# FAST, ACCURATE BARCODE DETECTION IN ULTRA HIGH-RESOLUTION IMAGES

Jerome Quenum, Kehan Wang, Avideh Zakhor

Department of Electrical Engineering and Computer Science University of California, Berkeley {jquenum, wang.kehan, avz}@berkeley.edu

## ABSTRACT

Object detection in Ultra High-Resolution (UHR) images has long been a challenging problem in computer vision due to the varying scales of the targeted objects. When it comes to barcode detection, resizing UHR input images to smaller sizes often leads to the loss of pertinent information, while processing them directly is highly inefficient and computationally expensive. In this paper, we propose using semantic segmentation to achieve a fast and accurate detection of barcodes of various scales in UHR images. Our pipeline involves a modified Region Proposal Network (RPN) on images of size greater than 10k×10k and a newly proposed Y-Net segmentation network, followed by a post-processing workflow for fitting a bounding box around each segmented barcode mask. The end-toend system has a latency of 16 milliseconds, which is  $2.5 \times$  faster than YOLOv4 and  $5.9 \times$  faster than Mask R-CNN. In terms of accuracy, our method outperforms YOLOv4 and Mask R-CNN by a mAP of 5.5% and 47.1% respectively, on a synthetic dataset. We have made available the generated synthetic barcode dataset and its code at http://www.github.com/viplabB/SBD/.

*Index Terms*— Barcode detection with deep neural networks, barcode segmentation, Ultra High-Resolution images.

## 1. INTRODUCTION

Barcodes are digital signs often made of adjacent and alternating black and white smaller rectangles that have become an intrinsic part of human society. In administration, for example, they are used to encode, save, and retrieve various users' information. At grocery stores, they are used to track sales and inventories. More interestingly in e-commerce, they are used to track and speed up processing time in warehouses and fulfillment centers.

In classical signal processing, filters used for detection are image-specific since input images are not all necessarily acquired with the same illumination, brightness, angle, or camera. Consequently, adaptive image processing algorithms are required, which can impact detection accuracy [1]. In addition, because classical signal processing methods often run on Central Processing Units, they tend to be much slower compared with deep learning implementations that are easily optimized on Graphics Processing Units (GPUs).

Over the years, a number of methods have been proposed to detect barcodes using classical signal processing [1, 2, 3, 4, 5], but nearly all of them take too long to process Ultra High-Resolution (UHR) images. More specifically, [5] used parallel segment detectors which improved on their previous work [6] of finding imaginary perpendicular lines in Hough space with maximal stable extremal regions to detect barcodes. Katona et al. [3] used morphological manipulation for barcode detection, but this method did not generalize well as different barcode types have varying detection performances. Similarly, [7] proposed using x and y derivative differences, but varying input images yields different outputs, and using such operation on UHR images often become highly inefficient.

With neural networks, though there has been much improvement in barcode detection tasks, few of them have addressed the fast and accurate detection problem in UHR images. Zamberletti et al. [8] paved the way for using neural networks to detect barcodes by investigating Hough spaces. This was followed by [9] which adapted the You Only Looked Once (YOLO) detector to find barcodes in Low Resolution (LR) images, but the YOLO algorithm is known to perform poorly with long shaped objects such as code 39 barcodes. Instance segmentation methods such as Mask R-CNN [10] perform better on images of size  $1024 \times 1024$  pixels but on smaller size images, the outputted Region of Interests (RoI) do not align well with long, 1D-barcode structures. This is because it typically predicts masks on  $28 \times 28$  pixels irrespective of object size, and thereby generates "wiggly" artifacts on some barcode predictions, losing spatial resolution. In the same way, dedicated object detection pipelines, such as YOLOv4 [11], though they perform well on lower Intersection over Union (IoU) thresholds, suffer accuracy at higher IoU thresholds. Among those using segmentation on LR images as a means for detection, [12] also tends to not perform well at higher IoU thresholds.

In this paper, we propose a pipeline for detecting barcodes using deep neural networks, shown in Fig. 1, which consists of two stages trained separately. When compared with classical signal processing methods, neural networks not only provide a faster inference time, but also yield higher accuracy because they learn meaningful filters for optimal feature extraction. As seen in Fig. 1, in the first stage, we expand on the Region Proposal Network (RPN) introduced in Faster R-CNN [13] to extract high definition regions of potential locations where barcodes might be. This stage allows us to significantly reduce inference computation time that would have been required otherwise in the second stage. In the second stage, we introduce Y-Net, a semantic segmentation network that detects all instances of barcodes in a given outputted RoI image ( $400 \times 400$ ). We then apply morphological operations on the predicted masks to separate and extract the corresponding bounding boxes as shown in Fig. 2.

One of the limitations of existing work on barcode detection is the insufficient number of training examples. ArTe-Lab 1D Medium Barcode Dataset [8] and the WWU Muenster Barcode Database [14] are two examples of existing available datasets. They contain 365 and 595 images respectively, with ground truth masks at a resolution of  $640 \times 480$ . Most of the samples in the ArTe-Lab dataset have only one EAN13 barcode per sample image, and few of them in the Muenster database have more than one barcode instance on a given

This work was supported by Amazon.com, Inc.



Fig. 1. Proposed Approach, the modified RPN is followed by Y-Net and the bounding box extractor.

image. To address this dataset availability problem, we have released 100,000 UHR and 100,000 LR synthetic barcode datasets along with their corresponding bounding boxes ground truths, and their ground truth masks to facilitate further studies. The outline of this paper is as follows: in Section 2, we describe details of our approach; in Section 3, we summarize our experimental results and in Section 4, we conclude and expend on our future work.

## 2. PROPOSED APPROACH

As seen in Fig. 1, our proposed method consists of three stages: the modified Region Proposal Network stage, our Y-Net<sup>1</sup> segmentation network stage, and the bounding box extraction stage.

### 2.1. Modified Region Proposal Network

Region proposals have been influential in computer vision and more so when it comes to object detection in UHR images. It is common in UHR images that barcodes are clustered in a small region of the image. To filter out most of the non-barcode backgrounds, we modified the RPN introduced in Faster R-CNN [13] to propose regions of barcodes for our next stages. By first transforming the UHR input image to an LR input image of size  $256 \times 256$ , the RPN was trained to identify blobs in LR images. Once a bounding box is placed around the identified blobs, the resulting proposed bounding box is remapped to the input UHR image by a perspective transformation, and the resulting regions are cropped out. The LR input to the RPN is chosen to be of size  $256 \times 256$  as a lower resolution results in the loss of pertinent information. Non-Max Suppression (NMS) is used on the predictions to select the most probable regions.

### 2.2. Y-Net Segmentation Network

As depicted in Fig. 3, Y-Net is made out of 3 main modules distributed in 2 branches: a Regular Convolutional Module shown in blue which constitutes the left branch, and a Pyramid Pooling Module shown in brown, along with a Dilated Convolution Mod-



**Fig. 2.** Sample outputs of our pipeline; yellow - segmented barcode pixels; purple - segmented background pixels; boxes - bounding box extracted; (a) synthetic barcode image; (b) real barcode image; (c) prediction results on (a); (d) prediction results on (b).

ule shown in orange which after concatenation and convolution constitute the right branch.

The **Regular Convolution Module** takes in  $400 \times 400$  output images of the RPN and consists of convolutional and pooling layers. It starts with 64 - channel  $3 \times 3$  kernels and doubles the number at each layer. We alternate between convolution and max-pooling until we reach a feature map size of  $25 \times 25$  pixels. This module allows the model to learn general pixel-wise information anywhere in the input image.

The *Dilated Convolution Module* takes advantage of the fact that barcodes have alternating black and white rectangles to learn sparse features in their structure. The motivation for this module comes from the fact that dilated convolution operators play a signif-

<sup>&</sup>lt;sup>1</sup>Our Y-Net architecture resembles the English alphabet letter "Y" and differs from [15] which used a pre-trained encoder network that is augmented with an untrained mirrored network and a decoder network.



Fig. 3. Y-Net Architecture.

icant role in the "algorithme a trous" for biorthogonal wavelet decomposition [16]. Therefore, the discontinuities in alternating patterns and sharp edges in barcodes are more accurately learned by such filters. In addition, they leverage a multiresolution and multiscale decomposition as they allow the kernels to widen their receptive fields with dilation rates from 1 up to 16. Here too a  $400 \times 400$ input image is used and we maintain 32 – channel  $3 \times 3$  kernels throughout the module while the dimensions of the layers are gradually reduced using a stride of 2 until a feature map of  $25 \times 25$  pixels is obtained.

The *Pyramid Pooling Module* allows the model to learn global information about potential locations of the barcodes at different scales and its layers are concatenated with the layers on the dilated convolution module in order to preserve the features extracted from both modules.

The resulting feature maps from the right branch are then added to the output of the Regular Convolution Module, which allows for the correction of features that would have been missed by either branch. In other words, the output of each branch constitutes a residual correction for the other thereby refining the result at each node as shown in white. The nodes are then up-sampled and concatenated with transposed convolution feature maps shown in red and yellow of the corresponding dimension. Throughout the network, we use ReLU as a non-linearity after each layer and add L<sub>2</sub> regularization to account for possible over-fitting scenarios that could have occurred during training. On all datasets, we use 80% for the training set, 10% for the validation set, and the remaining 10 % for the testing set. We use one NVIDIA Tesla V100 GPU for the training process. Since this is a segmentation network and we are interested in classifying background and barcodes, we use binary cross-entropy as loss function.

## 2.3. Bounding Box Extraction

Since some images contain barcodes that are really close to each other, their Y-Net outputs reflect the same configuration which makes the extraction of individual barcode bounding boxes complex as shown in Fig. 4(a). To separate them effectively, we perform an erosion, contour extraction, and bounding box expansion with a pixel correction margin. As shown in Fig. 4(b), the erosion stage allows the algorithm to widen gaps between segmented barcodes that may be separated by 1 or more pixels. The resulting mask is then used to infer individual barcode bounding boxes in the contour extraction stage in Fig. 4(c) through border following. A pixel correction margin is used to recover the original bounding boxes' dimensions during the expansion stage as shown in Fig. 4(d). This post-processing stage of our pipeline has an average processing time of 1.5 milliseconds (ms) because it is made of a set of Python matrix operations to efficiently extract bounding boxes from predicted masks.

## 3. DATASETS AND RESULTS

For the synthetic dataset, we use treepoem <sup>2</sup> and random-word <sup>3</sup> to generate UHR and LR barcode images. We use Code 39, Code 93, Code 128, UPC, EAN, PD417, ITF, Data Matrix, AZTEC, and QR among others. We model the number of barcodes in a given image using a Poisson process and a combination of perspective transforms is used to make the barcodes vary in shape and position from one image to the other. We have also added random black blobs at random

<sup>&</sup>lt;sup>2</sup>https://github.com/adamchainz/treepoem

<sup>&</sup>lt;sup>3</sup>https://github.com/vaibhavsingh97/random-word

	mAP	$AP_{50}$	$AP_{75}$	mAP	mAP	$AR_{50}$	$AR_{70}$	$AR_{80}$	$AR_{90}$	Latency	Resolution
	(all)	(all)	(all)	(small)	(medium)	(all)	(all)	(all)	(all)	(ms)	(px)
Mask R-CNN [10]	.466	.985	.317	.340	.489	.990	.740	.279	.023	94.8	$448 \times 448$
YOLOv4 [11]	.882	.990	.989	.815	.897	1.	1.	.995	.873	40.5	$320 \times 320$
Ours	.937	.990	.990	.903	.945	1.	1.	1.	.972	16.0	$400 \times 400$

Table 1. Average Precision for Max Detection of 100 and Average Recall for Max Detection of 10 computed using MS COCO API.

		Muenster	Dataset			ArTe Lab Dataset			
	DR	Precision	Recall	mIoU	DR	Precision	Recall	mIoU	
Creusot et al. [5]	.982	-	-	-	.989	-	-	-	
Hansen et al. [9]	.991	-	-	.873	.926	-	-	.816	
Namane et al. [1]	.966	-	-	.882	.930	-	-	.860	
Zharkov et al. [12]	.980	.777	.990	.842	.989	.814	.995	.819	
ours	1.	.984	1.	.921	1.	.974	1.	.934	

Table 2. Mean IoU (mIoU), Precison and Recall and Detection Rate (DR) at IoU threshold of 0.5 (Muenster and ArTe-Lab Dataset).

	Px Acc	Px mIoU	Px Prec	Px Rec
Mask R-CNN [10]	.993	.990	.989	.890
Ours	1.	1.	.999	.999

#### Table 3. Pixel-wise Metrics

locations on the original UHR and LR canvases. The real UHR barcodes dataset obtained from Amazon.com, Inc is made of 3.8 million UHR images of resolution up to  $30k \times 30k$  grayscale images and could not be released due to confidentiality reasons. Additionally, the Muenster and Artelab datasets are used with some data augmentation schemes for more samples.

For the RPN, we accumulated the number of bounding boxes inside the proposed regions and divided it by the total number of ground truth bounding boxes. Our implementation yields an accuracy of 98.03% on the synthetic dataset at 10 ms per image and 96.8% on the real dataset at 13 ms per image while the baseline [13] yields the same accuracies and an average latency over 2.5 seconds (s) per image for both datasets.

For Y-Net, we use the Microsoft (MS) COCO API, and Pixelwise metrics to evaluate against [10, 11]. By default, the MS COCO API configuration evaluates on *small, medium* and *large* areas objects but in our application, the largest detected barcode area is *medium*. Since Y-Net is a segmentation network and does not output confidence scores for each segmented barcode, we propose using *pseudo scores*, the ratio of the total number of nonzero pixels in a predicted mask to the total number of nonzero pixels in the corresponding ground truth mask at the location of a given object.

Table 1 shows mAP and mAR values of the models on the synthetic dataset. As seen, our pipeline outperforms [10], and [11] by a mAP of 47.1% and 5.5% and  $AP_{75}$  of 67.3% and 0.1% respectively. Also shown in Table 1, is a  $mAR_{90}$  improvement of 94.9% and 9.9% on [10] and, [11] respectively which highlights that Y-Net continues to yield better mAR results even at higher IoU thresholds. Both our approach and [11] achieve an  $AR_{50}$  of 100% and outperform [10] by 1%. For *small* area barcodes, Y-Net outperforms [10] and [11] by a mAP of 56.3% and 8.8% and for *medium* area barcodes, Y-Net displays a mAP increase of 45.6% and 4.8% on [10] and [11] respectively. In addition, Table 3 reveals that Y-Net a has much better semantic segmentation performance than [10]. Table 1 displays that Y-Net performs at least  $2.5 \times$  faster than the fastest of models [10] and, [11] on LR images.

Similarly, we have used the Detection Rate (DR), mIoU, Precision, and Recall, as described in [1, 5, 9, 12] on the Arte-Lab and Muenster datasets and as can be seen in Table 2, our method outperforms previous works on all of the mentioned metrics. This indicates that our bounding box extraction algorithm is working as expected to detect accurate bounding boxes. However, while it is successful in separating barcodes that are relatively close to each other, it has limitations when barcodes are overlapping as shown in Fig. 4(e). For those occlusion scenarios, the algorithm tends to group the overlapping barcodes into one bounding box instead of separate bounding boxes as shown in Fig. 4(f).



**Fig. 4.** (a) Y-Net output; (b) Y-Net output after erosion; (c) extracted bounding boxes –red, ground truth bounding boxes –green on eroded output; (d) final bounding boxes after pixel correction margin on Y-Net output; (e) Y-Net output of occluded barcodes scenarios; (f) final extracted bounding boxes are grouped after pixel correction margin due to overlaping barcodes in input image.

#### 4. CONCLUSION

In this paper, we showed that barcodes can be efficiently, accurately, and speedily detected using Y-Net on UHR images. With *pseudo scores* as *confidence scores*, our approach outperforms existing detection pipelines with a much better latency. In future work, we aim to extend this method to the multi-class detection task for small objects in UHR images and videos in a weakly supervised fashion.

### 5. REFERENCES

- A. Namane and M. Arezki, "Fast real time 1d barcode detection from webcam images using the bars detection method," in *Proceedings of the World Congress on Engineering (WCE)*, 2017, vol. 1.
- [2] L. Hock, H. Hanaizumi, and E. Ohbuchi, "Barcode readers using the camera device in mobile phones," in 2013 International Conference on Cyberworlds. 2004, pp. 260–265, IEEE Computer Society.
- [3] M. Katona and L. G Nyúl, "Efficient 1d and 2d barcode detection using mathematical morphology," in *Mathematical Morphology and Its Applications to Signal and Image Processing*. 2013, pp. 464–475, Springer Berlin Heidelberg.
- [4] G. Sörös and C. Flörkemeier, "Blur-resistant joint 1d and 2d barcode localization for smartphones," in *Proceedings of the* 12th International Conference on Mobile and Ubiquitous Multimedia. 2013, MUM '13, Association for Computing Machinery.
- [5] C. Creusot and A. Munawar, "Low-computation egocentric barcode detector for the blind," in 2016 IEEE International Conference on Image Processing (ICIP), 2016, pp. 2856–2860.
- [6] C. Creusot and A. Munawar, "Real-time barcode detection in the wild.," *IEEE Winter Conference on Applications of Computer Vision*, p. 239–245, 2015.
- [7] O. Gallo and R. Manduchi, "Reading 1d barcodes with mobile phones using deformable templates.," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 1834– 1843, 2011.
- [8] A. Zamberletti, I. Gallo, M. Carullo, and E. Binaghi, "Neural image restoration for decoding 1-d barcodes using common camera phones.," 01 2010, vol. 1, pp. 5–11.
- [9] D. K. Hansen, K. Nasrollahi, C. B. Rasmussen, and T. B. Moeslund, "Real-time barcode detection and classification using deep learning," *IJCCI*, vol. 1, pp. 321–327, 2017.
- [10] P. Dollár K. He, G. Gkioxari and R. Girshick, "Mask R-CNN," in 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2980–2988.
- [11] A. Bochkovskiy, C. Wang, and H. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," 2020.
- [12] A. Zharkov and I. Zagaynov, "Universal barcode detector via semantic segmentation," 2019 International Conference on Document Analysis and Recognition (ICDAR), pp. 837–843, 2019.
- [13] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [14] X. Jiang S. Wachenfeld, S. Terlunen, Robust 1-D Barcode Recognition on Camera Phones and Mobile Product Information Display, pp. 53–69, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [15] Ahmed Mohammed, Sule Yildirim, Ivar Farup, Marius Pedersen, and Øistein Hovde, "Y-net: A deep convolutional neural network for polyp detection," 2018.

[16] M. Holschneider, R. Kronland-Martinet, J. Morlet, and Ph. Tchamitchian, "A real-time algorithm for signal analysis with the help of the wavelet transform," *In Wavelets: Time-Frequency Methods and Phase Space. Proceedings of the International Conference*, 1987.