

DeNet: Scalable Real-time Object Detection with Directed Sparse Sampling

Lachlan Tychsen-Smith, Lars Petersson
CSIRO (Data61)
7 London Circuit, Canberra, ACT, 2601

Lachlan.Tychsen-Smith@data61.csiro.au, Lars.Petersson@data61.csiro.au

Abstract

We define the object detection from imagery problem as estimating a very large but extremely sparse bounding box dependent probability distribution. Subsequently we identify a sparse distribution estimation scheme, Directed Sparse Sampling, and employ it in a single end-to-end CNN based detection model. This methodology extends and formalizes previous state-of-the-art detection models with an additional emphasis on high evaluation rates and reduced manual engineering. We introduce two novelties, a corner based region-of-interest estimator and a deconvolution based CNN model. The resulting model is scene adaptive, does not require manually defined reference bounding boxes and produces highly competitive results on MSCOCO, Pascal VOC 2007 and Pascal VOC 2012 with real-time evaluation rates. Further analysis suggests our model performs particularly well when finegrained object localization is desirable. We argue that this advantage stems from the significantly larger set of available regions-of-interest relative to other methods. Source-code is available from: <https://github.com/lachlants/denet>

1. Introduction

Feed-forward neural networks exhibit good convergence properties given a random initialization under stochastic gradient descent (SGD) and, given an appropriate network design and training regime, can generalize well to previously unseen data [8]. In particular, convolutional neural networks (CNNs) built from interleaved convolution and pooling layers with ReLU activation functions have set numerous benchmarks in computer vision tasks [8] [6] [20]. A number of methodologies have been developed to map their state-of-the-art dense regression and classification capabilities to the problem of identifying axis aligned bounding boxes of object instances in images. In particular we highlight the relatively slow region based CNN approaches (R-CNN [4], Faster R-CNN [15]) and the more recent work on real-time detection (YOLO [14], SSD [12]).

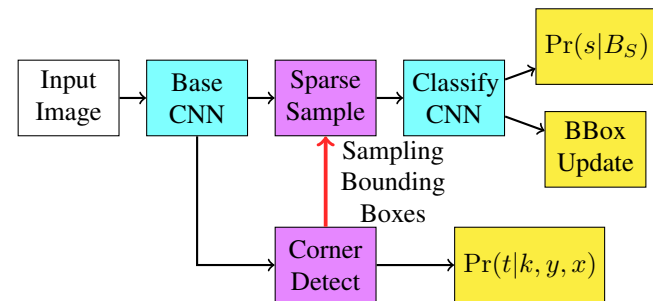


Figure 1. A high level flow diagram depicting the DeNet methodology. The CNN's are highlighted in blue, the novel components in purple and the outputs in yellow. The sampling bounding box dependency B_S (highlighted in red) is held constant during back propagation to produce an end-to-end trained model. The corner distribution and final classification distribution are jointly optimized using cross entropy loss.

Rather than focusing on obtaining state-of-the-art accuracy in a competition environment (i.e. computationally unconstrained) in this paper we emphasis the dual task of obtaining the best detection performance at a predefined evaluation rate i.e. 60 Hz and 30 Hz. The primary contributions made within this paper include:

- An improved theoretical understanding of modern detection methods and a generic framework in which to describe them i.e. Directed Sparse Sampling.
- A novel, fast, region-of-interest estimator which doesn't require manually defined reference bounding boxes.
- A novel application of deconvolution layers which greatly improves evaluation rates.
- Six implementations of our method demonstrating competitive detection performance on a range of benchmarks.
- An easily extended Theano based code release to facilitate the research community.

1.1. Related work

In region based CNN detection (R-CNN) [4] the image is first preprocessed with a region proposal algorithm e.g. selective search [21], region proposal network (RPN) [15], etc. This algorithm identifies image regions (i.e. bounding boxes) of *interest* (RoIs) which are then rescaled to fixed dimensions (normalizing scale and aspect ratio) and fed into a CNN based classifier. The CNN assigns a probability that the region bounds an object of interest or the *null* class and, via linear regression, identifies an improved bounding box. This approach has demonstrated state-of-the-art results, however, it is very expensive to train and evaluate, requiring multiple full CNN evaluations (one per region proposal) and an often expensive pre-processing step. Since the majority of CNN computation occurs in the first few layers, Fast R-CNN [3] addressed these issues by applying a shallow CNN to the image and then, for each region, extracting fixed sized features from the generated feature map for the final classification. In Faster R-CNN [15] the region proposal algorithm was integrated into the CNN providing an end-to-end solution, improved timings and demonstrating that both tasks (region proposal and classification) shared similar underlying features. Despite these improvements, to our knowledge, region based CNN's have not been demonstrated operating near real-time frequencies.

In You Only Look Once (YOLO) [14] they depart from the algorithmically defined *region* based approaches described above, opting instead for a predefined, regular grid of detectors. In effect they merged the region classification problem into the region proposal network (RPN) first proposed in Faster R-CNN. With this approach the CNN is only evaluated once to produce the outcomes for all detectors resulting in significantly reduced training and evaluation times. In Single Shot Detector (SSD) [12] this approach was further refined with an improved network design and training methodology to demonstrate comparable results to the region based methods. We note that the considerable improvements achieved with SSD required scene dependent engineering to manually predefine the most likely set of regions within the image to contain an object, a flaw shared with the Faster R-CNN region proposal network. In particular, SSD demonstrated an improvement of 2.7% MAP [12] by the addition of four aspect ratios to the predefined regions on the Pascal VOC2007 [1] dataset, highlighting the importance of manual engineering in modern state-of-the-art detector designs. Without going into too much detail, we note that in practice manually engineered solutions typically limit scalability and adaptiveness to different problem sets (without an expensive re-engineering process).

The primary differentiator between these methods lies in how each method identifies and treats the *regions* to be classified. R-CNN based methods sample regions sparsely based on an algorithmic preprocessing step and normalize

the region of interest while YOLO based approaches perform dense sampling with a manually defined grid of detectors without image normalization. Often dense methods are well suited to current implementations and, therefore offer a significant timing advantage over sparse methods. However, in this work, we demonstrate a novel model design which combines the ease of training, scene adaptability and classification accuracy of the sparse region-based approaches with the fast training and evaluation of the dense non-region based methods.

1.2. Probabilistic Object Detection

We formulate the probabilistic multiclass detection problem as first estimating the distribution $\Pr(s|B, I)$ where $s \in C \cup \{\text{null}\}$ is a random variable indicating the presence of an instance of class $c \in C$ or the *null* class (indicating no instances) which is *sufficiently* bounded by the box $B = \{x, y, w, h\}$ and I is the input image (omitted in subsequent derivations). This formulation incorporates the assumption that only a single instance of a class can occupy each bounding box. We note that this definition does not seek to perform instance assignment, but can be used as an input to an algorithm that does e.g. Non-Max Suppression.

Given a suitable neural network design we assert that $\Pr(s|B)$ can be estimated from training data with class bounding box annotations. However, since the number of unique bounding boxes is given by $|B| \propto XYWH$ where (X, Y) are the number of image positions and (W, H) the range of bounding box dimensions the naive solution quickly becomes intractable. For instance, assuming the most common settings for the ImageNet dataset, 1000 classes and 224×224 images, and considering all valid bounding boxes within the image, expressing this distribution requires approximately 629×10^9 values or 2.5TB in 32bit float format. Clearly this is an intractable problem with current hardware.

At the cost of localization accuracy, subsampling the output bounding boxes is a valid approach. For instance, by careful dataset dependent manual engineering, Faster R-CNN and YOLO based approaches subsample the distribution to the order of 10^4 to 10^5 bounding boxes [14] [15]. These boxes are then refined by estimating only the most likely bounding box in a local region via linear regression. As an alternative to large scale subsampling, we sought to exploit the fact that, due to occlusion and other factors, we expect a very small subset of bounding boxes to contain class instances other than the *null* class. Subsequently, we have developed a solution based on the state-of-the-art regression capabilities of a single end-to-end CNN which estimates the highly sparse distribution $\Pr(s|B)$ in a real-time (or computationally constrained) operational environment.

2. Directed Sparse Sampling (DSS)

We use the term *Directed Sparse Sampling* to refer to the method of applying a jointly optimized two stage CNN where one stage estimates the likely locations where user-defined *interesting* values occur and the other sparsely classifies the identified values e.g. in R-CNN based models (including R-FCN and DeNet) we estimate the bounding boxes which are most likely to include a non-*null* class assignment, then run a classifier over these bounding boxes.

2.1. Corner-based RoI Detector

Here we introduce the concept of bounding box corner estimation for efficient region-of-interest (RoI) estimation. In our methodology, this task is performed by estimating the likelihood that each position in the image contains an instance of one of 4 corner types i.e. $\Pr(t|k, y, x)$ where t is a binary variable indicating the presence of a corner of type $k \in \{\text{top left, top right, bottom left, bottom right}\}$ at position (x, y) in the input image. We assert that due to the natural translation invariance of the problem, estimating the corner distribution can be efficiently performed with a standard CNN design trained on bounding box annotated image data (e.g. MSCOCO [11], Pascal VOC [1], etc).

With the corner distribution defined we estimate the likelihood that a bounding box B contains an instance by applying a Naive Bayesian Classifier to each corner of the bounding box:

$$\Pr(s \neq \text{null}|B) \propto \prod_k \Pr(t|k, y_k, x_k) \quad (1)$$

where $(x_k, y_k) = f_k(B)$ indicates the bounding box position associated with each corner type k . For ease of implementation we define the $N \times N$ bounding boxes with the largest non-null probability $\Pr(s \neq \text{null}|B)$ as the *sampling bounding boxes* B_S . The user defined variable N balances the maximum number of detections the model can handle with the computational and memory requirements.

With the potentially non-null bounding boxes estimated, we pass a feature vector of predefined length from the corner detector model to the final classification stages. Therefore, the final classification stage is a function of the form $f : \bar{\alpha}_B \rightarrow \Pr(s|B)$ where $\bar{\alpha}_B$ is a feature vector uniquely identified by the sampling bounding box $B \in B_S$. It is important that the feature is uniquely associated with each bounding box, otherwise the classifier will have no information to distinguish between bounding boxes with the same $\bar{\alpha}_B$. Exactly how to construct the feature vector is still a matter of debate [9, 15] however we construct $\bar{\alpha}_B$ by concatenating together the *nearest neighbour sampling* features at predefined locations relative to each sampling bounding box (e.g. bounding box corners, center, etc) in addition to

the bounding box width and height. The bounding box center position was omitted from the feature vector such that the classifier would be agnostic to image offsets.

2.2. Training

During training, the model is initially forward propagated to generate the sampling bounding boxes B_S as described in the previous subsection. In addition, we augment the sampling bounding boxes with the ground truth bounding boxes and randomly generated samples. We then propagate the activations $\bar{\alpha}_B$ associated with the augmented set of sampling bounding boxes through the rest of the model to produce the final classification distribution $\Pr(s|B_S)$ and updated bounding box parameters. The set of sampling bounding boxes B_S is held constant during gradient estimation to enable end-to-end training, therefore the corner detector network is optimized in conjunction with the bounding box classification and estimation task. Since forward propagation is a necessary preprocessing step in the back propagation based SGD policy typically used to optimize neural networks, the DeNet method introduces no penalty to training time over a standard dense network.

The DeNet model jointly optimizes over the corner probability distribution, final classification distribution and bounding box regression cost, i.e.

$$\begin{aligned} \text{Cost} = & \frac{\lambda_t}{\Lambda_t} \sum_{k,y,x} \phi(t|k, y, x) \ln(\Pr(t|k, y, x)) + \\ & \frac{\lambda_s}{\Lambda_s} \sum_{B \in B_S} \phi(s|B) \ln(\Pr(s|B)) + \\ & \frac{\lambda_b}{\Lambda_b} \sum_i \text{SoftL}_1(\phi_{B,i} - \beta_i) \end{aligned} \quad (2)$$

where $\phi(\dots)$ are the ground truth corner and classification distributions, $\phi_{B,i} = \{x_i, y_i, w_i, h_i\}$ the ground truth bounding boxes, $(\lambda_s, \lambda_t, \lambda_b)$ are user defined constants indicating the relative strength of each component, $(\Lambda_s, \Lambda_t, \Lambda_b)$ are constants normalizing each component to 1 given the model initialization and $\text{SoftL}_1(x)$ is defined in [3]. The corner distribution $\phi(t|k, y, x)$ is identified by mapping each groundtruth instance's corners to a single position in the corner map, corners out of bounds are simply discarded. The detection distribution $\phi(s|B)$ is identified by calculating the *intersection over union* (IoU) overlap between the groundtruth bounding boxes and the sampling bounding boxes B_S . Following standard practice, the regression target bounding box ϕ_B is identified by selecting the ground truth bounding box with the largest IoU overlap.

2.3. Detection Model

Residual neural networks [6] have demonstrated impressive regression capabilities on a number of large scale datasets. In particular the 101 layer Residual Network model (ResNet-101) achieved state of the art performance on the ILSVRC2015 [16] and MSCOCO [11] datasets when combined with Faster R-CNN. As the base model to our networks we selected the 34 layer, 21M parameter ResNet-34 model (DeNet-34) and the 101 layer, 45M parameter ResNet-101 model (DeNet-101).

To each base model we modified the input size to 512×512 pixels, removed the final mean pooling and fully connected layers and appended two deconvolution [13] layers followed by a corner detector. The corner detector is responsible for generating the corner distribution and produces a *feature sampling map* via a learnt linear projection with F_s features at each spatial position. The deconvolution [13] layers efficiently reintroduce spatial information that was lost in the base model such that the feature map and corner probability distribution can be defined at a greater spatial resolution i.e. 64×64 compared to 16×16 without. This results in a 16×16 pixel minimum size for each sampling bounding box.

Following the corner detector is the *sparse layer* which observes the corners identified by the corner detector and generates a set of sampling bounding boxes (RoIs). The RoIs are used to extract a set of $N \times N$ feature vectors from the *feature sampling maps*. In this case, we are sparsely sampling N^2 bounding boxes from a set of 4.2M valid bounding boxes. A feature vector is constructed by extracting the *nearest neighbour* sampling features associated with a 7×7 grid plus the bounding box width and height. This produces a feature with $7 \times 7 \times F_s + 2$ values. We found that nearest neighbour sampling was sufficient because the feature sampling maps have the same, relatively high, spatial resolution as the bounding box corners. Finally, the feature vectors are propagated through a relatively shallow fully connected network to generate the final classification and fine tuned bounding box for each sampling RoI.

In Table 1 and 2 we describe the additional layers appended to the base models with the following definitions:

- **Conv:** Convolves a series of 2D filters over the input activations. Filter weights were initialized via the normal distribution $\mathcal{N}(0, \sigma)$ with $\sigma^2 = 2/(n_f n_x n_y)$ where n_f is the number of filters and (n_x, n_y) their spatial shape [5]. Following each convolution is batch normalization [7] then the ReLU activation function.
- **Deconv:** Applies a learnt deconvolution [13] (upsampling) operation followed by ReLU activation. In this case it is equivalent to upscaling both spatial dimensions then applying a **Conv** layer.

Model	F_0	F_1	F_2	F_3	F_4	F_5	F_6	F_7
DeNet-34	512	256	128	4706	1536	1024	768	512
DeNet-101	2048	384	192	6274	2048	1536	1024	768

Table 1. Filter parameters used for DeNet models. See Table 2

Layer	Input Shape	Filters	Shape	Stride
ResNet-34 or ResNet-101 [6] base model.				
Deconv	$16 \times 16 \times F_0$	F_1	3×3	2×2
Deconv	$32 \times 32 \times F_1$	F_2	3×3	2×2
Corner	$64 \times 64 \times F_2$	-	-	-
Sparse	-	-	-	-
Conv	$N \times N \times F_3$	F_4	1×1	1×1
Conv	$N \times N \times F_4$	F_5	1×1	1×1
Conv	$N \times N \times F_5$	F_6	1×1	1×1
Conv	$N \times N \times F_6$	F_7	1×1	1×1
Classifier	$N \times N \times F_7$	-	-	-

Table 2. DeNet: A ResNet derived model for DSS Object Detection with a 512×512 input image. Layers in the base models above the line are initialized with a pretrained ResNet-34 or ResNet-101 ImageNet 2012 classification model.

- **Corner:** Estimates a corner distribution via the softmax function and produces a sampling feature map. See Section 2.
- **Sparse:** Identifies sampling bounding boxes from corner distribution and produces a fixed size sampling feature from the sampling feature maps.
- **Classifier:** Maps activations to the desired probability distribution via the softmax function and generates bounding box targets.

For DeNet-34 we use a ResNet-34 base model and $F_s = 96$ to produce a feature vector of 4706 values and a total of 32M parameters. The DeNet-101 model uses a ResNet-101 base model and increased the number of filters by approximately $1.5 \times$ for the appended layers (See Table 1). These changes produce a sparse feature vector of 6274 values and a total of 69M parameters.

2.3.1 Skip Layer Variant

As an extension we considered augmenting the DeNet models with *skip* layers. In recent work, *skip* layers have demonstrated consistent improvements in classification [6], detection [10] and semantic segmentation [2] and, more generally, are an integral component to highway [18] and residual networks [6]. In this case, these layers connect the **Deconv** layer with the final layer in the base model which has the same spatial dimensions. Our implementation follows [10], each skip layer performs a linear projection of the source features to the destination feature dimensions and simply adds the resulting feature maps (before activation).

2.3.2 Wide Variant

In this model we modified the skip model variants to use a 128×128 spatial resolution for the corner and feature sampling maps by the addition of another **Deconv** and skip layer. We also increased N to 48 to produce 2304 RoIs. In the current implementation, this approach comes with a considerable timing cost due to the increased classification burden and the CPU bound algorithm for identifying RoIs. With further engineering (e.g. deduplication) we believe these costs could be reduced.

3. Implementation Details

Our models are implemented within our Theano based CNN library called DeNet. The source-code is available from: <https://github.com/lachlants/denet>

3.1. Training methodology

In all experiments we used Nesterov style SGD [19] with an initial learning rate of 0.1, momentum of 0.9 and weight decay of 0.0001 (only applied to weights). A batch size of 128 was employed for both models with 32 samples per GPU iteration. The learning rate was divided by 10 at epoch 30 and epoch 60 and a total of 90 training epochs were performed. Note that, apart from the batch size changes, these hyperparameters are identical to those used when training the original residual networks for classification [6]. No on-line hard negative mining [17] or other gradient optimization techniques were applied, however, we observed some instances of overtraining on Pascal VOC. In response, to increase exposure to negative samples, we introduced 10% randomly generated bounding box samples during training.

An augmentation strategy very similar to GoogLeNet [20] was employed to improve model generalization to different scales and translations. For each sample, a black border was added to the smallest dimension to produce a square image. At test time, this image was scaled to 512×512 pixels using bilinear sampling, during training a random crop was selected with an area between (0.08, 1.0) relative to the border image and an aspect ratio between (3/4, 4/3). The random crop was discarded and a new one generated if no ground truth objects overlapped with the crop by at least 50%. This process was repeated up to 10 times and, as a fallback, the entire bordered image was returned. As in testing, the resulting crop was scaled to 512×512 pixels. Random photometric (contrast, saturation and brightness) and mirror augmentation was also employed [20].

3.2. Identifying Sampling Bounding Boxes (RoIs)

A simple algorithm was developed to quickly search the corner distribution for non-null bounding boxes:

1. Search the corner distribution for corners $\{k, y, x\} \in C_\lambda$ where $\Pr(t = 1 | k, y, x) > \lambda$.

2. For each corner type, select the M corners with the greatest likelihood $C_M \subseteq C_\lambda$
3. Generate a set of unique bounding boxes by matching every corner within C_M of type top – left with every one of type bottom – right.
4. Calculate the probability of each bounding box being non-null via Equation 1.
5. Repeat steps 2 and 3 with corners of type top – right and bottom – left.
6. Sort bounding boxes by probability and keep the N^2 largest to produce the sampling bounding boxes B_S .

Since the vast majority of corners are culled in step 1 this method obtains a significant speed up beyond the naive brute force method i.e. testing every possible bounding box.

4. Results and Analysis

In this section we compare our design with previously published models. We note that in some cases, an *apples-to-apples* comparison is difficult due to the wide range of base models, data augmentation schemes and dataset merging. In particular, we note that SSD utilize larger batch sizes while R-CNN models have larger input resolutions (on average). All our DeNet timing results are provided for a single Titan X GPU (CuDNN 5110) with a batch size of 8x, the same settings used in SSD. For brevity we include only three flavours of the non real-time Faster R-CNN model, the original RPN (VGG), the ResNet-101 extension RPN+ (ResNet-101) and R-FCN for comparison (highlighted in grey in the tables). We note that due to implementation restrictions RPN based models are tested with a single image per batch.

Model	Max. Input	BS	L	Param.
RPN (VGG) [15]	1000×600	1	16	137M
RPN+ (ResNet) [6]	1000×600	1	100	45M
R-FCN [9]	1000×600	1	100	45M
Fast YOLO [14]	448×448	1	9	9M
YOLO [14]	448×448	1	26	60M
SSD300 [12]	300×300	8	25	27M
SSD512 [12]	512×512	8	25	27M
DeNet-34 (ours)	512×512	8	41	33M
DeNet-101 (ours)	512×512	8	107	69M

Table 3. Model overview detailing maximum input image sizes, batch size at test time (BS), number of activation layers (L) and approximate number of parameters.

In Table 3 we provide a broad overview of the baseline models. We note that despite an increased number of layers and parameters the DeNet models obtain improved evaluation rates (See Section 4.2).

λ_s	1	1	1	1	1	1	1
λ_t	1	10	10	100	100	100	500
λ_b	1	1	10	0.1	1	10	1
MAP (%)	49.6	63.3	64.9	71.8	72.8	72.1	70.6

Table 4. Optimizing cost hyperparameters $\{\lambda_s, \lambda_t, \lambda_b\}$, see Equation 2. MAP is provided for Pascal VOC 2007 `val` dataset.

Sample BBoxes	64	144	256	400	576	784	1024
Coverage (%)	81	87	91	93	95	96	97
Eval. Rate (Hz)	96	90	84	76	69	61	55

Table 5. Sample bounding boxes vs coverage over training dataset and evaluation rate.

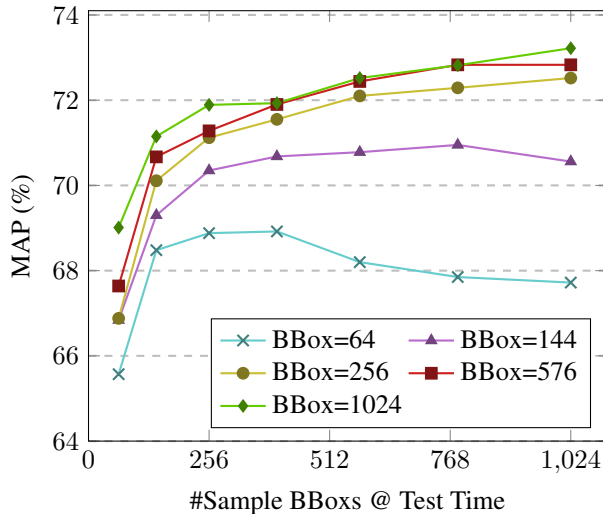


Figure 2. The MAP on the Pascal VOC 2007 validation dataset with varying number of bounding box samples during training (see legend) and testing (displayed on x-axis).

4.1. Hyperparameter Optimization

For the following we used the DeNet-34 model and trained it on Pascal VOC 2007 `train` and Pascal VOC 2012 `trainval` (14,041 images), for testing we used Pascal VOC 2007 `val` (2,510 images). The same DeNet-34 model initialization was used for all experiments. We applied the training procedure described in Section 3.1 except with a batch size of 96.

In Table 4 we performed a coarse search over the corner and bounding box regression cost parameters λ_t and λ_b . Best results were achieved by setting $\lambda_s = 1$, $\lambda_t = 100$ and $\lambda_b = 1$, these were applied in all subsequent experiments. Next, we investigated the model behaviour with varying numbers of sampling bounding boxes. In particular, we trained a set of models with $N = \{8, 12, 16, 24, 32\}$. At test time, we took each of these models and varied N from 8 to 32 to produce Figure 2. In Table 5, we provide the model evaluation rate and *coverage* (percentage of ground

	DeNet-34		DeNet-101	
Estimate corners	8.4 ms	68%	24.9ms	80%
Generate RoI	1.5 ms	12%	1.5 ms	5%
Classify RoI	2.4 ms	19%	4.5 ms	15%
Estimate instances	0.1 ms	1%	0.1 ms	0%
Total (per image)	12.4 ms	-	31.0 ms	-

Table 6. A coarse timing breakdown per image for DeNet models at test time on a Titan X GPU.

truth with a sampling bounding box with $\text{IoU} > 0.5$) over the training set that was obtained by the RoI estimator described in Section 2. As expected, we observed a consistently improving MAP with diminishing returns above 576 when training with a larger number of sampling bounding boxes. In general we observed an improved MAP with increased testing bounding boxes at the cost of evaluation rates. For subsequent experiments we set $N = 24$ for both training and testing.

4.2. Timing Breakdown and Evaluation Rates

In Table 6 we present a coarse analysis of the timing for both DeNet models. We broke the timing into 4 sequentially executed stages:

- Estimate corners:** Images are uploaded to the GPU and fed through the base network generating the corner distribution and sampling feature maps. The corner distribution is transferred from GPU to CPU memory.
- Generate RoI:** The sampling bounding boxes (RoIs) are generated from the corner distribution.
- Classify RoI:** The final classification CNN is executed, the classification distribution and bounding box regression outputs are transferred from GPU to CPU.
- Estimate instances:** Non-Max Supression is run over the resulting detection *hits* producing a de-duplicated list of detections for each image.

We observe that the vast majority of time is spent evaluating the base network to generate corners. Also, note that the CPU bound **Generate RoI** stage timing can vary substantially between different samples and may require additional tuning depending on application. Furthermore, we wish to emphasize a number of important features of the DeNet model which makes it significantly faster than most other baseline models:

- Deconvolution:** Spatial information is increased via deconvolution layers as opposed to the *atrous* modified models used in R-FCN and SSD. This method introduces spatial information significantly later in the model, greatly improving evaluation rates.

- **Fast RoI Features:** Features are extracted via a simple *nearest neighbour* sampling method, limiting the the number of feature reads to 49 per RoI. Some RPN variants use pooling which varies from 49-580 per RoI.
- **Input Image Dimensions:** DeNet scales all images to 512x512 pixels, whereas RPN based methods use a varying input size up to 1000x600 pixels.
- **Batching:** Our models are tested with 8x samples per batch (same as SSD). This improves GPU utilization.

With these improvements in timing we are able to use a more expressive base model for the same evaluation rate.

4.3. RoI Coverage Comparison

In Table 7. we provide the coverage obtained by the top 300 RoIs for RPN, R-FCN and DeNet methods. We observe that given a relatively low number of RoIs, RPN (VGG) and R-FCN provide better coverage at low IoU thresholds, however with increasing IoU the DeNet models provide significantly improved coverage.

Model	Top 300 Coverage@IoU (%)				
	0.5	0.6	0.7	0.8	0.9
RPN (VGG)	89.24	81.99	65.18	27.79	2.66
R-FCN	91.41	86.72	77.37	45.37	5.73
DeNet-34	80.98	77.00	71.63	63.39	46.37
DeNet-101	82.47	78.69	73.81	65.36	47.83

Table 7. Coverage on Pascal VOC 2007 *test* using 300 sample bounding boxes (RoI proposals).

We note that RPN / R-FCN utilize bounding box regression and deduplication methods in their RoI proposal networks, these factors improve coverage with low numbers of proposals. As demonstrated in the following sections, the DeNet RoI coverage results do not necessarily translate to a reduced MAP for the full model, which includes NMS and bounding box regression, at lower IoU thresholds.

4.4. MSCOCO

The Microsoft Common Object in Context [11] dataset consists of 82K training and 40K validation images distributed across 80 classes. For testing, the dataset includes an 80K test dataset from which a user known subset of 20K images forms the *test-dev2015* set and an unknown subset of 20K images forms *test2015* allowing only 5 evaluations. Due to the dataset size, number of classes and relatively small size of the object instances within the images, MSCOCO is a considerably more difficult dataset compared to the Pascal VOC challenges. The primary evaluation metric for MSCOCO is the integral of the MAP over the detection matching parameter IoU=0.5 to IoU=0.95.

This metric places a greater emphasis on localization performance compared to the Pascal datasets. We found that setting $\lambda_t = 50$ for DeNet-101 was necessary for convergence, this is likely due to the greater number of corners present within each image on average compared to the validation experiments. Training took ~ 4 days with $2 \times$ Tesla P100 GPUs for DeNet-34 and ~ 6.5 days with $4 \times$ Tesla P100 GPUs for DeNet-101.

In Table 8 we provide the precision and recall results for our models on *test-dev2015*. The DeNet models demonstrate a clear advantage over other high evaluation rate implementations e.g. our real-time DeNet-34 model beats SSD300 by 6.2% MAP at the same evaluation rate and SSD512 by 2.6% at more than twice the evaluation rate. The DeNet-101 model furthers this advantage and is only beaten by the very slow *competition* style RPN+ model utilizing multi-scale evaluation and bounding box refinement. At the time of writing, the DeNet-101 model obtains a result good enough to be in the top-10 on the MSCOCO competition leaderboard which doesn't consider evaluation time. The *skip* model variants consistently improved performance on small and medium sized objects (see AR and AP for small and medium area objects in table) with a minor cost to large objects and evaluation rate. The *wide* variants further improved small object detection and fine object localization at the cost of evaluation rate. Near identical results were obtained on MSCOCO *test-std2015* e.g. we obtained a MAP@[0.5:0.95] of 29.3% and 31.7% for DeNet-34 and DeNet-101 respectively. Analysis suggests our advantage stems from improved large object detection and *finegrain* object localization performance, as reflected in the MAP@IoU=0.75 result. We argue this is an outcome of the much larger range of candidate RoI's our method produces e.g. the vanilla DeNet models can select from a possible set of 4.2×10^6 bounding boxes while SSD utilizes 2.5×10^4 . Utilizing such a large set of candidate bounding boxes would likely be intractable with the dense evaluation methods used in the YOLO and RPN derived models.

4.5. Pascal VOC 2007

We combined the *trainval* samples from Pascal VOC 2007 and 2012 [1] (denoted 07+12 in table) to produce 16,551 training samples. For testing we used Pascal VOC 2007 *test* containing 4,991 samples. We note that this dataset is considerably smaller than MSCOCO and therefore more susceptible to overtraining and image augmentation methods. Training time was ~ 13 hours for DeNet-34 using $2 \times$ Tesla P100s and ~ 20 hours for DeNet-101 using $4 \times$ Tesla P100s. In Table 9, we provide the MAP and timing results. We observed the skip layer variant DeNet-34 improving upon SSD300's peak MAP by 1.6% and 20Hz. In the near real-time domain DeNet-101 matches SSD512 at a higher evaluation rate.

Model	Eval. Rate	AP@IoU (%)			AP@Area (%)			AR@Dets (%)			AR@Area (%)		
		0.5:0.95	0.5	0.75	S	M	L	1	10	100	S	M	L
RPN (VGG)	7 Hz	21.9	42.7	-	-	-	-	-	-	-	-	-	-
RPN+ (ResNet)	<1 Hz	34.9	55.7	-	15.6	38.7	50.9	-	-	-	-	-	-
R-FCN	9 Hz	29.9	51.9	-	10.8	32.8	45.0	-	-	-	-	-	-
SSD300	58 Hz	23.2	41.2	23.4	5.3	23.2	39.6	22.5	33.2	35.3	9.6	37.6	56.5
SSD512	23 Hz	26.8	46.5	27.8	9.0	28.9	41.9	24.8	37.5	39.8	14.0	43.5	59.0
DeNet-34	83 Hz	29.4	46.2	31.2	7.8	30.8	47.4	26.9	38.0	38.5	11.2	41.9	63.0
DeNet-34 (skip)	82 Hz	29.5	47.9	31.1	8.8	30.9	47.0	26.9	38.0	38.6	13.2	41.7	61.6
DeNet-34 (wide)	44 Hz	30.0	48.9	31.8	10.1	30.9	45.7	27.3	39.5	40.3	17.0	42.8	60.9
DeNet-101	34 Hz	31.9	50.5	34.2	9.7	34.9	50.6	28.4	39.8	40.3	13.1	44.8	64.1
DeNet-101 (skip)	33 Hz	32.3	51.4	34.6	10.5	35.1	50.9	28.5	40.2	40.8	14.7	44.9	63.8
DeNet-101 (wide)	17 Hz	33.8	53.4	36.1	12.3	36.1	50.8	29.6	42.6	43.5	19.3	46.9	64.3

Table 8. MSCOCO average precision (AP) and average recall (AR) results evaluated on test-dev2015 dataset.

Model	Dataset	Eval. Rate	MAP
RPN (VGG)	07+12	7 Hz	73.2%
RPN (ResNet)	07+12	2 Hz	76.4%
R-FCN	07+12	9 Hz	80.5%
Fast YOLO	07+12	155 Hz	52.7%
YOLO	07+12	45 Hz	63.4%
SSD300	07+12	58 Hz	74.3%
SSD512	07+12	23 Hz	76.8%
DeNet-34	07+12	83 Hz	75.3%
DeNet-34 (skip)	07+12	82 Hz	75.9%
DeNet-101	07+12	34 Hz	77.0%
DeNet-101 (skip)	07+12	33 Hz	77.1%

Table 9. Pascal VOC 2007 mean average precision and timing.

4.6. Pascal VOC 2012

In this experiment we combine `trainvaltest` from Pascal VOC 2007 and `trainval` from Pascal VOC 2012 [1] (denoted 07++12 in table) to produce 21,503 training samples. Test scores are evaluated on 10,991 samples by the Pascal VOC 2012 testing server. For this dataset the DeNet-34 model matches SSD300, however, for reasons unknown, DeNet-101 demonstrates results below SSD512. For reference, we note that DeNet-101 obtains results near identical to the other ResNet-101 based model, RPN (ResNet) with an order of magnitude improvement in evaluation rate. Training time was ~ 18 hours for DeNet-34 with $2\times$ Tesla P100s and ~ 28 hours for DeNet-101 with $4\times$ Tesla P100s.

5. Conclusion

In this work, we describe a framework for sparse estimation with CNNs and present a novel region-of-interest detector and classification model which reduces manual engineering and improves state-of-the-art detection performance with real-time and near real-time evaluation rates.

Model	Dataset	Eval. Rate	MAP
RPN (VGG)	07++12	7 Hz	70.4%
RPN (ResNet)	07++12	2 Hz	73.8%
R-FCN	07++12	9 Hz	77.6%
YOLO	07++12	45 Hz	57.9%
SSD300	07++12	58 Hz	72.4%
SSD512	07++12	23 Hz	74.9%
DeNet-34	07++12	83 Hz	72.3%
DeNet-101	07++12	33 Hz	73.9%

Table 10. Pascal VOC 2012 mean average precision and timing.

Utilizing deconvolution and skip layers first described in the context of semantic segmentation, we demonstrated a highly computationally efficient model with tightly coupled RoI, class prediction and bounding box regression. We provide further evidence that skip connections consistently improved detection rates for small and medium sized objects. While the *wide* model variant highlighted the importance of corner map resolution for small and medium sized objects, and provides a natural pathway for future development.

Analysis suggests our model performs particularly well when finer object localization is desirable. We propose that the improved localization is due to the much larger set of possible sampling bounding boxes that are feasible with our sparse sampling method i.e. 4.2×10^6 compared to less than 2.5×10^4 for SSD512 and RPN. This feature allows the model to potentially select a bounding box (before bounding box regression) which is significantly closer to the ground truth. Furthermore, since we no longer define a set of reference bounding boxes, this approach has reduced manual engineering requirements and can adapt well to problems which utilize bounding boxes with a very large range of aspect ratios and scales e.g. rotationally variant or non-rigid objects.

References

- [1] M. Everingham, L. Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vision*, 88(2):303–338, June 2010. [2](#), [3](#), [7](#), [8](#)
- [2] G. Ghiasi and C. C. Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *European Conference on Computer Vision*, pages 519–534. Springer, 2016. [4](#)
- [3] R. Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015. [2](#), [3](#)
- [4] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Region-based convolutional networks for accurate object detection and segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 38(1):142–158, 2016. [1](#), [2](#)
- [5] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [4](#)
- [6] 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. [1](#), [4](#), [5](#)
- [7] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. [4](#)
- [8] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998. [1](#)
- [9] Y. Li, K. He, J. Sun, et al. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in Neural Information Processing Systems*, pages 379–387, 2016. [3](#), [5](#)
- [10] T. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie. Feature pyramid networks for object detection. *CoRR*, abs/1612.03144, 2016. [4](#)
- [11] 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. [3](#), [4](#), [7](#)
- [12] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C. Fu, and A. C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015. [1](#), [2](#), [5](#)
- [13] 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. [4](#)
- [14] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016. [1](#), [2](#), [5](#)
- [15] 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](#), [2](#), [3](#), [5](#)
- [16] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li. Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575, 2014. [4](#)
- [17] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 761–769, 2016. [5](#)
- [18] R. K. Srivastava, K. Greff, and J. Schmidhuber. Training very deep networks. In *Advances in neural information processing systems*, pages 2377–2385, 2015. [4](#)
- [19] I. Sutskever, J. Martens, G. E. Dahl, and G. E. Hinton. On the importance of initialization and momentum in deep learning. *ICML (3)*, 28:1139–1147, 2013. [5](#)
- [20] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. [1](#), [5](#)
- [21] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 104(2):154–171, 2013. [2](#)