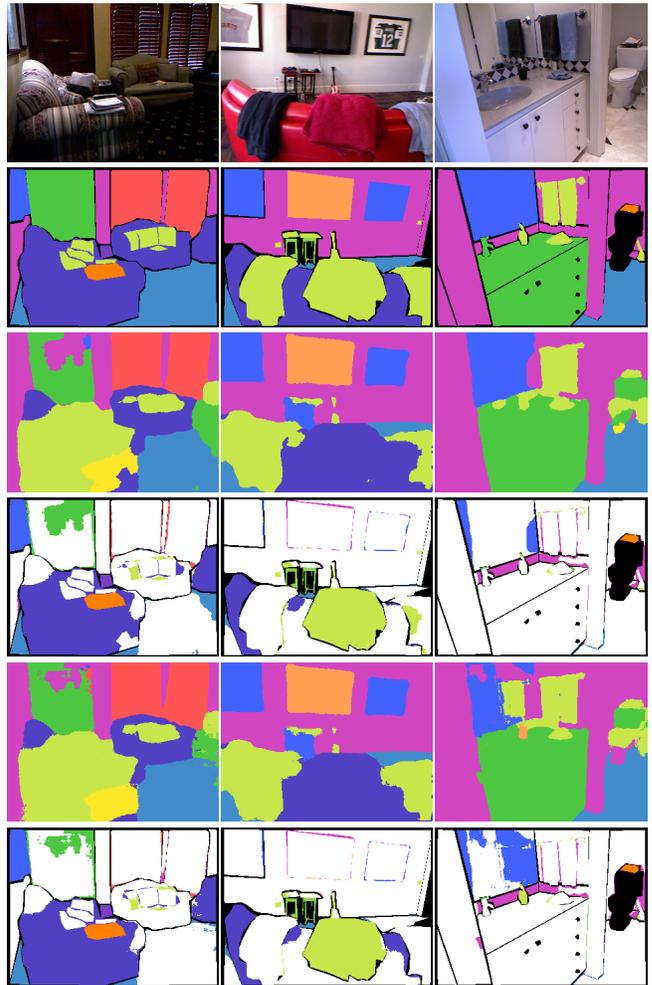


Fig. 5: Challenging cases for MVCNet-MaxPool-F (top to bottom: RGB image, ground-truth, single-view prediction on keyframe, multi-view prediction fused in keyframe). On the left, the network fails to classify the objects for all frames. In the middle, the network makes some errors in single-view prediction, but through multi-view fusion, some mistakes are corrected. On the right, multi-view fusion degenerates performance due to the mirror reflections.

from one view point to another. Multi-view max-pooling of feature maps overall provides the best performance gains in single-view segmentation and fusion of multiple views.

We demonstrate the superior performance of multi-view consistency training and Bayesian fusion on the NYUDv2 13-class and 40-class semantic segmentation benchmark. All multi-view consistency training approaches outperform single-view trained baselines. They are key to boosting segmentation performance when fusing network predictions from multiple view points during testing. On NYUDv2, our model sets a new state-of-the-art performance using an end-to-end trained network for single-view predictions as well as multi-view fused semantic segmentation without further postprocessing stages such as dense CRFs. In future work, we want to further investigate integration of our approach in a semantic SLAM system, for example, through coupling of pose tracking and SLAM with our semantic predictions.

TABLE III: NYUDv2 13-class semantic segmentation IoU scores. Our method achieves best per-class accuracy and average IoU.

method		bed	objects	chair	furniture	ceiling	floor	decorat.	sofa	table	wall	window	books	TV	average accuracy
class frequency		4.08	7.31	3.45	12.71	1.47	9.88	3.40	2.84	3.42	24.57	4.91	2.78	0.99	
single-view	Eigen et al. [8]	56.71	38.29	50.23	54.76	64.50	89.76	45.20	47.85	42.47	74.34	56.24	45.72	34.34	53.88
	FuseNet-SF3 [1]	61.52	37.95	52.67	53.97	64.73	89.01	47.11	57.17	39.20	75.08	58.06	37.64	29.77	54.14
	MVCNet-Mono	65.27	37.82	54.09	59.39	65.26	89.15	49.47	57.00	44.14	75.31	57.22	49.21	36.14	56.88
	MVCNet-Augment	65.33	38.30	54.15	59.54	67.65	89.26	49.27	55.18	43.39	74.59	58.46	49.35	38.84	57.18
	MVCNet-Bayesian	65.76	38.79	54.60	59.28	67.58	89.69	48.98	56.72	42.42	75.26	59.55	49.27	36.51	57.26
MVCNet-MaxPool	65.71	39.10	54.59	59.23	66.41	89.94	49.50	56.30	43.51	75.33	59.11	49.18	37.37	57.33	
multi-view	FuseNet-SF3 [1]	64.95	39.62	55.28	55.90	64.99	89.88	47.99	60.17	42.40	76.24	59.97	39.80	30.91	56.01
	MVCNet-Mono	67.11	40.14	56.39	60.90	66.07	89.77	50.32	59.49	46.12	76.51	59.03	48.80	37.13	58.29
	MVCNet-Augment	68.22	40.04	56.55	61.82	67.88	90.06	50.85	58.00	45.98	75.85	60.43	50.50	39.89	58.93
	MVCNet-Bayesian	68.38	40.87	57.10	61.84	67.98	90.64	50.05	59.70	44.73	76.50	61.75	51.01	36.99	59.04
	MVCNet-MaxPool	68.09	41.58	56.88	61.56	67.21	90.64	50.69	59.73	45.46	76.68	61.28	50.60	37.51	59.07

REFERENCES

- [1] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "Fusenet: incorporating depth into semantic segmentation via fusion-based cnn architecture," in *Asian Conf. on Computer Vision (ACCV)*, 2016.
- [2] C. Couprie, C. Farabet, L. Najman, and Y. Lecun, *Indoor semantic segmentation using depth information*. 2013.
- [3] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik, "Learning rich features from RGB-D images for object detection and segmentation," in *Europ. Conf. on Computer Vision (ECCV)*, 2014.
- [4] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [5] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [6] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy layer-wise training of deep networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- [7] H. Noh, S. Hong, and B. Han, "Learning deconvolution network for semantic segmentation," in *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 1520–1528, 2015.
- [8] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [9] G. Lin, A. Milan, C. Shen, and I. Reid, "RefineNet: Multi-path refinement networks for high-resolution semantic segmentation," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [10] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," in *Int. Conf. on Learning Representations (ICLR)*, 2016.
- [11] Z. Wu, C. Shen, and A. van den Hengel, "Wider or deeper: Revisiting the resnet model for visual recognition," *CoRR abs/1611.10080*, 2016.
- [12] Z. Li, Y. Gan, X. Liang, Y. Yu, H. Cheng, and L. Lin, "LSTM-CF: Unifying context modeling and fusion with LSTMs for RGB-D scene labeling," in *Europ. Conf. on Computer Vision (ECCV)*, 2016.
- [13] G. Lin, C. Shen, A. van den Hengel, and I. D. Reid, "Exploring context with deep structured models for semantic segmentation," *CoRR*, vol. abs/1603.03183, 2016.
- [14] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. Kelly, and A. J. Davison, "SLAM++: Simultaneous localisation and mapping at the level of objects," *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [15] A. Hermans, G. Floros, and B. Leibe, "Dense 3D semantic mapping of indoor scenes from rgb-d images," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 2631–2638, 2014.
- [16] J. Stückler, B. Waldvogel, H. Schulz, and S. Behnke, "Dense real-time mapping of object-class semantics from RGB-D video," *J. of Real-Time Image Processing*, 2015.
- [17] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese, "3D semantic parsing of large-scale indoor spaces," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [18] G. Riegler, A. O. Ulusoy, and A. Geiger, "OctNet: Learning deep 3d representations at high resolutions," *CoRR*, vol. abs/1611.05009, 2016.
- [19] J. McCormac, A. Handa, A. J. Davison, and S. Leutenegger, "SemanticFusion: Dense 3D semantic mapping with convolutional neural networks," *CoRR*, vol. abs/1609.05130, 2016.
- [20] T. Whelan, R. F. Salas-Moreno, B. Glocker, A. J. Davison, and S. Leutenegger, "ElasticFusion: Real-time dense SLAM and light source estimation," *Int. J. of Robotics Research (IJRR)*, 2016.
- [21] Y. He, W. Chiu, M. Keuper, and M. Fritz, "STD2P: Rgb-d semantic segmentation using spatio-temporal data driven pooling," in *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [22] A. Kundu, V. Vineet, and V. Koltun, "Feature space optimization for semantic video segmentation," in *IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] H. Su, S. Maji, E. Kalogerakis, and E. G. Learned-Miller, "Multi-view convolutional neural networks for 3d shape recognition," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [24] C. Lee, S. Xie, P. W. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in *Proc. of the 18th Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [25] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox, "Flownet: Learning optical flow with convolutional networks," in *The IEEE Int. Conf. on Computer Vision (ICCV)*, December 2015.
- [26] M. Zeiler and R. Fergus, *Stochastic pooling for regularization of deep convolutional neural networks*. 2013.
- [27] C. Kerl, J. Sturm, and D. Cremers, "Dense visual SLAM for RGB-D cameras," in *IEEE/RSJ Intelligent Robots and Systems (IROS)*, 2013.
- [28] M. Jaderberg, K. Simonyan, A. Zisserman, and k. kavukcuoglu, "Spatial transformer networks," in *Advances in Neural Information Processing Systems (NIPS)* (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, eds.), 2015.
- [29] P. K. Nathan Silberman, Derek Hoiem and R. Fergus, "Indoor segmentation and support inference from RGBD images," in *Europ. Conf. on Computer Vision (ECCV)*, 2012.
- [30] S. Gupta, P. Arbelaez, and J. Malik, "Perceptual organization and recognition of indoor scenes from RGB-D images," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [31] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.
- [32] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [34] L. Bottou, "Stochastic gradient descent tricks," in *Neural Networks: Tricks of the Trade*, pp. 421–436, Springer, 2012.
- [35] A. Handa, V. Patraucean, V. Badrinarayanan, S. Stent, and R. Cipolla, "Scenenet: Understanding real world indoor scenes with synthetic data," in *IEEE Comp. Vision and Pattern Recognition (CVPR)*, 2016.
- [36] P. Krähenbühl and V. Koltun, "Efficient inference in fully connected CRFs with Gaussian edge potentials," in *Advances in Neural Information Processing Systems (NIPS)*, 2011.