



Research article

Boundary distribution estimation for precise object detection

Peng Zhi, Haoran Zhou, Hang Huang, Rui Zhao, Rui Zhou* and Qingguo Zhou*

School of Information Science and Engineering, Lanzhou University, Lanzhou, China

* **Correspondence:** Email: zr@lzu.edu.cn, zhouqg@lzu.edu.cn; Tel: +869318912025; Fax: +869318912025.

Abstract: In the field of state-of-the-art object detection, the task of object localization is typically accomplished through a dedicated subnet that emphasizes bounding box regression. This subnet traditionally predicts the object's position by regressing the box's center position and scaling factors. Despite the widespread adoption of this approach, we have observed that the localization results often suffer from defects, leading to unsatisfactory detector performance. In this paper, we address the shortcomings of previous methods through theoretical analysis and experimental verification and present an innovative solution for precise object detection. Instead of solely focusing on the object's center and size, our approach enhances the accuracy of bounding box localization by refining the box edges based on the estimated distribution at the object's boundary. Experimental results demonstrate the potential and generalizability of our proposed method.

Keywords: object detection; deep learning; boundary estimation; box refinement

1. Introduction

As a combination of classification and localization, object detection is intended to spot desired objects in the image and categorize them respectively. In recent years, the field of object classification has made significant advancements due to the progress in deep learning frameworks. Classifiers [1–4] now demonstrate impressive performance on diverse and difficult benchmarks [5–8]. However, the aspect of localization, which involves estimating the precise position and size of objects, still lags behind and poses limitations on the overall detection performance.

In state-of-the-art object detectors such as faster region-based convolutional neural network (R-CNN) [9], RetinaNet [10] and CenterNet [11], the task of object localization is carried out by a box subnet that focuses on bounding box regression. This subnet typically predicts bounding boxes by regressing the box's center position (x, y) and scaling factors (w, h) . While this design has shown effective results in conventional frameworks, it suffers from inherent flaws when it comes to precise

localization. As illustrated in Figure 1, each variable of (x, y, w, h) impacts multiple edges when adjusting the box's center position or size. Consequently, all edges of the box shift together, which compromises the accuracy of localization. Moreover, the center-focus representation method introduces a feature imbalance, wherein detectors disproportionately emphasize internal features that are less crucial for localization.

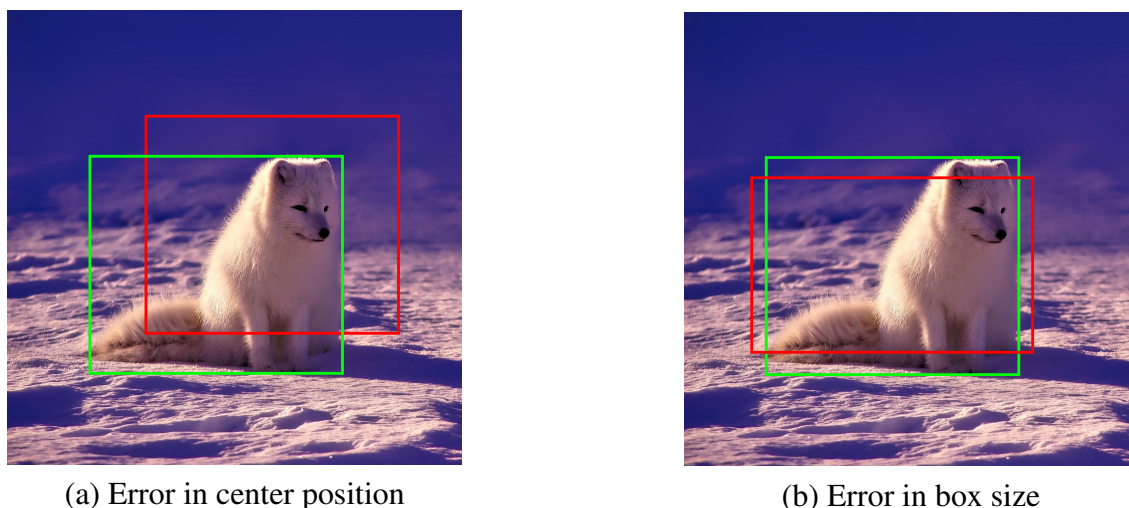


Figure 1. Illustration of the center-size box tuning method: the green box represents the ground truth, while the red box represents the predicted box with either center error or size error.

The aim of achieving precise localization is to minimize the disparity between predictions and the ground truth. When humans annotate ground truth boxes, they typically align each side of the box individually with the corresponding boundary, which is the most efficient approach. Taking inspiration from this manual annotation method, we introduce a novel solution for precise localization by refining previous localization results using estimations of the actual boundary distribution. Our proposed method follows a two-step process. As shown in Figure 2, it first predicts the boundary within the proposed area of the existing pipeline. Next, the coarse boundary is determined based on the boundary probability map. For each side of the box, we estimate the corresponding fine boundary distribution by considering the context of the coarse boundary. This fine boundary distribution is then utilized to refine the final localization results.

We apply our proposed method to various frameworks and assess their performance on the common objects in context (COCO) test-dev dataset [5]. Our method demonstrates significant potential through straightforward functional estimation. Specifically, we achieved an improvement of 2.0% in average precision (AP) without incurring additional computational costs or requiring additional annotation data, as compared to the competent Mask R-CNN [12] baseline.

Our main contributions can be summarized as follows:

- We propose a universal method for precise object detection that enhances the accuracy of bounding box localization by refining the box edges based on estimations of the object's boundary distribution. This novel approach improves the overall detection performance.
- In order to accurately generate each side of the bounding box, we evaluated various representation

methods and conclude that the edge representation method is better suited for precise localization. This insight further enhances the effectiveness of our method.

- To leverage boundary features effectively, we introduce coarse-to-fine boundary distribution estimation modules. These modules demonstrate a significant improvement over previous cascade architectures, leading to more accurate and reliable object detection results.

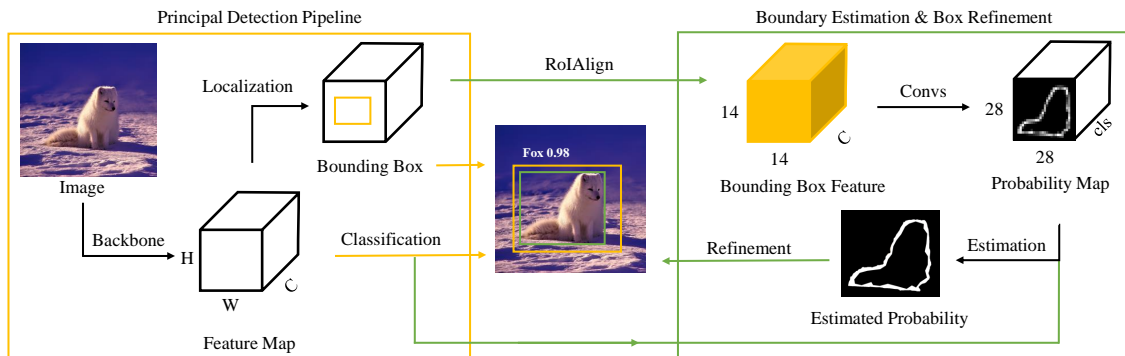


Figure 2. Illustration of principal detection pipelines with boundary estimation and box refinement.

2. Related works

2.1. Localization in detection

Due to the remarkable advancements in deep learning, convolutional neural network (CNN) based methods have emerged as the dominant solutions in detection applications. These deep learning detectors can be categorized into three main types based on their structures: two-stage detectors [9, 12–14], one-stage detectors [15–18] and anchor-free detectors [11, 19, 20].

While different in their implementation, all of these methods adhere to the classification-localization paradigm. At its core, representing a bounding box requires four independent variables. In anchor-based detectors like Faster R-CNN [9] and RetinaNet [10], the bounding box is regressed from proposals or predefined anchors and represented as offsets of the center ($\delta x, \delta y$) and relative scaling factors ($\delta w, \delta h$). Similarly, anchor-free methods such as fully convolutional one-stage (FCOS) [20] and CenterNet [11] predict the object's center and its corresponding size. These center-size box tuning methods effectively locate objects but are still prone to inherent errors. To address this issue, CornerNet [21] introduces a novel approach by utilizing the object's top-left corner and bottom-right corner to form the bounding box. However, this method introduces additional computational costs due to corner pair grouping. In contrast, our method combines the strengths of these two approaches by refining the edges of previous localization results. By doing so, we can effectively enhance the accuracy of object detection without incurring the computational overhead associated with grouping corner pairs.

2.2. Boundary in computer vision

The semantic importance of edges and boundaries has long been recognized in various tasks at different levels. In Figure 3, we present two representative applications to highlight the similarities and differences between object detection and these tasks.

Edge detection, typically considered a low-level task, focuses on extracting visually prominent edges and object boundaries from natural images. It can be categorized into three main approaches: early filter methods [22–24], information theory methods [25,26] and learning-based methods [27,28]. Recent CNN-based methods [29–31] aim to automatically learn hierarchical features, but they still face challenges in accurately distinguishing object boundaries from detection results. Utilizing low-level cues, such as edges, to perform object-level tasks proves to be a challenging endeavor.

Instance segmentation aims to provide dense inference by generating pixel-wise masks for each instance of objects belonging to the same class. Some instance segmentation methods [12,32,33] build upon existing detectors. They first utilize the detector to obtain object proposals and subsequently predict masks for these regions. However, these methods typically rely on pixel-level annotations, which are often unavailable in the context of object detection. While object detection shares some similarities with these tasks, such as the utilization of edges and boundaries, it also presents distinct challenges and requirements, such as precise localization, classification and the handling of multiple instances of objects. Overcoming these challenges often requires tailored approaches and adaptations specific to object detection scenarios.

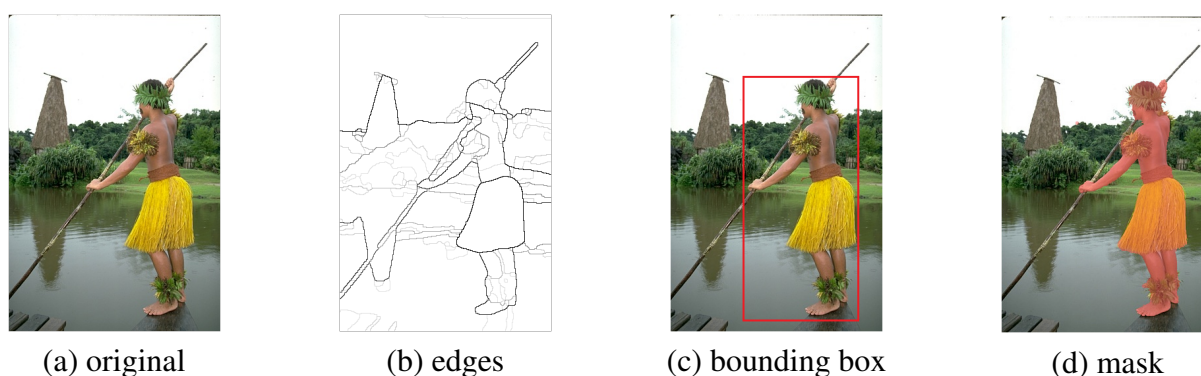


Figure 3. Illustration of some boundary-related applications in computer vision (edge detection, object detection, and instance segmentation).

2.3. Detection with boundary features

Boundary context, including the surroundings and shapes of objects, contains valuable information that can benefit detection applications. Traditional detectors, such as Sobel [24], Canny [22], and Marr-Hildreth [23], typically rely on smoothing filters or image gradients to extract local features at a low level. Edge boxes [34] introduces a proposed algorithm based on edges in the input image, which works well for single objects but struggles to accurately separate multiple objects. These image-based detectors have limitations as they primarily consider the low-level visual context of objects and are challenging to apply to modern detectors. To address these limitations, recent deep learning-based

methods have introduced novel techniques to incorporate context information. For example, the deformable R-FCN [35] utilizes deformable convolution to include context information in object detection. The BAN [36] includes a boundary-aware network that leverages boundary context for improved performance. Side-aware boundary localization (SABL) [37] aims to predict precise object boundaries using bucket-based approaches. Inspired by these deep learning-based methods, we referred to them and developed our estimation module to effectively leverage the boundary context of objects in our detection approach. By doing so, we enhance the overall accuracy and performance of object detection.

3. Method

In this research, our primary focus is to enhance the localization module in modern object detectors. Taking inspiration from the manual annotation method, we introduce an edge-focused box representation as a key component for achieving precise localization. Furthermore, we devise a novel boundary estimation method that effectively leverages boundary features. The overall pipeline of our proposed method is depicted in Figure 2.

Our approach aligns with the fundamental detection pipeline, which considers object detection as a combination of localization and classification. Initially, features corresponding to the object are cropped and resized based on the previous localization results. These features are then utilized to predict the coarse boundary, which is subsequently refined to estimate the fine boundary. Each side of the bounding box is refined using the corresponding estimated fine boundary. Importantly, our method is not dependent on a specific detection pipeline and does not require any additional annotations. By employing this approach, we significantly improve the localization accuracy in object detection, without the need for extra annotations or modifications to the existing detection pipeline.

3.1. Bounding box representation

The objective of precise localization is to minimize the discrepancy between predictions and the ground truth. In manual annotation, the most efficient approach is to align each side of the box with the object boundary individually. However, conventional detectors typically choose to predict the centers and sizes of bounding boxes instead of adjusting the edges directly.

As depicted in Figure 1, the center-size tuning method inherently faces challenges when attempting to adjust individual edges. Let us consider a scenario in which each side of the box can be represented by its respective vertical and horizontal lines. The vertical and horizontal edges of the rectangular box (E_l, E_r, E_t, E_b) can be represented by a set of scalars (l, r, t, b), or alternatively, by the box center and size (x, y, w, h):

$$E_l = f_l(l) = g_l(x, w/2) \quad (3.1)$$

$$E_r = f_r(r) = g_r(x, w/2) \quad (3.2)$$

$$E_t = f_t(t) = g_t(y, h/2) \quad (3.3)$$

$$E_b = f_b(b) = g_b(y, h/2) \quad (3.4)$$

Each edge is influenced by both the center and size of the box. Consequently, adjusting a single edge becomes more difficult when employing the center-size method. This dilemma is prevalent in

modern detectors, whether they predict the center and size directly or use the offset from the proposal relative to the ground truth box. A similar notion for precise localization is also reflected in the corner representation method [21]. The top-left corner and bottom-right corner can be viewed as the intersection points of the corresponding edges, denoted as $C_{tl}(l, t)$ and $C_{br}(r, b)$ respectively.

3.2. Coarse boundary localization

As depicted in Figures 2 and 4, our method can be broken down into a two-step scheme: coarse boundary localization and fine boundary estimation. Initially, we utilize the region of interest align (RoIAlign) [12] operation to extract and resize features within the proposal area. The resulting bounding box features, with dimensions of $14 \times 14 \times C$ (where $C = 256$ in our experiments), are subjected to four 3×3 convolutional layers with a stride of 1, followed by a 2×2 deconvolutional layer with a stride of 2, and finally a 1×1 convolutional layer to produce the $28 \times 28 \times cls$ (where $cls = 80$ in the COCO dataset [5]) boundary probability map. In conjunction with the classification result obtained from the previous detector, we extract a single-channel probability map for a specific class. Subsequently, we compute the maximum value along each row and column, resulting in two compressed vectors. These vectors serve as inputs to the coarse boundary localization module.

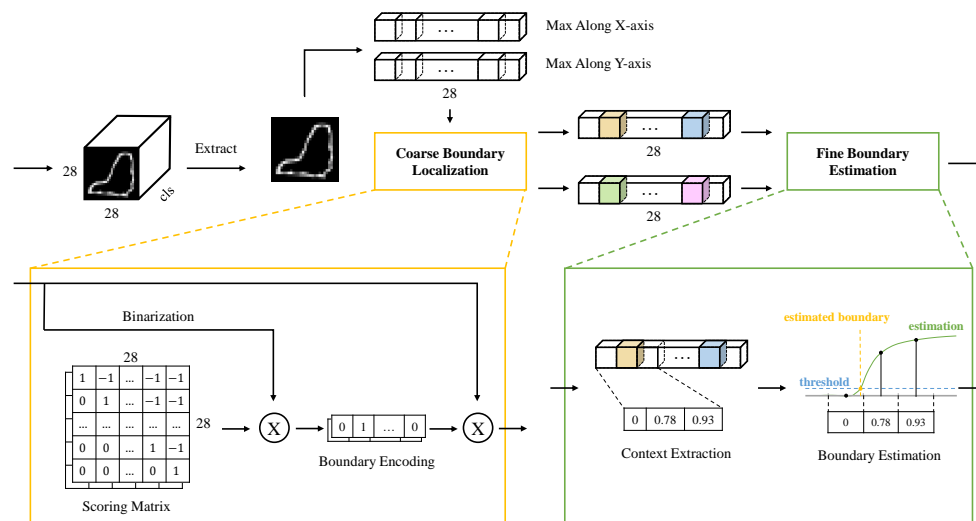


Figure 4. Illustration of part of the boundary estimation and box refinement module.

Within the coarse boundary localization module, we first binarize the vector using a threshold (in our case, the threshold is set to 1×10^{-4}) and multiply it by a scoring matrix and its transposed matrix. The scoring matrix, denoted as M_s in Figure 4, undergoes matrix multiplication with the binarized vector v_b according to the following equation:

$$\vec{v}_b \times M_s = [\vec{v}_b m_{c1}, \dots, \vec{v}_b m_{c28}] \quad (3.5)$$

And for each $\vec{v}_b \vec{m}_{ci}$, it follows that

$$\vec{v}_b \vec{m}_{ci} = \begin{cases} 0, & i < i_{fp} \\ 1, & i = i_{fp} \\ 1 - k', & i_{fp} < i \leq i_{lp} \\ -k, & i > i_{lp} \end{cases} \quad (3.6)$$

i_{fp} and i_{lp} represent the indices of the first and last positive elements in the vector, respectively. Additionally, k and k' are integers that satisfy $k, k' \in [0, 28]$. These matrices are designed to identify the approximate positions of the left/top and right/bottom boundaries within the probability map. The resulting product is then activated using the rectified linear unit (ReLU) function.

3.3. Fine boundary estimation

As shown in Figure 4, the operation of the fine boundary estimation module begins by extracting the contextual information of the coarse boundary. The estimated fine boundary is then used to refine the bounding box. Assuming that the pixel values of the image are uniformly sampled from a continuous distribution, and that the variations near the object boundary can be described by an elementary function, we derive the precise boundary points based on this assumption and the pixel values near the object boundary. Our estimation method differs from the conventional interpolation approaches because it focuses on the distribution of the boundary rather than estimating values between known points. This allows us to select a relatively low threshold (in our experiments, we used a threshold of 0) to minimize the outward expansion of the boundary during the interpolation process.

For better understanding, let us assume that the one-dimensional vector obtained from the coarse-boundary localization process is denoted as V^{lr} . By considering V^{lr} as a uniform sampling of a continuous distribution, with sampling points located at the pixel centers, we can transform the discrete V^{lr} into a continuous representation. Assuming that the distribution near the left boundary of the object can be approximated by an elementary function $f(x)$, we utilize limit calculations and indices to demonstrate that the elementary function $f(x)$ used to approximate the distribution near the boundary needs to pass through the coordinates $(0, 0)$ and $(1, 1)$ in the Cartesian coordinate system. The calculation process for the right boundary and the upper and lower boundaries of the target box follows a similar procedure, akin to the calculation based on V^{lr} . Consequently, several key requirements must be fulfilled during the process. First, the boundary distribution function $f_B : [0, 1] \rightarrow [0, 1]$ should be a monotonically increasing and differentiable function. Additionally, the gradient of the function should be smooth to alleviate numerical instability issues. For instance, the function $f(x) = \sqrt{x}$ is unsuitable as a distribution function because its derivative $f'(x)$ tends to infinity as x approaches 0, rendering it untrainable. In this study, we experimented with three elementary functions that meet the aforementioned requirements to fit the distribution of objects near the boundary, as detailed in Section 4.2. Ultimately, a linear function was chosen as the fitting solution for this research.

4. Experiments

4.1. Evaluation

We present the evaluation results for our proposed estimation method, which were obtained by training and evaluating the networks on the COCO dataset [5]. To ensure fairness and consistency, we adhere to the universal splitting method for the training, validation and testing of our model on the COCO dataset. The detection results were evaluated using the standard COCO-style AP metric.

In our experiments, we chose Mask R-CNN [12] as the baseline model. The parameters of the mask head in Mask R-CNN were used to initialize the network described in Section 3.2. This is because the mask head has an inherent advantage in predicting the boundary probability map.

We observed that most elements inside the object are similar (close to 1), which renders them insignificant for estimation purposes. Consequently, we employed a functional estimation approach, denoted as $f(x)$, with the value x at the coarse boundary serving as the variable.

The detection results are presented in Figure 5, where the blue box represents the baseline method and the red box represents our method. As shown in Table 1 and Figure 6, our method demonstrates improvements in overall AP, individual AP metrics and category-wise AP metrics as compared to the ResNet-50-FPN baseline with linear estimation $f(x) = x$. Notably, our method exhibits higher precision than the baseline even when both results are classified as positive examples. It is also noteworthy that our method outperforms the baseline when using a deeper ResNet-101-FPN backbone, despite the increased computational cost associated with this backbone.

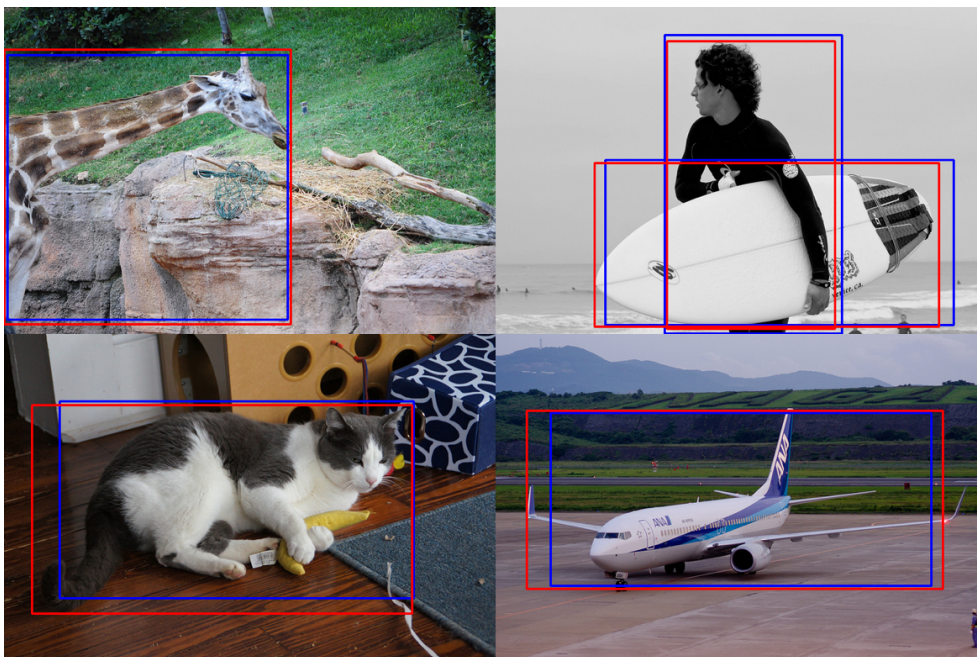


Figure 5. Detection results on the COCO test-dev dataset, with the baseline model shown in blue and our model shown in red.

each measurement based on the processing time for 100 images of resolution 1024×1024 . Mask R-CNN achieved an average FPS of 6.00, while our model achieved a slightly lower average FPS of 5.97, indicating a marginal 0.5% reduction in FPS compared to Mask R-CNN. Therefore, the impact on speed can be considered negligible.

4.2. Additional experiments and analysis

Comparison between different estimation functions. In order to investigate the impact of different boundary distribution functions, we conducted comparative experiments with three cases (in Table 2): *linear*, *exponential* and *logarithmic*, represented by the functions $f(x) = x$, $f(x) = x^2$, and $f(x) = \ln((e - 1)x + 1)$, respectively. From the experiment, we observed that the linear method achieved better overall performance, particularly on large objects. The convex logarithmic function performed better on smaller objects, exhibiting a higher AP_{50} and AP_{75} . On the other hand, the concave exponential function did not demonstrate significant advantages.

Table 2. Comparison between different estimation functions on COCO val dataset.

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Baseline	37.5	58.6	40.8	21.8	41.0	48.9
W/ lin	39.5	59.0	42.4	22.4	43.0	52.0
W/ exp	39.3	58.9	41.9	22.4	42.9	51.7
W/ log	39.4	59.2	42.6	22.8	43.0	51.9

Influence of extra mask annotation. Although our method does not use annotation data during training, the initialization of the boundary prediction network with parameters from Mask R-CNN [12] introduces additional mask information. To assess the impact of mask annotation, we replaced the parameters in the boundary prediction network with randomly initialized parameters. As shown in Table 3, our method remains effective, with an overall AP improvement of 1.8%. These results indicate that extra annotation, such as masks, can provide benefits but are not essential for the proposed method.

Table 3. Ablation study on extra mask annotation.

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Baseline	37.5	58.6	40.8	21.8	41.0	48.9
W/o mask	38.5	58.3	41.1	21.6	42.3	50.5
W/ mask	39.5	59.0	42.4	22.4	43.0	52.0

Comparison with extra regression. Cascading methods have demonstrated competence in precise localization. To highlight the superiority of our boundary estimation method, we compare it to baseline methods with an additional regression stage. As shown in Table 4, our approach consistently outperforms the Mask R-CNN baselines, leading by up to 0.8% in overall AP. These results validate the effectiveness of our method in achieving precise localization.

Table 4. Comparison between extra regression method and our edge refinement method.

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Baseline	37.5	58.6	40.8	21.8	41.0	48.9
W/ reg	38.7	58.1	41.5	21.9	42.2	51.2
W/ ref	39.5	59.0	42.4	22.4	43.0	52.0

5. Conclusions

In this paper, we propose a novel approach to achieve precise object detection by estimating object boundary distributions. Our primary goal is to improve detection accuracy. Drawing inspiration from manual annotation techniques, we harness the advantages of the edge-focus box representation method. By predicting the coarse boundary and refining it through fine boundary estimation, our method significantly enhances the accuracy of bounding box edges based on previous localization results. Extensive experiments conducted on the COCO dataset demonstrate the potential of our boundary estimation method across various detection pipelines, eliminating the need for additional annotation. In future research, we aim to extend the application of our boundary detection procedure to diverse domains, including medical imaging. This extended application holds significant value, particularly in optical coherence tomography imaging, where our procedure enables the precise detection of boundaries in retina tissue layers with high accuracy. This advancement has the potential to facilitate the accurate diagnosis of ocular or neuropathological disorders [41, 42]. Furthermore, we plan to explore the utilization of a lighter ResNet18 backbone or a more complex ResNet101 backbone to address the specific requirements of speed and accuracy in different scenarios.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

This work was partially supported by the National Key R&D Program of China under Grant No. 2020YFC0832500, Gansu Province Science and Technology Major Project - Industrial Project under Grant No. 22ZD6GA048, Gansu Province Key Research and Development Plan - Industrial Project under Grant No. 22YF7GA004, Fundamental Research Funds for the Central Universities under Grant Nos. lzujbky-2022-kb12, lzujbky-2021-sp43, lzujbky-2020-sp02, lzujbky-2019-kb51 and lzujbky-2018-k12, National Natural Science Foundation of China under Grant No. 61402210, Science and Technology Plan of Qinghai Province under Grant No.2020-GX-164 and Supercomputing Center of Lanzhou University. We also gratefully acknowledge the support of NVIDIA Corporation for the donation of the Jetson TX1 used for this research. This work was also partially supported by the Gansu Provincial Science and Technology Major Special Innovation Consortium Project (21ZD3GA002), from the Gansu Province Green and Smart Highway Transportation Innovation Consortium as part of the Gansu Province Green and Smart Highway Key Technology Research and Demonstration. We would like to express our gratitude to our co-authors, Mr. Hang Huang and Mr.

Haoran Zhou, for their dedicated efforts during their postgraduate studies [43]. We acknowledge their hard work and appreciate their significant contributions to the research.

References

1. R. Kaur, S. Singh, A comprehensive review of object detection with deep learning, *Digital Signal Process.*, **132** (2023), 103812. <https://doi.org/10.1016/j.dsp.2022.103812>
2. P. Jiang, D. Ergu, F. Liu, Y. Cai, B. Ma, A Review of Yolo algorithm developments, *Proc. Comput. Sci.*, **199** (2022), 1066–1073. <https://doi.org/10.1016/j.procs.2022.01.135>
3. W. Liu, G. Wu, F. Ren, X. Kang, DFF-ResNet: An insect pest recognition model based on residual networks, *Big Data Min. Anal.*, **3** (2020), 300–310. <https://doi.org/10.26599/BDMA.2020.9020021>
4. A. Mughees, L. Tao, Multiple deep-belief-network-based spectral-spatial classification of hyperspectral images, *Tsinghua Sci. Technol.*, **24** (2019), 183–194. <https://doi.org/10.26599/TST.2018.9010043>
5. T. Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, et al., Microsoft COCO: Common objects in context, in *European Conference on Computer Vision*, (2014), 740–755. https://doi.org/10.1007/978-3-319-10602-1_48
6. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, et al., ImageNet large scale visual recognition challenge, *Int. J. Comput. Vis.*, **115** (2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
7. Y. Fan, D. Ni, H. Ma, HyperDB: a hyperspectral land class database designed for an image processing system, *Tsinghua Sci. Technol.*, **22** (2017), 112–118. <https://doi.org/10.1109/TST.2017.7830901>
8. M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, A. Zisserman, The PASCAL visual object classes challenge: A retrospective, *Int. J. Comput. Vis.*, **111** (2015), 98–136. <https://doi.org/10.1007/s11263-014-0733-5>
9. S. Ren, K. He, R. Girshick, J. Sun, Faster R-CNN: Towards real-time object detection with region proposal networks, *IEEE Trans. Pattern Anal. Mach. Intell.*, **39** (2017), 1137–1149. <https://doi.org/10.1109/TPAMI.2016.2577031>
10. T. Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in *2017 IEEE International Conference on Computer Vision (ICCV)*, (2017), 2999–3007. <https://doi.org/10.1109/ICCV.2017.324>
11. X. Zhou, D. Wang, P. Krähenbühl, Objects as points, preprint, arXiv:1904.07850.
12. K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask R-CNN, *IEEE Trans. Pattern Anal. Mach. Intell.*, **42** (2020), 386–397. <https://doi.org/10.1109/TPAMI.2018.2844175>
13. M. Chen, F. Bai, Z. Gerile, Special object detection based on Mask RCNN, in *2021 17th International Conference on Computational Intelligence and Security (CIS)*, (2021), 128–132. <https://doi.org/10.1109/CIS54983.2021.00035>
14. Z. Ou, Z. Wang, F. Xiao, B. Xiong, H. Zhang, M. Song, et al., AD-RCNN: Adaptive dynamic neural network for small object detection, *IEEE Int. Things J.*, **10** (2023), 4226–4238. <https://doi.org/10.1109/JIOT.2022.3215469>

15. L. Yang, Y. Xu, S. Wang, C. Yang, Z. Zhang, B. Li, et al., PDNet: Toward better one-stage object detection with prediction decoupling, *IEEE Trans. Image Process.*, **31** (2022), 5121–5133. <https://doi.org/10.1109/TIP.2022.3193223>
16. J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2016), 770–778. <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.91>
17. W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Y. Fu, et al., SSD: Single shot multiBox detector, in *European Conference on Computer Vision*, (2016), 21–37. https://doi.org/10.1007/978-3-319-46448-0_2
18. J. Redmon, A. Farhadi, YOLOv3: An incremental improvement, preprint, arXiv:1804.02767.
19. G. Wang, J. Wu, B. Tian, S. Teng, L. Chen, D. Cao, et al., CenterNet3D: An anchor free object detector for point cloud, *IEEE Trans. Intell. Transp. Syst.*, **23** (2022), 12953–12965. <https://doi.org/10.1109/TITS.2021.3118698>
20. Z. Tian, C. Shen, H. Chen, T. He, FCOS: Fully convolutional one-stage object detection, in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, (2019), 9626–9635. <https://doi.ieeecomputersociety.org/10.1109/ICCV.2019.00972>
21. H. Law, J. Deng, CornerNet: Detecting objects as paired keypoints, *Int. J. Comput. Vis.*, **128** (2020), 642–656. <https://doi.org/10.1007/s11263-019-01204-1>
22. J. Canny, A computational approach to edge detection, *IEEE Trans. Pattern Anal. Mach. Intell.*, **8** (1986), 679–698. <https://doi.org/10.1109/TPAMI.1986.4767851>
23. D. Marr, E. Hildreth, Theory of edge detection, *Proc. R. Soc. Lond. B*, **207** (1980), 187–217. <https://doi.org/10.1098/rspb.1980.0020>
24. J. Kittler, On the accuracy of the Sobel edge detector, *Image Vis. Comput.*, **1** (1983), 37–42. [https://doi.org/10.1016/0262-8856\(83\)90006-9](https://doi.org/10.1016/0262-8856(83)90006-9)
25. D. R. Martin, C. C. Fowlkes, J. Malik, Learning to detect natural image boundaries using local brightness, color, and texture cues, *IEEE Trans. Pattern Anal. Mach. Intell.*, **26** (2004), 530–549. <https://doi.org/10.1109/TPAMI.2004.1273918>
26. P. Arbeláez, M. Maire, C. Fowlkes, J. Malik, Contour detection and hierarchical image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.*, **33** (2011), 898–916. <https://doi.org/10.1109/TPAMI.2010.161>
27. J. J. Lim, C. L. Zitnick, P. Dollár, Sketch tokens: A learned mid-level representation for contour and object detection, in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, (2013), 3158–3165. <https://doi.org/10.1109/CVPR.2013.406>
28. P. Dollár, C. L. Zitnick, Structured forests for fast edge detection, in *2013 IEEE International Conference on Computer Vision*, (2013), 1841–1848. <https://doi.org/10.1109/ICCV.2013.231>
29. S. Xie, Z. Tu, Holistically-nested edge detection, in *2015 IEEE International Conference on Computer Vision (ICCV)*, (2015), 1395–1403. <https://doi.org/10.1109/ICCV.2015.164>
30. G. Bertasius, J. Shi, L. Torresani, DeepEdge: A multi-scale bifurcated deep network for top-down contour detection, in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2015), 4380–4389. <https://doi.org/10.1109/CVPR.2015.7299067>
31. W. Shen, X. Wang, Y. Wang, X. Bai, Z. Zhang, DeepContour: A deep convolutional feature learned by positive-sharing loss for contour detection, in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2015), 3982–3991. <https://doi.org/10.1109/CVPR.2015.7299024>

32. Z. Huang, L. Huang, Y. Gong, C. Huang, X. Wang, Mask scoring R-CNN, in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (2019), 6402–6411. <https://doi.org/10.1109/CVPR.2019.00657>
33. K. Chen, J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, et al., Hybrid task cascade for instance segmentation, in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (2019), 4974–4983. <https://doi.org/10.1109/CVPR.2019.00511>
34. C. L. Zitnick, P. Dollár, Edge boxes: Locating object proposals from edges, in *European Conference on Computer Vision*, (2014), 391–405. https://doi.org/10.1007/978-3-319-10602-1_26
35. J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, et al., Deformable convolutional networks, in *2017 IEEE International Conference on Computer Vision (ICCV)*, (2017), 764–773. <https://doi.org/10.1109/ICCV.2017.89>
36. Y. Kim, T. Kim, B. N. Kang, J. Kim, D. Kim, BAN: Focusing on boundary context for object detection, in *Asian Conference on Computer Vision*, (2018), 555–570. https://doi.org/10.1007/978-3-030-20876-9_35
37. J. Wang, W. Zhang, Y. Cao, K. Chen, J. Pang, T. Gong, et al., Side-aware boundary localization for more precise object detection, in *European Conference on Computer Vision*, (2020), 403–419. https://doi.org/10.1007/978-3-030-58548-8_24
38. C. Y. Fu, W. Liu, A. Ranga, A. Tyagi, A. C. Berg, DSSD: Deconvolutional single shot detector, preprint, arXiv:1701.06659.
39. R. Araki, T. Onishi, T. Hirakawa, T. Yamashita, H. Fujiyoshi, MT-DSSD: Deconvolutional single shot detector using multi task learning for object detection, segmentation, and grasping detection, in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, (2020), 10487–10493. <https://doi.org/10.1109/ICRA40945.2020.9197251>
40. C. Y. Fu, M. Shvets, A. C. Berg, RetinaMask: Learning to predict masks improves state-of-the-art single-shot detection for free, preprint, arXiv:1901.03353.
41. R. K. Meleppat, M. V. Matham, L. K. Seah, Optical frequency domain imaging with a rapidly swept laser in the 1300nm bio-imaging window, in *International Conference on Optical and Photonic Engineering*, (2015), 721–729. <https://doi.org/10.1117/12.2190530>
42. R. K. Meleppat, C. R. Fortenbach, Y. Jian, E. S. Martinez, K. Wagner, B. S. Modjtahedi, et al., In Vivo Imaging of Retinal and Choroidal Morphology and Vascular Plexuses of Vertebrates Using Swept-Source Optical Coherence Tomography, *Transl. Vis. Sci. Technol.*, **11** (2022), 11. <https://doi.org/10.1167/tvst.11.8.11>
43. H. Huang, *Research on Object Detection Based on Improved MASK R-CNN*, Master's degree, Lanzhou University in Lanzhou, 2021. <https://doi.org/10.27204/d.cnki.glzhu.2021.001818>



AIMS Press

©2023 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)