# In pixels we trust: From Pixel Labeling to Object Localization and Scene Categorization

Carlos Herranz-Perdiguero<sup>1</sup>, Carolina Redondo-Cabrera<sup>1</sup> and Roberto J. López-Sastre<sup>1</sup>

Abstract-While there has been significant progress in solving the problems of image pixel labeling, object detection and scene classification, existing approaches normally address them separately. In this paper, we propose to tackle these problems from a bottom-up perspective, where we simply need a semantic segmentation of the scene as input. We employ the DeepLab architecture, based on the ResNet deep network, which leverages multi-scale inputs to later fuse their responses to perform a precise pixel labeling of the scene. This semantic segmentation mask is used to localize the objects and to recognize the scene, following two simple yet effective strategies. We evaluate the benefits of our solutions, performing a thorough experimental evaluation on the NYU Depth V2 dataset. Our approach achieves a performance that beats the leading results by a significant margin, defining the new state of the art in this benchmark for the three tasks comprising the scene understanding: semantic segmentation, object detection and scene categorization.

#### I. INTRODUCTION

Visual semantic scene understanding has become a crucial capability in many applications, such as autonomous driving [1], indoor navigation [2] or robotic manipulation [3]. In fact, it is the capability that enables robots to truly interact with the environment with a minimal intervention or number of sensors, due to the rich information contained in images.

Its goal consists in obtaining as much semantic knowledge of a given scene as possible. This includes scene categorization (labeling the whole scene), object detection (predicting object locations by bounding boxes), and semantic segmentation (labeling each pixel with a category).

In recent years, there has been large progress in solving all these tasks (*e.g.* [4], [5], [6], [7]). However, most current state-of-the-art approaches propose to solve these problems separately. In this paper, we demonstrate that it results beneficial to address these problems following a bottomup approach, which mostly relies on a precise semantic segmentation of the scene.

As it is shown in Figure 1, our solution starts with a pixel labeling deep network. We base our design on the DeepLab architecture [4], using ResNet deep networks [8], enhancing its performance with a multi-scale supervision. We then propose two simple yet precise approaches to perform



Fig. 1. We train a multi-scale DeepLab architecture (ResNet based) to perform a precise pixel labeling of the scene. The key innovation is to leverage on this semantic segmentation to perform both the object localization and scene recognition, offering a fast and precise integral solution, which can be embedded in any mobile robotic platform.

both the object localization and the scene recognition tasks considering only the scene semantic segmentation as input.

Overall, the main contributions of our work are as follows:

- In principle, the core of the proposed architecture relies on a semantic segmentation network. We follow an intuitive extension of the DeepLab pixel labeling model [4], which consists in fusing the response of a pool of multi-scale ResNet based networks.
- 2) Starting from our precise semantic segmentation, we demonstrate that it is possible to successfully classify the scenes. Technically, we introduce two classifiers based solutions. First, we design a simple approach that uses SVMs with additive kernels, using as features the histograms of class labels computed from the semantic segmentation. Second, by viewing each semantic segmentation as a one-hot binary vector, indicating the presence/absence of the corresponding classes, we propose to perform the categorization of the scene with a Linear SVM. We show that these two surprisingly simple and fast solutions can surpass prior state-of-theart results reported by complex deep networks trained to directly classify the scenes.
- 3) We finally showcase the potential of our precise semantic segmentations via the task of object localization. We just need to identify in the pixel labeling mask for the whole scene, the connected components that enclose the objects, allowing us to localize them.
- 4) Without bells and whistles, all our models surpass all previous state-of-the-art results on the challenging NYU-Depth V2 (NYUD2) dataset [9] for the three tasks comprising scene understanding: scene categorization, object detection and semantic segmentation.

<sup>\*</sup>This work is supported by project PREPEATE, with reference TEC2016-80326-R, of the Spanish Ministry of Economy, Industry and Competitiveness. We gratefully acknowledge the support of NVIDIA Corporation with the donation of a GPU used for this research. Cloud computing resources were kindly provided through a Microsoft Azure for Research Award.

<sup>&</sup>lt;sup>1</sup>The authors are with GRAM research group, Department of Signal Theory and Communications, University of Alcalá, 28805, Alcalá de Henares, Spain {c.herranz,carolina.redondoc}@edu.uah.es, robertoj.lopez@.uah.es

The rest of the paper is organized as follows. Section II reviews the previous works. We technically detail our approaches in Section III. Section IV includes the experimental evaluation. We conclude in Section V.

### II. RELATED WORK

Recently, CNNs have become the state-of-the-art approach in order to solve scene understanding related tasks separately. Increasingly better results have been achieved in semantic segmentation (*e.g.* [2], [4], [10]), scene classification (*e.g.* [8], [11], [12], [13]) and object detection (*e.g.* [5], [14]).

However, not many works deal with the scene understanding problem as a whole. Among those that have done it, we find the earlier works which use handcrafted features, to capture some specific properties considered as representative. Probably, one of the first approaches in trying to get a major scene understanding can be found in [15]. They propose a hierarchical generative model that is able to classify the image (scene classification) and localize the objects that are part of it. Moreover, based on the relationship between the type of scene and the objects, the model provides pixel level segmentations (semantic segmentation). In [16], a holistic model capable of accomplishing all the three tasks is presented too. They make use of a hierarchical graphical model, a Conditional Random Field, with variables representing aspects such as the presence or absence of a class, the class labels, or the correctness of a certain candidate for object detection. That way, the model can learn jointly about all the different tasks comprising scene understanding. Gupta et al. [17], [18] propose a method to detect contours on depth images for semantic segmentation, and subsequently quantize the segmentation outputs as local features for scene classification. More recently, the use of multi-layered neural networks makes possible to learn features directly from large amounts of data to solve the scene understanding problem. For instance, Gupta et al. [19] use the R-CNN [20] detector on depth images to localize objects in indoor scenes. Then, the output from their object detectors is used with a superpixel classification framework to obtain a semantic segmentation of the scene.

Our work clearly differs from the references above. First, unlike [17], [19], [18], [11], [13], our approach does *not* need to use depth images, we solve the whole scene understanding task with just one modality, *i.e.* RGB images only. Second, we base the solution to the problem of scene understanding on an accurate semantic segmentation. Neither contour detectors [17], [18], nor object detections [19] are needed. Instead, we directly obtain the pixel labeling model with a multi-scale deep network. We then solve the tasks of object localization and scene categorization by simply using the semantic segmentation mask as a feature.

### **III. SCENE UNDERSTANDING**

# A. Scene Labeling Model

The use of Fully Convolutional Networks (FCNs), first introduced by Long *et al.* [21], has achieved great results by adapting image classification deep networks to the task

of semantic segmentation. In a nutshell, FCNs are mainly based on the replacement of the last fully-connected layers with convolutional ones.

Here, we propose the use of such networks, following the state-of-the-art publicly available model for scene labeling known as DeepLab [4]. This model introduces three mayor contributions: 1) atrous convolutions, as a way of increasing the field of view (FOV) of the convolutional filters; 2) Atrous Spatial Pyramid Pooling (ASPP) technique, which makes use of several parallel atrous convolutions at different rates, to extract features at various scales before merging them; and 3) a post-processing CRF to refine the scene labeling results given by the CNN under the premise that nearby and similar appearance pixels should be given the same label.

Technically, for our scene labeling model, we adopt this DeepLab architecture, incorporating a ResNet-101 based network [8] to it. For the indoor scene labeling problem we want to address with the NYUD2 dataset, we follow a joint multi-scale learning procedure. For that purpose, we process the input image with scaling factors  $\{0.5, 0.75, 1\}$ of the original resolution. Thereby, we actually have three complete ResNet-based networks working in parallel, as branches of the whole architecture (see Figure 1). At the end, we fuse the results at each pixel position by taking the maximum response given by the network for each scale. This procedure can be seen as a way of increasing and improving the contextual information that the network obtains from the images. The effect when processing smaller scales is similar to increasing the FOV of the filters since each filter covers a bigger part of the image, while larger scales are more detail-oriented. Therefore, the multi-scale supervision, together with the atrous convolution and ASPP, allows our approach to enhance its performance, by having different information of the surrounding area of each of the pixels in the image.

Technically, we are given a set of images,  $\{I_1, I_2, \ldots, I_N\}$ , with their corresponding semantic segmentation annotations,  $\{\phi_1, \phi_2, \ldots, \phi_N\}$ . We use the deep network described, and represent  $\hat{\theta}_i$  as the score map at the output of our network, and  $\hat{\phi}_i$  as our estimation for the semantic segmentation.

The multi-scale learning process solves the following optimization. For each training image, during learning, we apply a soft max function to map the scores given by our network to a probability distribution over the complete set of C classes for each of the k positions in the CNN output map:

$$\hat{p}_{kc} = \frac{\exp(\theta_{kc})}{\sum_{c'} \exp(\hat{\theta}_{kc'})}.$$
(1)

The result,  $\hat{p}_{kc}$ , is used to obtain the cross-entropy classification loss,  $\zeta$ , with which our model is optimized according to the following equation,

$$\zeta = \frac{-1}{K} \sum_{k=1}^{K} \log(\hat{p}_k, \phi), \tag{2}$$

so that our estimation for each pixel in the image,  $\hat{\phi}$ , will be the class for which the response of the CNN output map is maximum:

$$\hat{\phi} = \underset{C}{\operatorname{argmax}}(\hat{\theta}). \tag{3}$$

### B. Scene Classification

Since the breakthrough of deep learning and CNNs, scene classification has been carried out using deep networks, originally designed for the object recognition problem, which achieve state-of-the-art results (*e.g.* [8], [11], [12], [13]). On the contrary, in this work, we aim to use the semantic segmentation results to directly classify the scenes. Our results show that our surprisingly simple and fast approach can surpass the complex deep models trained for scene categorization.

Technically, we explore two classifier based solutions, which take as input features our scene labeling output. First, we propose to use the histograms of class labels obtained from the semantic segmentation. Thus, as descriptors for each image,  $\bar{F}_i$ , we take a *C*-bin histogram, hist, representing the number of pixels of each of the *C* classes in our semantic segmentation estimation  $\hat{\phi}$ :

$$F_i = \text{hist}(\hat{\phi}) = [f_{i1}, f_{i2}, ..., f_{iC}],$$
 (4)

that we normalize as follows,

$$\bar{F}_i = \operatorname{norm}(F_i), \tag{5}$$

where  $f_{ic}$  is the number of pixels of class c in image i, and norm corresponds to an L2 normalization. This idea is based on the fact that the objects that appear in an image should define the scene in which they are: the parts *form* the whole. With this histogram-like features, we propose to use SVMs with additive kernels [22], because their good performance working with this type of features. In the experiments, we try several of them, such as, intersection kernel,  $\chi^2$  or Jensen-Shannon.

On top of that, when constructing the histograms, we also explore the use of a two-level spatial pyramid pooling. With the aim of adding localization information to our descriptors, we divide the image in four parts and extract a histogram, following the procedure described in Equation 5, for each of them. Thereby, the new descriptors will be a concatenation of five C-size vectors: the original one, representing the full image, and one for each of the four second-level parts.

Our second model for scene categorization consists in learning a linear classifier with *one-hot vectors* computed from the semantic segmentation. For each image, we define a *one-hot vector*,  $G_i$ , as

$$G_i = [u(f_{i1}), u(f_{i2}), ..., u(f_{iC})] = [g_{i1}, g_{i2}, ..., g_{iC}], \quad (6)$$

with *u* being the function:

$$u(x) = \begin{cases} 1 & x > \delta \\ 0 & x \le \delta \end{cases}.$$
 (7)

The threshold parameter,  $\delta$ , is the minimum number of pixels needed to consider that a class appears in the image.

Therefore,  $g_1, ..., g_C \in \{0, 1\}$ . One-hot vectors just take into account whether a category is present or not in the image. For the categorization, we simply use a linear SVM. The use of one-hot vectors but with the additive kernels do not improve the results, according to our experiments.

### C. Object Detection Integration

Given the semantic segmentation, we present a model to classify the scene, but, can we also solve the object localization task? In this section we describe how to obtain the bounding boxes (BBs) from the pixel labeling masks, in order to evaluate the object detection accuracy of our model, closing, this way, the whole scene understanding problem.

Our solution is embarrassingly simple and fast. Given a semantic segmentation mask, and for each object category  $c \in C$ , we simply proceed to obtain a binary mask with the pixels that belong to the class of interest. We then trace region boundaries putting a tight BB around each connected component identified in the mask.

We need to assign a detection score to each of these BBs, and we again rely on the result of the semantic segmentation for this purpose. Given an estimated BB,  $BB_i^c$ , for the object category c, its corresponding object detection score  $s_i$  is obtained computing the mean of the confidence scores of the semantic segmentation output pixels within the BB that belong to the object class c. In other words, we technically follow this equation:

$$s_{i} = \frac{1}{P_{c}^{BB_{i}^{c}}} \sum_{j=1}^{BB_{i}^{c}} \hat{\theta}_{jc},$$
(8)

where  $P_c^{BB_i^c}$  is the total number of pixels of class c in bounding box  $BB_i^c$ .  $\hat{\theta}_{jc}$  represents the score values of class c of the output map provided by our semantic segmentation network.

#### **IV. EXPERIMENTS**

To evaluate the effectiveness of our models, we perform scene labeling, scene classification and object detection experiments on the challenging scene understanding NYU Depth V2 dataset [9] (NYUD2). The details of the experiments and the corresponding results are provided in the following sections.

## A. Experimental Setup

1) Implementation Details: We start from the publicly available DeepLab model implementation [4], which employs the Caffe deep learning framework [23]. We implement the ResNet-101 [8] deep architecture as the base network for the DeepLab model. This network is pretrained on the MS COCO dataset [24]. We replace the last layer with a soft max classifier layer, which has as many targets as the number of semantic classes of our task in the NYUD2 dataset. Our loss function is the sum of cross-entropy terms for each spatial position in the CNN output map (subsampled by 8 compared to the original image). All positions and labels are equally weighted in the overall loss function (except for unlabeled pixels in the ground truth annotations, which are ignored). We optimize the objective function using SGD. For fine tuning, the learning rate is initialized at 0.001 and we apply a *poly* learning rate policy (the learning rate is multiplied by  $(1 - \frac{iter}{max.iter})^{power})$  with *power* parameter fixed to 0.9. As data augmentation techniques, we apply mirroring, with a probability of 50%, and cropping: we feed our network with patches of size 385x385 instead of using the complete 480x640 images of the NYUD2 dataset. Our model for semantic segmentation is trained during 20K iterations with a batch size of 1.

With respect to *one-hot vectors* in scene classification, we set the threshold parameter,  $\delta$ , to 0.5% of the total number of pixels in the image, since small objects do not tend to be the most representative of a particular type of scene.

2) Dataset and Evaluation Metrics: In our experiments, we use the NYUD2 dataset [9], which consists of 1449 images. We use the publicly available split, which has 795 images for training and 654 images for testing.

For semantic segmentation, the original task proposed in [9] consists of assigning the pixels of the images in just four categories. But, we follow the more complex setup described in [17], and we study a more fine-grained 40 class discrimination task where the scene is divided in *structure categories* like walls, floors, ceiling, windows or doors; *furniture items* like beds, chairs, tables or sofa; and *objects* like lamps, bags, towels or boxes. The complete list is given in Table II.

We measure the performance of our model in the segmentation task using the pixel accuracy (pixels correctly predicted divided by the total number of pixels annotated), the mean class accuracy (mean of the pixel accuracy when computed for each class separately) and the Jaccard index (true predictions divided by the union of predictions and true labels), which is also known as the intersection over union (IoU).

For the scene categorization task, we follow the experimental setup detailed in [17], where the original 27 categories are reorganized into 10 classes: the 9 most common categories plus the category *other* for the images in the remaining classes. For evaluation, we report the classification accuracy, which is the precision over all test images.

Finally, for the object detection task, we again follow the setup proposed in [17], where 17 of 40 object categories are used. They include items such as toilet, bed, sofa, chair, etc. The ground truth BBs for the objects are obtained from the ground truth semantic segmentations, discarding the ignored pixels (pixels with label 255). We follow the standard PASCAL VOC [25] metric of average precision (AP) for measuring the detection performance. That is, for each object class c, we first sort the BBs by their score. Then, we define an overlap criterion of 0.5 between a predicted BB and the ground truth to consider a detection as true positive, and we calculate the precision/recall curve from the ranking obtained. The AP, for each class c, is defined as the mean

precision at eleven equally spaced recall levels:

$$AP = \frac{1}{11} \sum_{r \in (0,0.1,\dots,1)} p(r),$$
(9)

where p(r) is the precision at recall r, but interpolated by taking the maximum precision measured for a method for which the corresponding recall exceeds r.

# B. Scene Labeling

We here compare our model with other state-of-the-art methods in the task of scene labeling. Several approaches have reported semantic segmentation results on this dataset recently, and some of which make use of depth information to enhance their results. Table I shows that our model, simply from RGB images, and using the ResNet-101 as a base network with the aforementioned multi-scale input procedure for the DeepLab architecture, beats them all in all the three evaluation metrics. A set of qualitative results are shown in Figure 2.

Finally, in Table II, we offer detailed IoU results for each of the 40 classes. Our approach outperforms all other methods that have reported these detailed results, obtaining a particularly considerable improvement in the most difficult classes. Therefore, we conclude that the solution proposed for semantic segmentation, defines the new state-of-the-art results for the NYUD2 dataset.

TABLE I Scene labeling performance for the NYUD2 dataset 40-class classification task

Method	Input	pix acc	class acc	IoU
RCNN [19]	RGB-HHA	60.3	35.1	28.6
FCN-16 [21]	RGB-HHA	65.4	46.1	34.0
Eigen et al. [26]	RGB-D-N	65.6	45.1	34.1
FuseNet-SF3 [27]	RGB-D	66.4	44.2	34.0
Context-CRF [28]	RGB	67.4	49.6	37.1
MVCNet [2]	RGB-D	69.1	50.1	38.0
Ours	RGB	70.9	52.2	41.8

# C. Scene Classification

Table III shows the results of our methods when we compare them with other state-of-the-art models for scene recognition in the NYUD2 dataset. We incorporate to the evaluation, two baselines methods, which are a VGG-16 [12] and a ResNet-101 [8] architecture, which we train to directly classify the scenes. Technically, we develop a fine tuning procedure for two pretrained models of these networks, where the objective now is correctly classify the scenes in the NYUD2 dataset.

Using the proposed 40-classes scene labeling histograms as features to feed a linear SVM, already improves all previously reported methods, and stays very close to the ResNet-101 baseline model, specifically trained for the problem of



Fig. 2. Qualitative results fro the scene labeling problem. For each group of three images, we show, from left to right: the RGB image, the ground truth and our results. We represent the best (first 3 rows) and worst (last 2 rows) test images in terms of pixel accuracy.

 TABLE II

 IOU FOR EACH OF THE CLASSES IN THE 40-CLASS SEGMENTATION TASK.

	wall	floor	cabinet	bed	chair	sofa	table	door	window	book sh.	picture	counter	blinds	desk	shelves	curtain	dresser	pillow	mirror	floor mat
SC [9]	60.7	77.8	33.0	40.3	32.4	25.3	21.0	5.9	29.7	22.7	35.7	33.1	40.6	4.7	3.3	27.4	13.3	18.9	4.4	7.1
KDES [29]	60.0	74.4	37.1	42.2	32.5	28.2	16.6	12.9	27.7	17.3	32.4	38.6	26.5	10.1	6.1	27.6	7.0	19.7	17.9	20.1
SVM+Scene [17]	68	81	48	55	40	44	30	8.3	33	20	40	47	44	10	5.1	34	22	28	19	22
RCNN-Depth [19]	68.0	81.3	44.9	65.0	47.9	47.9	29.9	20.3	32.6	18.1	40.3	51.3	42.0	11.3	3.5	29.1	34.8	34.4	16.4	28.0
Det+Scene [18]	67.9	81.5	45.0	60.1	41.3	47.6	29.5	12.9	34.8	18.1	40.7	51.7	41.2	6.7	5.2	26.9	25.0	32.8	21.2	30.7
Ours	74.0	80.6	55.9	64.6	55.0	61.0	40.9	33.1	45.0	39.9	54.5	53.0	57.1	14.5	10.9	46.6	40.2	37.0	34.0	28.4
	clothes	ceiling	books	fridge	tv	paper	towel	shower cur	box	wh. board	person	night stand	toilet	sink	lamp	bath tub	bag	other str.	other fur.	other prop
SC [9]	clothes 6.5	ceiling 73.2	books 5.5	fridge	tv 5.7	paper 12.7	<b>towel</b> 0.1	shower cur 3.6	<b>box</b> 0.1	wh. board	person 6.6	night stand 6.3	toilet 26.7	sink 25.1	lamp 15.9	bath tub 0	bag 0	other str. 6.4	other fur. 3.8	other prop 22.4
SC [9] KDES [29]	clothes 6.5 9.5	ceiling 73.2 53.9	books 5.5 14.8	fridge 1.4 1.9	tv 5.7 18.6	paper 12.7 11.7	towel 0.1 12.6	shower cur 3.6 5.4	box 0.1 3.3	<b>wh. board</b> 0 0.2	person 6.6 13.6	night stand 6.3 9.2	toilet 26.7 35.2	sink 25.1 28.9	lamp 15.9 14.2	<b>bath tub</b> 0 7.8	bag 0 1.2	other str. 6.4 5.7	other fur. 3.8 5.5	other prop 22.4 9.7
SC [9] KDES [29] SVM+Scene [17]	clothes 6.5 9.5 6.9	ceiling 73.2 53.9 59	books 5.5 14.8 4.4	fridge 1.4 1.9 15	tv 5.7 18.6 9.3	paper 12.7 11.7 1.9	towel 0.1 12.6 14	shower cur 3.6 5.4 18	box 0.1 3.3 4.8	wh. board 0 0.2 37	person 6.6 13.6 16	night stand 6.3 9.2 20	toilet 26.7 35.2 50	sink 25.1 28.9 26	lamp 15.9 14.2 6.8	bath tub 0 7.8 33	bag 0 1.2 0.65	other str. 6.4 5.7 6.9	other fur. 3.8 5.5 2	other prop           22.4           9.7           22
SC [9] KDES [29] SVM+Scene [17] RCNN-Depth [19]	clothes 6.5 9.5 6.9 4.7	ceiling 73.2 53.9 59 60.5	books 5.5 14.8 4.4 6.4	fridge 1.4 1.9 15 14.5	tv 5.7 18.6 9.3 31.0	paper 12.7 11.7 1.9 14.3	towel 0.1 12.6 14 16.3	shower cur 3.6 5.4 18 4.2	box 0.1 3.3 4.8 2.1	<b>wh. board</b> 0 0.2 <b>37</b> 14.2	person 6.6 13.6 16 0.2	night stand 6.3 9.2 20 27.2	toilet 26.7 35.2 50 55.1	sink 25.1 28.9 26 37.5	lamp 15.9 14.2 6.8 34.8	bath tub 0 7.8 33 38.2	bag 0 1.2 0.65 0.2	other str. 6.4 5.7 6.9 7.1	other fur.           3.8           5.5           2           6.1	other prop           22.4           9.7           22           23.1
SC [9] KDES [29] SVM+Scene [17] RCNN-Depth [19] Det+Scene [18]	clothes 6.5 9.5 6.9 4.7 7.7	ceiling 73.2 53.9 59 60.5 61.2	books 5.5 14.8 4.4 6.4 7.5	fridge 1.4 1.9 15 14.5 11.8	tv 5.7 18.6 9.3 31.0 15.8	paper           12.7           11.7           1.9           14.3           14.7	towel 0.1 12.6 14 16.3 20.0	shower cur           3.6           5.4           18           4.2           4.2	box           0.1           3.3           4.8           2.1           1.1	wh. board           0           0.2           37           14.2           10.9	person 6.6 13.6 16 0.2 1.4	night stand 6.3 9.2 20 27.2 17.9	toilet 26.7 35.2 50 55.1 48.1	sink 25.1 28.9 26 37.5 45.1	lamp 15.9 14.2 6.8 34.8 31.1	bath tub 0 7.8 33 38.2 19.1	bag 0 1.2 0.65 0.2 0.0	other str. 6.4 5.7 6.9 7.1 7.6	other fur.           3.8           5.5           2           6.1           3.8	other prop           22.4           9.7           22           23.1           22.6

scene classification. If instead we use one-hot vectors as described in Section III-B, we increase the accuracy by 1,3%. Using a Jensen-Shannon kernel for our SVM, with normalized histograms, obtains another 0,3%. Finally, applying a two-level spatial pyramid pooling technique improves the performance by an extra 0,76%, attaining a 68,96% classification accuracy, the new state-of-the-art for this dataset.

# D. Object Detection

Here, we are interested in investigating the task of detecting furniture like objects in indoor scenes. The NYUD2 dataset [9] was originally proposed to study semantic segmentation and scene recognition. However, we can derive BBs annotations for the objects directly from the semantic segmentation ground truth, following the experimental setup of [17]. Hence, we can also use this dataset to evaluate the performance of the models in an object localization task.

Table IV summarizes the results achieved by our simple object localization solution and the proposed approaches described by Gupta *et al.* in [17], [18], [19]. We observe that we are able to consistently outperform all these previous works. Note that our simple approach, using only RGB im-

#### TABLE III

Scene classification accuracy for NYUD2. Comparison with the state of the art.

	Method	Input	Acc.
State of the art	RCNN-Depth [19]	RGB-D	45.4
	FV+CNN [11]	RGB-D	63.9
	D-CNN [13]	RGB-D	65.8
Baselines	VGG-16 classification	RGB	64.37
	ResNet-101 classification	RGB	67.73
	Histograms + Linear SVM	RGB	66.51
Ours	One-hot + Linear SVM	RGB	67.89
	Histograms + Additive Kernel SVM	RGB	68.20
	Histograms + Spatial Pyramid + Additive Kernel SVM	RGB	68.96

#### TABLE IV

OBJECT DETECTION PERFORMANCE, USING THE AP METRIC, FOR THE NYUD2 DATASET. COMPARISON WITH THE STATE OF THE ART.

Input	SVM+Scene [17] RGBD	Det+Scene [18] RGBD	RCNN-Depth [19] RGBD	Ours RGB
Bed	52.1	56.0	66.5	58.7
Chair	6.4	23.5	40.8	43.9
Sofa	17.5	34.2	42.8	52.7
Counter	32.7	24.0	37.6	51.3
Lamp	1.4	26.7	29.3	27.7
Pillow	3.3	20.7	37.4	19.2
Sink	14.0	22.8	24.2	46.1
Table	9.3	17.2	24.3	23.4
Bathtub	28.4	19.3	22.9	33.7
Television	3.1	19.5	37.2	42.7
Bookshelf	6.7	17.5	21.8	35.1
Toilet	13.3	41.5	53.0	74.0
Box	0.7	0.6	3.0	2.5
Desk	0.8	6.2	10.2	7.0
Door	5.0	9.5	20.5	23.3
Dresser	13.3	16.4	26.2	26.5
Night-stand	-	32.6	39.5	43.5
Mean over 16 classes	13	24.3	31.1	35.5
Mean over 17 classes	-	23.0	31.6	36.0

ages, outperforms all these methods which jointly exploit the depth and appearance information. Importantly, we achieve an improvement of more than 5% for the AP compared to the best results detailed in [19]. We provide some qualitative detection results in Figure 3.

#### V. CONCLUSION

This paper proposes a highly effective *bottom-up* approach to perform scene understanding. From an accurate multi-scale semantic segmentation deep network, using the DeepLab model with ResNets, we present and develop a variety of simple but robust strategies to use scene labeling-based descriptors for performing both scene classification and object detection tasks.

We compare our solutions with other state-of-the-art methods using a variety of metrics in the challenging NYUD2 dataset. The obtained results confirm that our methods are able to outperform *all* other previously reported models in all the three tasks comprising scene understanding, by only using RGB images.

Overall, in this work we demonstrate that to rely on a precise scene labeling, as the one given by our model, for acquiring a complete scene understanding, is an extremely effective solution. Therefore, *in pixels we trust*.

#### REFERENCES

- Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, "MFNet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes," in *IROS*, 2017.
- [2] L. Ma, J. Stückler, C. Kerl, and D. Cremers, "Multi-view deep learning for consistent semantic mapping with RGB-D cameras," in *IROS*, 2017.
- [3] J. M. Wong, V. Kee, T. Le, S. Wagner, G. L. Mariottini, A. Schneider, L. Hamilton, R. Chipalkatty, M. Hebert, D. M. S. Johnson, J. Wu, B. Zhou, and A. Torralba, "SegICP: Integrated deep semantic segmentation and pose estimation," in *IROS*, 2017.
- [4] L.-C. 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," arXiv:1606.00915, 2016.
- [5] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in CVPR, 2017.
- [6] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "SUN database: Large-scale scene recognition from abbey to zoo," in *CVPR*, 2010.
- [7] P. Espinace, T. Kollar, A. Soto, and N. Roy, "Indoor scene recognition through object detection," in *ICRA*, 2010.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [9] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from RGBD images," in ECCV, 2012.
- [10] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in CVPR, 2017.
- [11] A. Wang, J. Cai, J. Liu, and T. Cham, "Modality and component aware feature fusion for RGB-D scene classification," in CVPR, 2016.
- [12] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations*, 2015.
- [13] X. Song, L. Herranz, and S. Jiang, "Depth CNNs for RGB-D scene recognition: Learning from scratch better than transferring from RGB-CNNs," in AAAI, 2017.
- [14] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *ECCV*, 2016.
- [15] L.-J. Li, R. Socher, and F.-F. Li, "Towards total scene understanding: Classification, annotation and segmentation in an automatic framework," in *CVPR*, 2009.
- [16] J. Yao, S. Fidler, and R. Urtasun, "Describing the scene as a whole: Joint object detection, scene classification and semantic segmentation," in CVPR, 2012.
- [17] S. Gupta, P. Arbelaez, and J. Malik, "Perceptual organization and recognition of indoor scenes from RGB-D images," in CVPR, 2013.
- [18] S. Gupta, P. Arbeláez, R. Girshick, and J. Malik, "Indoor scene understanding with rgb-d images: Bottom-up segmentation, object detection and semantic segmentation," *IJCV*, 2015.
- [19] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik, "Learning rich features from RGB-D images for object detection and segmentation," in *ECCV*, 2014.
- [20] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in CVPR, 2014.
- [21] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *CVPR*, 2015.
- [22] A. Vedaldi and A. Zisserman, "Efficient additive kernels via explicit feature maps." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 3, pp. 480–492, 2012.
- [23] 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.
- [24] 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 *ECCV*, 2014.
- [25] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results," http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html.
- [26] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," *ICCV*, 2014.



Fig. 3. Qualitative Results for the Object Detection Task. The ground truth bounding boxes are shown in green and our predicted bounding boxes in blue.

- [27] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "FuseNet: Incorporating depth into semantic segmentation via fusion-based CNN architecture," in Asian Conference on Computer Vision (ACCV), November 2016.
- [28] G. Lin, C. Shen, A. van den Hengel, and I. D. Reid, "Exploring context with deep structured models for semantic segmentation," *TPAMI*, vol. 40, no. 6, pp. 1352–1366, 2018. [Online]. Available: http://arxiv.org/abs/1603.03183
- [29] X. Ren, L. Bo, and D. Fox, "RGB-D scene labeling: Features and algorithms," in CVPR, 2012.