

A Systematic Evaluation of Object Detection Networks for Scientific Plots

Pritha Ganguly*, Nitesh S Methani*, Mitesh M. Khapra, Pratyush Kumar

Robert Bosch Centre for Data Science and AI (RBC-DSAI),
Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India.
{prithag, nmethani, miteshk, pratyush}@cse.iitm.ac.in

Abstract

Are existing object detection methods adequate for detecting text and visual elements in scientific plots which are arguably different than the objects found in natural images? To answer this question, we train and compare the accuracy of Fast/Faster R-CNN, SSD, YOLO and RetinaNet on the PlotQA dataset with over 220,000 scientific plots. At the standard IOU setting of 0.5, most networks perform well with mAP scores greater than 80% in detecting the relatively simple objects in plots. However, the performance drops drastically when evaluated at a stricter IOU of 0.9 with the best model giving a mAP of 35.70%. Note that such a stricter evaluation is essential when dealing with scientific plots where even minor localisation errors can lead to large errors in downstream numerical inferences. Given this poor performance, we propose minor modifications to existing models by combining ideas from different object detection networks. While this significantly improves the performance, there are still two main issues: (i) performance on text objects which are essential for reasoning is very poor, and (ii) inference time is unacceptably large considering the simplicity of plots. To solve this open problem, we make a series of contributions: (a) an efficient region proposal method based on Laplacian edge detectors, (b) a feature representation of region proposals that includes neighbouring information, (c) a linking component to join multiple region proposals for detecting longer textual objects, and (d) a custom loss function that combines a smooth ℓ_1 -loss with an IOU-based loss. Combining these ideas, our final model is very accurate at extreme IOU values achieving a mAP of 93.44%@0.9 IOU. Simultaneously, our model is very efficient with an inference time 16x lesser than the current models, including one-stage detectors. Our model also achieves a high accuracy on an extrinsic plot-to-table conversion task with an F1 score of 0.77. With these contributions, we make a definitive progress in object detection for plots and enable further exploration on automated reasoning of plots.

Introduction

Object detection is one of the fundamental problems in computer vision with the aim of answering *what objects are where* in a given input image. Most of the object detection research in the past few years has been on natural images

*The first two authors have contributed equally
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Demonstrating sensitivity to IOU on images from (a) PASCAL VOC and (b) PlotQA dataset. Localising on the natural images in (a) is accurate even at low IOU thresholds of 0.5 and 0.75. For the plots in (b), the comparison of plotted values for the two years is incorrect at lower IOU values of 0.5 (left) and 0.75 (centre), and is correct for 0.9 (right).

with real-life objects. For instance, in the PASCAL VOC dataset (Everingham et al. 2010), the four major classes of objects are people, animals, vehicles, and indoor objects such as furniture. In this work, we study object detection for a very different class of images, namely computer-generated scientific plots. Fig. 1b shows an example of a scientific plot: It is a bar plot depicting the number of neonatal deaths in Bulgaria and Cuba over two years. Object detection on this plot would be required to identify different visual and textual elements of the plot such as bars, legend previews and labels for ticks, axes & legend entries. Such object detection can then enable a question-answering task. For instance, for Fig. 1b we could ask “In which year does Cuba have lower neonatal deaths?”. Clearly, this has use-cases in data analytics, and has been studied in recent research (Kafle et al. 2018; Kahou et al. 2018; Methani et al. 2020).

It should be clear that scientific plots differ significantly from natural images. Firstly, plots have a combination of textual and visual elements interspersed throughout. The text can either be very short (such as numerical tick labels) or span multiple lines (such as in plot-titles). Secondly, objects in a plot have a large range of sizes and aspect ratios. Depending on the value represented, a bar in a bar-plot can be short or long, while in a line-plot a thin line could depict the data. Thirdly, plots impose structural relationships between some of the objects. For instance, the legend preview and the close-by title text denote a correspondence map. This also

applies to a tick label and its corresponding bar in a bar-plot.

Given these differences, it needs to be seen if existing object detection methods are adequate for scientific plots. In particular, are they capable of (a) detecting short and long pieces of text, (b) detecting objects with large data-dependent range of sizes and aspect ratios, and (c) localising objects accurately enough to extract structural relationships between objects? To answer this question, we first evaluate state-of-the-art object detection networks on the PlotQA dataset (Methani et al. 2020) which has over 220,000 scientific plots sourced from real-world data thereby having a realistic vocabulary of axes variables and complex data patterns which mimic the plots found in scientific documents and reports. We observe that, across these models, the average of mAP@0.5 is only around 87%, indicating success in detecting the relatively simple objects in the plot.

While the above results appear positive, a closer manual inspection revealed that these models make critical errors which lead to large errors in downstream numerical inference on the plots. This disparity is because we use an IOU of 0.5 while computing the mAP scores. While IOU values in the range of 0.5 and 0.7 are acceptable for natural images where large relative areas are covered by foreground objects, such values are unacceptably low for scientific plots. This is demonstrated in Fig. 1a with two example images from the PASCAL VOC dataset where the predicted box (red) is very different from the ground-truth box (cyan), but still acceptable as the IOU is within range. Contrast this with the case for the plot in Fig. 1b. For an IOU setting of 0.5 (left) and 0.75 (middle), the estimated values of the data-points would incorrectly identify that neonatal deaths in Cuba are lower in 2002 than the actual value in 2003. Only at the high IOU value of 0.9, this is correctly resolved. Thus, downstream numerical reasoning on plots requires much stricter IOU settings in comparison to object detection on natural images. For the PlotQA dataset if we use a stricter IOU of 0.9, then the mAP scores for all models drop drastically with the best model giving a mAP of only 35.70%. In particular, one-stage detectors such as SSD (Liu et al. 2016) and YOLO-v3 (Redmon and Farhadi 2018) have a single-digit mAP@0.9.

The poor performance of current models at high IOU settings motivate improvements in the models. We first propose minor modifications to existing models. In particular, we propose a hybrid network which combines Faster R-CNN (Ren et al. 2015) with the Feature Pyramid Network (Lin et al. 2017a) from RetinaNet (Lin et al. 2017b) and the ROIAlign idea from Mask R-CNN (He et al. 2017). This significantly improves the performance giving an overall mAP of 77.22%@0.9 IOU. However, careful analysis reveals two major limitations: (i) accuracy on text objects is very low which can lead to errors in downstream analytics tasks, and (ii) the inference time is very high (374 ms) which is unacceptable given the lower visual complexity of plots.

To further improve the speed and performance, we propose an architecture, named PlotNet, which contains (i) a fast and conservative region proposal method based on Laplacian edge detectors, (ii) an enhanced feature representation of region proposals that includes local neighbouring information, (iii) a linking component that combines mul-

tipple region proposals for better detection of longer textual objects, and (iv) a custom regression loss function that combines smooth ℓ_1 -loss with an IOU-based loss designed for improving localisation at higher IOUs. This significantly improves accuracy with a mAP of 93.44%@0.9 IOU. Further, it is 16x faster than its closest competitor and has 3x lower FLOPs. We also evaluate PlotNet on an extrinsic plot-to-table conversion task where we extract the plot’s underlying data into a table. Specifically, we replace Faster R-CNN with PlotNet in the Visual Element Detection stage of the pipeline proposed in PlotQA (Methani et al. 2020). This results in a relative improvement of 35.08% in the table F1 score.

In summary, the contributions of this paper are as follows:

1. We motivate the need for object detection at extreme IOU values for the specific dataset of scientific plots which require accurate localization.
2. We evaluate the robustness of nine different object detection networks (SSD, YOLO-v3, RetinaNet, variants of Fast and Faster R-CNN) to increase in IOU and identify that Feature Pyramid Network (FPN) and ROIAlign (RA) are good design choices for higher accuracy.
3. We propose PlotNet that improves on mAP by over 16 points while reducing execution time by over 16 times from its closest competitor. Thus, PlotNet is faster than one-stage detectors and simultaneously more accurate than the best two-stage detectors.

The rest of the paper is organised as follows: We first discuss the different datasets and detail the experimental setup for evaluating existing object detection models on scientific plots. We also report the results and critique their performance on an example plot image from the PlotQA dataset. We then detail the architecture of PlotNet and compare it with other networks in terms of accuracy and speed. Note that, we follow an unconventional organisation of describing partial experimental results first as this lays the groundwork for motivating the design and evaluation of PlotNet.

Evaluation of Existing Models

Dataset: Automated visual analysis and subsequent question-answering on scientific plots was first proposed in FigureQA (Kahou et al. 2018). There are three publicly available datasets, namely FigureQA (Kahou et al. 2018), DVQA (Kafle et al. 2018), and PlotQA (Methani et al. 2020). These datasets contain scientific plots with bounding boxes and labels for different plot elements including bars, lines, tick labels, legend entries, and plot labels. We run our experiments on the PlotQA dataset (Methani et al. 2020), as it is based on real-world data while both FigureQA and DVQA are based on synthetic data. For instance, in synthetic datasets, the label names are selected from a limited vocabulary such as colour names in FigureQA and top-1000 nouns in the Brown corpus in DVQA. On the other hand, PlotQA has datasets collected from public data repositories. This impacts object detection as the text labels show large variability in PlotQA dataset. Secondly, synthetic datasets use limited range of plotted values such as integers in a fixed range, while PlotQA plots real data. This impacts object de-

Existing Models	IOU @0.9									@0.75		@0.5
	bar	dot-line	legend label	legend preview	plot title	x-axis label	x-axis ticks	y-axis label	y-axis ticks	mAP	mAP	mAP
(a) SSD	1.39	0.60	2.18	0.39	0.04	3.39	0.44	5.14	0.20	1.53	39.78	82.33
(b) YOLO-v3	15.51	8.72	7.15	11.70	0.02	4.39	8.08	9.59	1.70	7.43	73.31	96.27
(c) RetinaNet	16.51	18.5	77.26	29.74	16.58	67.62	28.40	3.14	17.31	30.56	81.13	90.13
(d) FRCNN	53.38	1.68	12.59	14.06	0.03	42.13	25.49	11.68	31.98	21.45	63.68	72.83
(e) FrRCNN	6.92	1.68	1.39	1.45	0.00	4.35	6.10	3.57	5.18	4.08	50.51	88.49
(f) Mask R-CNN	47.54	5.36	50.83	32.43	0.33	40.20	33.72	80.53	30.31	35.70	82.45	93.72
(g) FRCNN (FPN+RA)	87.59	31.62	79.05	66.39	0.22	69.78	88.29	46.63	84.60	61.57	69.82	72.18
(h) FrRCNN (RA)	63.89	14.79	70.95	60.61	0.18	83.89	60.76	93.47	50.87	55.49	89.14	96.80
(i) FrRCNN (FPN+RA)	85.54	27.86	93.68	96.30	0.22	99.09	96.04	99.46	96.80	77.22	94.58	97.76

Table 1: Comparison of existing and hybrid models on the PlotQA dataset with mAP scores (in %) at IOUs of 0.9, 0.75, and 0.5. For IOU@0.9, the class-wise average precision (in %) is shown for all the classes.

tection as the size of bars in a bar-plot and the slopes in a line-plot show large variability in the PlotQA dataset.

The PlotQA dataset (Methani et al. 2020) contains over 220,000 scientific plots across three categories of bar (both horizontal and vertical), line, and scatter plots. The dataset includes ground-truth bounding boxes for bars, dot-lines, legend labels, legend previews, plot-title, axes ticks and their labels. The underlying data is from data sources such as World Bank Open Data containing natural variable names such as mortality rate, crop yield, country names, and so on.

Training Setup: We used the existing implementations for the R-CNN family, YOLO-v3, SSD and RetinaNet. ResNet-50 (R-50) pre-trained on ImageNet (Deng et al. 2009) dataset is the backbone feature extractor for Fast R-CNN, Faster R-CNN, Mask R-CNN and RetinaNet. For SSD and YOLO-v3, InceptionNet (Szegedy et al. 2015) and DarkNet-53 were the backbone feature extractors, respectively.

These models were trained with an initial base learning rate of 0.025 with momentum stochastic gradient descent algorithm. The network’s classification and regression heads use a batch-size of 512 ROIs. RetinaNet and SSD models were trained with a batch-size of 32 with a learning rate of 0.004. Based on evaluation on the validation dataset, we modified the parameters in the focal loss for RetinaNet to $\alpha = 0.75$ and $\gamma = 1.0$, against recommended values of $\alpha = 0.25$ and $\gamma = 1.0$. The model was trained with a batch-size of 64 and a learning rate of 0.001.

Results & Comparative Analysis: For each of the models, mAP score at three different IOU values of 0.9, 0.75, and 0.5 are shown in Table 1. Here are the important observations:

- mAP@0.5 is fairly high (average over 87%) across models indicating that the relatively simple visual elements of the plots are being identified with high accuracy.
- mAP@0.75 falls markedly in comparison to mAP@0.5, with an average drop of about 22 points across models. Specifically for SSD, Faster R-CNN the drop is very high at about 40 points.
- mAP@0.9 drops to remarkably low values; on average mAP@0.9 is less than half of mAP@0.75. Specifically, SSD, Faster R-CNN, and YOLO have single-digit

mAP@0.9 values.

- For an IOU of 0.9, the AP for individual classes shows large variability across models. Relatively, plot-title and dot-line classes have the lowest AP values across models.

To better illustrate the performance of each model, we exemplify the bounding box outputs of the different models on specific parts of an example plot shown in Fig. 2. We make the following observations, model-wise.

- **SSD** glaringly misses detecting one of the bars, and also has low localisation accuracy as evidenced in the misaligned bounding for the second bar in Fig. 2a. However, it correctly detects small tick labels, perhaps due to proposal generation performed at multiple resolutions.
- **YOLO-v3** detects all objects (including both bars), but with lower localisation accuracy. For instance in Fig. 2b, the upper bar and plot title have misaligned bounding boxes. To see if this problem could have been solved by imposing priors on aspect ratios of bounding boxes, we plotted aspect ratios of all objects across plots and found no distinct clusters, *i.e.*, aspect ratios of bars, texts, *etc.* vary in a continuum. This makes it hard to choose appropriate priors for bounding boxes.
- **RetinaNet** which is based on SSD also misses out on detecting one of the bars and also the y-tick label (Fig. 2c). The bounding box of the detected bar is more accurate than that in SSD, indicating the benefit of the lateral connections in generating features for the regression head. Across the three one-stage detectors, which have much higher speed than the two-stage detectors, RetinaNet is the clear winner (row (c) in Table 1). While not apparent in the illustrated example, RetinaNet’s focal loss with custom tuned parameters (α, β) instead of hard suppression, may also be contributing to its higher performance.
- **Fast R-CNN (FRCNN)** breaks up one of the bars into smaller objects (aligning with lines on the background grid of the plot). It also misses several objects including a legend item, the plot title, and a tick label (Fig. 2d). This could be attributed to the proposal generation method which uses selective search (SS). This under-performance

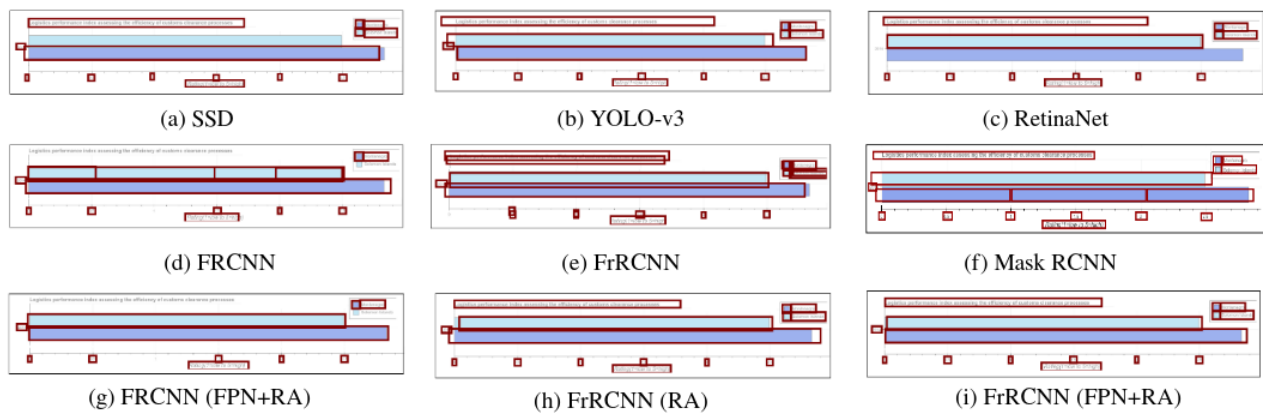


Figure 2: Detected bounding boxes on an example plot from PlotQA for different models corresponding to Table 1 at IOU 0.9.

is also visible at low IOUs: mAP@0.5 is lowest for Fast R-CNN (row (d) in Table 1) potentially due to poorly performing SS which remains unaffected by IOU.

- **Faster R-CNN (FrRCNN)** improves over the recall of SS by detecting most objects due to more complex region proposal network (RPN). However, RPN creates multiple overlapping proposals, even after non-maximal suppression (NMS) (Fig. 2e). This lowers the bulky model’s mAP to just 4%, which is second-lowest.
- **Mask R-CNN** uses Faster R-CNN as the backbone architecture but uses ROIAlign instead of ROIpool gives mixed results when compared to Faster R-CNN. It is able to detect the longer textual elements (*e.g.*, title) but has poorer localisation accuracy on the bars. It also breaks up one of the bars into smaller objects. However, its localisation accuracy on the tick labels is better than Faster R-CNN.

In summary, most state-of-the-art models have low robustness to high IOU values on this different class of images.

A hybrid network combining existing ideas: The above discussion clearly establishes the need for better models for object detection over scientific plots. However, before we do so, we wanted to examine if combining ideas from existing models can help in improving the performance. Among the one-stage detectors, RetinaNet (row(c) in Table 1) gave the best performance. Similarly, among the two-stage detectors Mask R-CNN (row(f) in Table 1) gave the best performance. However, the qualitative analysis presented in Fig. 2 suggested that Faster R-CNN has some advantages over Mask R-CNN. Taking all of this into consideration we decided to combine the relative merits of Faster R-CNN, RetinaNet and Mask R-CNN. In particular, we retain the overall architecture of Faster R-CNN but use FPN as the feature extractor (as in RetinaNet) and replace ROIpool with ROIAlign (as in Mask R-CNN). The results obtained by making these modifications are summarised in rows (g), (h) and (i) of Table 1. For the sake of completeness we also present results obtained by combining Fast R-CNN (FRCNN) with FPN and ROIAlign. We observe that across all three IOU values, the highest mAP values are obtained by Faster R-CNN with FPN and ROIAlign (row(i) in Table 1). This mAP of 77.22%

is the best number reported in Table 1.

We do the same qualitative analysis as before for the three hybrid models and make the following observations:

- **Fast R-CNN with FPN and ROIAlign (FRCNN (FPN+RA))** improves on Fast R-CNN by not breaking up the bar into smaller objects, due to the FPN which enables improved feature extraction for the regression head. However due to the use of selective search (SS), many objects continue to go undetected (Fig. 2g). Notably, for the detected objects, the localisation accuracy is high, perhaps due to the use of ROIAlign’s bilinear interpolation when mapping proposals into smaller cells.
- **Faster R-CNN with ROIAlign (FrRCNN (RA))** improves on Faster R-CNN due to the substitution of ROIpool with ROIAlign. This leads to more accurate localisation and thereby removal of multiple proposals by NMS (Fig. 2h). Interestingly, this model under-performs FRCNN (FPN+RA) as evident in the mAP values and in the example by the difference in localisation accuracy. This illustrates the importance of FPN in being able to nullify the limitations of SS and improve localisation.
- **Faster R-CNN with FPN and ROIAlign (FrRCNN (FPN+RA))** performs the best amongst all existing models. It combines the RPN of Faster R-CNN for better region proposals, with FPN which provide better features, and ROIAlign which provides better mapping of features to scaled cells (Fig. 2i). There is an additive effect of combining these three important ideas in object detection, as evidenced by the significant difference between this model and the rest.

In summary, the model with the highest performance is the one that combines the best ideas in object detection.

Our Proposed Model

Given the low mAP scores of existing models at the requisite IOU of 0.9, we propose a new network (shown in Fig. 3) which is designed bottom up based on three key observations. First, we observe that networks which use existing region proposal methods such as selective search, RPN, and

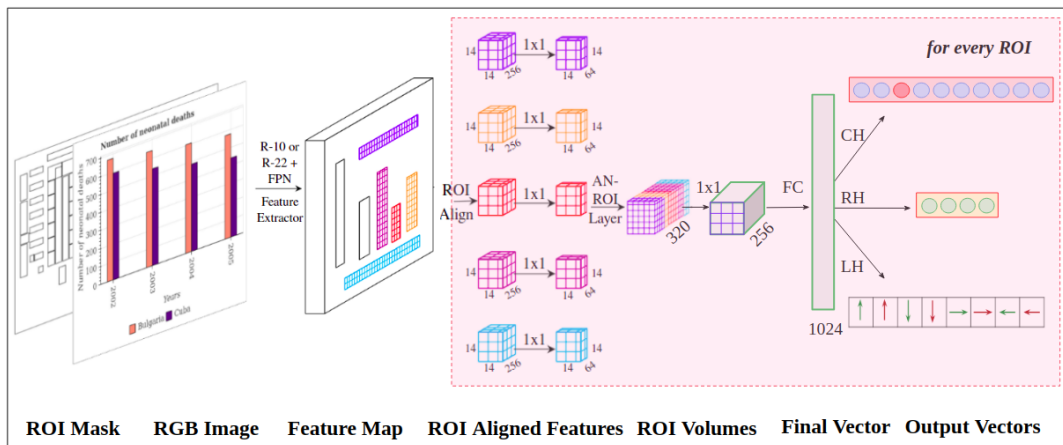


Figure 3: Architecture diagram of our proposed model, PlotNet.

anchor based methods have low mAP@0.9. In particular, these methods either generate too many proposals or miss out some objects. We contrast this with the apparent low visual complexity of these plots which suggests that detecting region proposals should be easier. Based on this insight, we propose a region proposal method which relies on traditional CV-based methods such as edge detection followed by contour detection. We however retain ROIAlign and FPN components which improved the performance of models. However, we note that FPN adds a significant computational cost and its addition needs to be carefully evaluated. While doing so, we make our second observation that longer textual elements such as titles and legend labels get detected as multiple proposals which need to be linked. We propose a separate linking component which decides whether a given proposal needs to be merged with any of its neighbours. None of the existing models perform such linking. Third, we notice a sharp decline of mAP scores on increase in the IOU. To address this, we design a custom loss function, which has non-negligible loss values for high IOU (> 0.8).

In summary, in the design of a custom model which we refer to as PlotNet, we (i) use a computationally efficient CV-based region proposal method, (ii) supplement it with a link prediction method to detect contiguous text objects, (iii) use ROIAlign and neighbouring information to better map proposals into smaller cells, (iv) evaluate the necessity of FPN, and (v) evaluate the need for IOU-based loss functions. We now describe the different components of our model.

Region Proposal: As an alternative to Region Proposal Network (RPN), we propose a combination of CV methods to generate region proposals (Fig. 4(a)). Specifically, we (a) draw edges with a Laplacian edge detector on the image of the plot, (b) extract continuous closed contours from edges based on uniform color and/or intensity, (c) convert contours to a bounding-box by finding the minimal up-right bounding rectangle for each of the identified contours. These boxes serve as regions of interest (ROI) which are passed as input to the network. These proposals are very small in number (90 proposals on an average) in comparison to selective search which gives around $2k$ proposals. Further, while se-

lective search takes ~ 6740 ms per image on average to generate proposals, our method takes only 34ms. We refer the readers to the supplementary section for more details.

Feature Extraction: We use ResNet (He et al. 2016) for extracting features from the input image after resizing the image to 650×650 . To exploit structural information present in the image, we add the ROIs proposed earlier as the 4^{th} channel to the RGB input image. We tried different number of layers in the ResNet model & found that even with 10 layers we were able to get a good performance. We also consider FPN as the feature extractor with ResNet-22 as the backbone architecture as a potential trade-off between performance and cost. Once the image features are extracted, the externally generated ROIs are projected onto the feature map. To extract a fixed sized ROI feature, we pass them through the ROIAlign layer (He et al. 2017) which outputs the fixed size feature map of size $14 \times 14 \times 256$. We further reduce the depth of each ROI feature map to $14 \times 14 \times 64$ by using 1×1 convolution layers.

Accumulating Neighbouring ROI Information (AN-ROI layer): To incorporate local neighbouring information into each ROI feature, we create an AN-ROI volume by concatenating the ROI's immediate left, right, top and bottom features along the depth, resulting in a feature volume of size $14 \times 14 \times 320$. We then apply convolutional layers on this AN-ROI volume resulting in a feature volume of size $14 \times 14 \times 256$. We hypothesise that such neighbouring features would increase the accuracy of classifying, regressing, or linking individual ROIs.

Classification (CH), Regression (RH) & Linking heads (LH): The ROI features extracted above are passed through two fully-connected layers with 2048 and 1024 neurons, respectively. Each ROI feature vector, is independently passed through the CH which uses the softmax function to output a probability distribution over the 9 classes of objects in the images and a tenth background class. These 9 classes are bar, dot-line, legend-preview, legend-label, plot-title, x-axis label, y-axis label, x-axis ticks and y-axis ticks. The same ROI feature vector is also fed to the RH which predicts the 4 co-ordinates (top-left and bottom-right) of the bounding




Figure 4: (a) The proposals generated by our CV method are shown in red. (b) Detected bounding boxes by PlotNet on an example plot from PlotQA dataset at IOU 0.9.

box. Lastly, the same ROI feature vector is passed to the LH which predicts whether this ROI needs to be merged with none, one, or more of its immediate 4 (top, left, right and bottom) neighbours.

IOU-based Loss functions: Most state-of-the-art object detection models use ℓ_n loss (e.g., ℓ_2 -loss, smooth ℓ_1 -loss) for performing bounding box regression. However, several studies (Yu et al. 2016; Berman, Triki, and Blaschko 2018) suggest that there are some disadvantages of doing so and instead an IOU-based loss function which better correlates with the final evaluation metric should be used. Indeed, some studies have showed that using $-\log IOU$ (Yu et al. 2016) and $1 - IOU$ (Berman, Triki, and Blaschko 2018) as loss functions give better results by ensuring that (i) the training objective is aligned with the evaluation metric, and (ii) all the 4 coordinates of the bounding box are considered jointly. However, these loss functions fail to learn anything in case of non-overlapping boxes. To overcome this, more generalized loss functions like GIOU (Rezatofighi et al. 2019), DIOU, and CIOU (Zheng et al. 2020) have been proposed which add an additional penalty term to deal with the non-overlapping boxes.

In existing loss functions, the penalty is negligible for boxes which have a large IOU overlap with the ground-truth box. To enable the network to learn tighter bounding boxes, motivated by focal loss (Lin et al. 2017b), we propose a Focal IOU (FIOU) loss that more gradually decreases the penalty as IOU overlap with ground-truth box increases. Formally, FIOU is defined as:

$$\mathcal{L}_{FIOU} = -(1 + IOU)^\gamma * \log(IOU)$$

FIOU focuses on higher IOU values by dynamically scaling the $\log(IOU)$ loss (Yu et al. 2016), where the scaling factor increases as the IOU increases. The higher the hyper-parameter γ , the larger is the penalty at medium IOU levels. We experimented with a couple of γ values and have found $\gamma = 2$ to work best for the PlotQA dataset. We include a plot comparing these loss functions graphically in the supplementary section. We then define a custom loss function

combining smooth ℓ_1 -loss (SL1) with FIOU: $\mathcal{L}_{Custom} = \mathcal{L}_{SL1} + \mathcal{L}_{FIOU}$. This custom loss function achieves state-of-the-art results as we report in the comparison across loss functions in Table 2.

Training & Implementation Details: While training, for every proposed ROI, we need to assign a ground-truth class for the 9 object classes and the background class. We identify the ROI’s center and identify if it lies in any of the ground-truth bounding boxes. There would be at most one such box, since objects in plots do not overlap unlike natural images. If no such ground-truth bounding box is found, the ROI is considered to be in the background class and is ignored by the regression head.

Similarly, for every proposed ROI, the coordinates of the parent ground-truth box identified above are assigned as the regression targets. In particular, for visual ROIs, the regression target is set to the co-ordinates of the parent box. For textual objects, it is difficult to regress the ROI to match the entire span of the parent box. For example, in Fig. 4(a), for the ROI containing the word “Number” in the title, the ground-truth box would be the entire title spanning all the words (cyan box). To avoid this large difference from the proposal, we create the regression targets for “Number” by clipping the ground-truth box to have the same boundary as the proposed box along the horizontal direction. The task then is to grow the proposed ROI vertically and then later link it to its neighbour thereby creating the entire title box.

Lastly, for creating the ground-truth for linking, we assign a binary value to each ROI for each of the 4 directions. These 4 values indicate whether the ROI needs to be linked to its left, right, top, or bottom neighbours. In order to find the neighbours, we consider an area of 50×50 around a ROI and check if any of the neighbouring ROIs have the same parent box. If so, we assign 1 to the link corresponding to the direction (top, right, bottom, left) of that neighbour.

We trained our model for 10 epochs using Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.0001. We experimented with different loss functions for bounding box regression, and used cross-entropy loss for classification as well as link prediction.

Discussion

We now discuss the performance of different variants of our model as reported in Table 2. Note that all the variants in Table 2 use FPN as we always get better results with FPN. The variants mentioned in rows (a) and (b) use smooth ℓ_1 -loss (SL1) as the regression loss and do not have the Linking Head (LH) and the AN-ROI layer, respectively. The variants mentioned in rows (c) to (j) comprise the LH and the AN-ROI layer but only differ with respect to the regression loss. We make the following observations:

Ablation Studies: Comparing rows (a) and (b) with (c) of Table 2, we observe that (1) adding linking module allows us to handle longer textual elements (e.g., AP for plot-title improved from 0.00% to 99.44%), and (2) adding neighbourhood information using AN-ROI leads to a significant improvement in the performance (e.g., mAP@0.9 improved from 83.89% to 91.77%). Rows (d) to (i) suggest that when we use any of the IOU-based loss functions, the mAP@0.9 is

S.L	PlotNet Variants	IOU @0.9									@0.75	@0.5	
		bar	dot-line	legend label	legend preview	plot title	x-axis label	x-axis ticks	y-axis label	y-axis ticks	mAP	mAP	mAP
(a)	No LH	85.30	52.85	29.64	94.30	0.00	10.36	80.77	1.47	81.59	48.48	53.71	54.03
(b)	No AN-ROI	91.02	31.69	97.08	81.57	99.36	96.06	85.33	82.00	90.95	83.89	97.21	98.11
(c)	\mathcal{L}_{SL1}	92.16	61.18	98.38	93.46	99.44	97.21	94.21	95.45	94.42	91.77	97.74	98.24
(d)	\mathcal{L}_{1-IOU}	91.79	41.86	93.74	94.64	98.29	83.11	85.69	89.32	49.36	80.87	96.38	98.20
(e)	$\mathcal{L}_{-\log IOU}$	91.83	45.78	91.48	94.15	98.95	74.24	87.19	89.34	50.11	80.34	96.97	98.26
(f)	\mathcal{L}_{GIOU}	91.71	49.30	95.99	93.55	98.42	95.03	89.77	94.08	86.06	88.21	96.37	98.16
(g)	\mathcal{L}_{DIOU}	91.35	52.22	96.31	93.45	96.82	96.18	89.63	95.46	94.07	89.50	97.17	98.22
(h)	\mathcal{L}_{CIOU}	91.15	55.03	97.89	92.99	99.46	96.33	91.30	90.40	95.48	90.00	97.27	98.28
(i)	\mathcal{L}_{FIOU}	91.88	61.44	96.44	95.58	99.27	97.19	90.64	97.55	87.66	90.88	97.30	98.31
(j)	\mathcal{L}_{Custom}	92.80	70.11	98.47	96.33	99.52	97.31	94.29	97.66	94.48	93.44	97.93	98.32

Table 2: Comparison of variants of PlotNet on the PlotQA dataset with mAP scores (in %) at IOUs of 0.9, 0.75, and 0.5.

Models	Precision	Recall	F1-score
FRCNN (FPN+RA)	0.63	0.02	0.04
FrRCNN (RA)	0.62	0.12	0.20
FrRCNN (FPN+RA)	0.62	0.52	0.57
PlotNet - Ours	0.78	0.76	0.77

Table 3: Comparison of different models on the plot-to-table conversion task on the PlotQA dataset.

lower than what we obtain by using only the smooth ℓ_1 -loss (row (c)). However, among the IOU-based loss functions, FIOU gives the best performance. Further, using our custom loss (last row), we get the best performance with an overall mAP@0.9 of 93.44%.

Comparison to other models: In Fig. 5, we compare the mAP@0.9 and inference time of different models. We observe that PlotNet lies in the most favorable regime, *i.e.*, high mAP and low latency. In particular, PlotNet has the smallest inference time, beating one-stage detectors. Further, it improves upon its closest competitor (FrRCNN (FPN+RA)) by 16.22 absolute points in mAP. We also refer to Fig. 4(b) which shows that PlotNet detects accurate boxes. We note that this example is representative of the overall performance. Amongst individual classes, PlotNet majorly improves the accuracy on plot-titles which have long texts. We attribute this to combining simple region proposals with an explicit linking method. The improved accuracy of PlotNet on small objects like dot-line can be attributed to the additional neighbouring information present in each ROI feature.

Extrinsic Evaluation: Once the objects in a plot are accurately detected they can be used for inferences in a downstream task. For example, the data encoded in the plot can be represented as a structured table and then QA can be performed on this structured table. We follow the same procedure as outlined in (Methani et al. 2020) to construct a structured table from the objects identified in the plot. The quality of the generated structured table clearly depends on accurate localisation and classification of objects in the plot. The original PlotQA dataset also provides the gold standard structured tables associated with each plot. We can compute




Figure 5: mAP (in %) v/s Inference Time per image (in ms) for different object detection models on PlotQA at an IOU setting of 0.9. (x, y) represents the tuple (mAP, time).

the F1-score by comparing the tables generated after object detection with the ground truth tables. We report these numbers in Table 3. We observe that PlotNet improves the F1-score by a significant 20 points w.r.t its closest competitor - Faster R-CNN (FPN+RA). This signifies that the improved results with PlotNet enable automated reasoning with plots.

Conclusion

Existing object detection networks do not work for scientific plots - they have very low accuracy at the high IOU values required for reasoning over plots. Our proposed PlotNet makes a series of contributions across region proposal, model design, feature extraction, and loss function. These contributions together give a significant improvement of 16.22 points at mAP@0.9 IOU. PlotNet is also much faster with a 16x speedup in comparison with existing networks, including one-stage detectors. On the extrinsic challenging task of plot-to-table, PlotNet provides an improvement of 20 points in the F1 score. These significant results enable further exploration of automated reasoning over plots.

Acknowledgments

We are extremely thankful to Google for funding this work. Such extensive experimentation would not have been possible without their invaluable support.

References

- Berman, M.; Triki, A. R.; and Blaschko, M. B. 2018. The Lovász-Softmax Loss: A Tractable Surrogate for the Optimization of the Intersection-Over-Union Measure in Neural Networks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 4413–4421. IEEE Computer Society. doi:10.1109/CVPR.2018.00464.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Li, F. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, 248–255. doi:10.1109/CVPRW.2009.5206848. URL <https://doi.org/10.1109/CVPRW.2009.5206848>.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88(2): 303–338. ISSN 1573-1405. doi:10.1007/s11263-009-0275-4. URL <https://doi.org/10.1007/s11263-009-0275-4>.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. B. 2017. Mask R-CNN. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, 2980–2988. doi:10.1109/ICCV.2017.322. URL <https://doi.org/10.1109/ICCV.2017.322>.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, 770–778. doi:10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Kafle, K.; Price, B. L.; Cohen, S.; and Kanan, C. 2018. DVQA: Understanding Data Visualizations via Question Answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 5648–5656. doi:10.1109/CVPR.2018.00592.
- Kahou, S. E.; Michalski, V.; Atkinson, A.; Kádár, Á.; Trischler, A.; and Bengio, Y. 2018. FigureQA: An Annotated Figure Dataset for Visual Reasoning. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. URL <https://openreview.net/forum?id=H1mz0OyDz>.
- Kingma, D.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*.
- Lin, T.; Dollár, P.; Girshick, R. B.; He, K.; Hariharan, B.; and Belongie, S. J. 2017a. Feature Pyramid Networks for Object Detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 936–944. doi:10.1109/CVPR.2017.106. URL <https://doi.org/10.1109/CVPR.2017.106>.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017b. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, 2999–3007. doi:10.1109/ICCV.2017.324. URL <https://doi.org/10.1109/ICCV.2017.324>.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S. E.; Fu, C.; and Berg, A. C. 2016. SSD: Single Shot MultiBox Detector. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, 21–37. doi:10.1007/978-3-319-46448-0_2. URL https://doi.org/10.1007/978-3-319-46448-0_2.
- Methani, N.; Ganguly, P.; Khapra, M. M.; and Kumar, P. 2020. PlotQA: Reasoning over Scientific Plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Redmon, J.; and Farhadi, A. 2018. YOLOv3: An Incremental Improvement. *CoRR* abs/1804.02767. URL <http://arxiv.org/abs/1804.02767>.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, 91–99. URL <http://papers.nips.cc/paper/5638-faster-r-cnn-towards-real-time-object-detection-with-region-proposal-networks>.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I. D.; and Savarese, S. 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, 658–666. Computer Vision Foundation / IEEE. doi:10.1109/CVPR.2019.00075.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S. E.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, 1–9. IEEE Computer Society. doi:10.1109/CVPR.2015.7298594. URL <https://doi.org/10.1109/CVPR.2015.7298594>.
- Yu, J.; Jiang, Y.; Wang, Z.; Cao, Z.; and Huang, T. S. 2016. UnitBox: An Advanced Object Detection Network. In Hanjalic, A.; Snoek, C.; Worring, M.; Bulterman, D. C. A.; Huet, B.; Kelliher, A.; Kompatsiaris, Y.; and Li, J., eds., *Proceedings of the 2016 ACM Conference on Multimedia Conference, MM 2016, Amsterdam, The Netherlands, October 15-19, 2016*, 516–520. ACM. doi:10.1145/2964284.2967274.
- Zheng, Z.; Wang, P.; Liu, W.; Li, J.; Ye, R.; and Ren, D. 2020. Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The*

Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, 12993–13000. AAAI Press. URL <https://aaai.org/ojs/index.php/AAAI/article/view/6999>.