# **Real-Time Indoor Object Detection Based on Hybrid** CNN-Transformer Approach

## Salah-eddine LAIDOUDI<sup>a,b,\*</sup>, Madjid MAIDI<sup>b</sup> and Samir OTMANE<sup>a</sup>

<sup>a</sup>IBISC, Université Paris-Saclay, Univ Evry, Evry-Courcouronnes 91020, France <sup>b</sup>ESME Research Lab, Ivry-sur-Seine, 94200, France

Abstract. Real-time object detection in indoor settings is a challenging area of computer vision, faced with unique obstacles such as variable lighting and complex backgrounds. This field holds significant potential to revolutionize applications like augmented and mixed realities by enabling more seamless interactions between digital content and the physical world. However, the scarcity of research specifically fitted to the intricacies of indoor environments has highlighted a clear gap in the literature. To address this, our study delves into the evaluation of existing datasets and computational models, leading to the creation of a refined dataset. This new dataset is derived from OpenImages v7[14], focusing exclusively on 32 indoor categories selected for their relevance to real-world applications. Alongside this, we present an adaptation of a CNN detection model. incorporating an attention mechanism to enhance the model's ability to discern and prioritize critical features within cluttered indoor scenes. Our findings demonstrate that this approach is not just competitive to existing state-of-the-art models in accuracy and speed but also opens new avenues for research and application in the field of real-time indoor object detection.

## 1 Introduction

Object detection, a cornerstone of computer vision, has experienced transformative growth with the advent of deep learning technologies. Traditional techniques such as the Viola-Jones detector[31] and Histogram of Oriented Gradients (HOG)[5] laid the early groundwork by enabling systems to recognize objects through feature detection and machine learning classifiers.

However, these methods struggled with high variability in object appearances and were generally limited in their ability to scale with complexity and diversity of input data. The deep learning era introduced a paradigm shift with the development of convolutional neural networks (CNNs), significantly enhancing the ability and efficiency of object detection systems. This evolution continued with the introduction of advanced architectures, including single-stage detectors like YOLO (You Only Look Once)[25][23][24][2][29][18][32][13][33][8] and SSD (Single Shot MultiBox Detector)[21], and two-stage detectors such as R-CNN[10] and its variants[9][26], which improved detection accuracy by refining proposals through a secondary classification step.

More recently, the integration of transformer-based models[30], which employ self-attention mechanisms to capture global dependencies within the image data, has started to set new benchmarks in the field.

Despite these advancements, real-time object detection in indoor environments remains fraught with challenges. Indoor settings are characterized by varying lighting conditions—from natural light flooding a room to multiple artificial light sources—which can drastically affect the visibility and appearance of objects. Furthermore, object occlusions, where items are partially or fully hidden behind others, add an additional layer of complexity. Lastly, the demand for low latency in applications such as augmented reality (AR) and mixed reality (MR) means that the detection system not only needs to be accurate but also exceedingly fast. These conditions present unique hurdles that are not fully addressed by current object detection models.

The primary goal of our research is to develop an object detection system optimized for real-time application in indoor settings. This system is intended to significantly enhance user interactions in multiple fields such as augmented reality (AR) and mixed reality (MR), where seamless integration of digital content with the physical world is what makes these experiences more immersive.

Our work introduces a novel hybrid architecture that combines the robustness of CNNs with the sophisticated spatial reasoning capabilities of transformers. This approach is designed to be lightweight, catering to the needs of real-time processing without sacrificing accuracy. By employing this hybrid model, our system not only addresses the typical challenges associated with indoor object detection—such as variable lighting and occlusions—but also shows promising results that are competitive with, and in some cases superior to, current state-of-the-art technologies. Our contribution is a significant step forward in computer vision, pushing the boundaries of what is achievable in real-time indoor object detection.

## 2 Literature review

Object detection has significantly evolved over the past decade due to advances in neural networks and machine learning. Despite this progress, research specifically targeting indoor environments is limited. Typically, object detectors are evaluated using the COCO dataset[19], which is comprehensive but does not fully address the challenges of indoor settings, such as complex lighting, varied object placements, and diverse backgrounds.

As we previously mentioned on the introduction, single-stage detectors have garnered attention for their efficiency, with a particular emphasis on speed, making them ideal candidates for real-time ap-

<sup>\*</sup> Corresponding Author. Email: salah-eddine1.laidoudi@esme.fr



Figure 1. Most common object detection methods throughout the time

plications. Within this category, we observe two primary types of single-stage detectors:

Anchor-Based Detectors: These detectors, exemplified by earlier versions of YOLO, SSD, and RetinaNet ..., rely on predefined bounding boxes, known as anchors, which serve as priors to guide the detection process. They operate by pre-establishing a variety of bounding boxes with different widths and heights tailored to the common aspect ratios observed for each class type. During detection, these anchors are tiled across the image, and the model predicts the likelihood of an object being present within these preset tiles.

**Anchor-Free Detectors:** The recent shift towards anchor-free detectors signifies a promising trend in the field of real-time object detection, with notable examples including CornerNet[17], CenterNet[7], FCOS[28], and the latest iterations of YOLO (v5-v9). These modern detectors abstain from the reliance on prior bounding box knowledge. Instead, they directly predict critical points—specifically the top-left and bottom-right corners of the bounding box, thereby streamlining the detection process. This direct prediction approach circumvents the computational burden associated with anchor manipulation, leading to a more efficient detection framework suitable for rapid deployment in real-world scenarios.

Our research primarily aligns with the progressive trend of anchorfree detectors, we adopt and further this approach, augmenting it with the capabilities of a transformer.

Several noteworthy efforts in adjacent areas include innovations in low-textured object detection[15], where lighter SSD architectures have been proposed[16]. For instance, a notable study suggested modifications to the SSD framework to better handle objects with minimal textural information, enhancing detection in specific but limited scenarios. Furthermore, the RT-DETR[39] model stands out as a pioneering real-time detection transformer, integrating the rapid inference capabilities necessary for real-time applications with the advanced spatial reasoning of transformers.

Additionally, the YOLO series has significantly impacted the field with its high-speed processing and increasingly accurate detection rates. These models have consistently pushed the boundaries of what is achievable in real-time object detection, setting high benchmarks for both speed and accuracy. However, these systems are often calibrated and tested against datasets that do not adequately mimic the challenges specific to indoor environments.

In this last part of the literature review we'll examine some of the models from our comparative study more closely.

**Yolov5**[29], developed by Glen Jocher of Ultralytics in 2020 shortly after YOLOv4[2], is an advanced object detection model built on Pytorch rather than Darknet. It includes a pre-training tool called AutoAnchor that optimizes anchor boxes using a k-means function and a Genetic Evolution algorithm, which improves detection by evolving anchors over 1000 generations with CIoU loss and Best Possible Recall metrics. The architecture incorporates a modified CSPDarknet53 backbone, SPPF[11] (spatial pyramid pooling fast) for efficient feature processing at various scales, and various advanced augmentations like Mosaic, MixUp[37], and HSV changes to enhance training. YOLOv5[29] offers multiple model sizes (from nano to extra large) to accommodate different device capabilities and performance needs. It is open-source, supported by a large community, and includes tools for easy integration and deployment on mobile devices.

**YOLOv6** [18], developed by Meituan Vision AI Department and released in September 2022, advances real-time object detection with significant improvements in speed and accuracy over predecessors like YOLOv5[29] and YOLOX[8]. Its architecture features a new EfficientRep[35] backbone for enhanced parallel processing, a PAN [20]topology neck, and an efficient decoupled head with hybrid-channel strategy. Notable innovations include Task Alignment Learning for label assignment, advanced VariFocal [38] and SIoU/GIoU[18] loss functions for precise detection, and optimized quantization techniques for faster performance. On the MS COCO dataset[19], YOLOv6-L achieved an AP of 52.5% and AP50 of 70% at about 50 FPS on an NVIDIA Tesla T4, demonstrating its effective-ness in high-speed object detection applications.

**YOLOv7**[32], released in July 2022, significantly advances object detection with speeds ranging from 5 to 160 FPS. Developed using only the MS COCO dataset, it introduces key architectural improvements:

Optimizes deep learning models by managing gradient paths efficiently.

• Model Scaling for Concatenation-based Models: Adjusts block depth and width to maintain optimal structure and efficiency.

YOLOv7 also features performance enhancements like revised reparameterized convolution (RepConvN) for better network structure, dual label assignment for precision in training and output, and advanced batch normalization for improved inference efficiency. These innovations set new benchmarks in object detection technology.

**YOLOv8**[13], released by Ultralytics in January 2023, expands the YOLO family with versions ranging from nano to extra-large. It supports diverse vision tasks like detection, segmentation, and classification. Enhancements include an updated C2f module in the backbone for better feature integration and an anchor-free, decoupled head design that improves accuracy by processing objectness, classification, and regression separately. It employs advanced loss functions like CIoU[34] and DFL[12], optimizing detection, especially for smaller objects.

Additionally, YOLOv8 introduces the YOLOv8-Seg for semantic segmentation, achieving top-tier results in benchmarks while maintaining high efficiency and speed. Available via CLI or as a PIP package, YOLOv8 offers easy integration for various applications. In testing, YOLOv8x reached an AP of 53.9% and a speed of 280 FPS on an NVIDIA A100, surpassing previous models like YOLOv5.

**RT-DETR** or Real-Time DEtection TRansformer (RT-DETR)[39], an innovative real-time end-to-end object detector that efficiently addresses the limitations of YOLO models and Transformer-based detectors like DETR[3]. RT-DETR optimizes performance by focusing first on enhancing speed without sacrificing accuracy, and then on improving accuracy while maintaining speed. Key advancements include an efficient hybrid encoder that rapidly processes multi-scale features, and an uncertainty-minimal query selection that enhances decoder accuracy.

RT-DETR also features adaptable speed tuning by varying the number of decoder layers, allowing customization to different operational scenarios without retraining. Performance tests on the COCO dataset show RT-DETR models achieving 53.1% and 54.3% AP for RT-DETR-R50 and RT-DETR-R101, respectively, surpassing prior models in both speed and accuracy. Moreover, post pre-training with Objects365[27], the models achieve up to 56.2% AP.

Despite these developments, the field of indoor object detection remains underexplored. The predominant use of the COCO dataset for model evaluation introduces a significant limitation, as it does not adequately represent the specific challenges posed by indoor settings, such as varied lighting conditions and complex spatial arrangements. This mismatch between training environments and real-world applications indicates a critical gap in the literature.

Our research seeks to address these shortcomings by focusing specifically on real-time indoor object detection on a custom dataset made from the OpenImages v7 dataset. We aim to develop a model that not only meets the generic criteria of accuracy and speed but is also fine-tuned to the specific demands and conditions of indoor environments. By utilizing a novel hybrid architecture that combines the strengths of CNNs and transformers, our work is poised to offer new insights and substantial improvements over the current state-of-the-art methods, specifically made for real-world indoor applications.

## 3 Methodology

## 3.1 Data collection

In this section, we'll describe the selection and preparation of a specialized subset of the OpenImages v7 dataset. We carefully curated 32 indoor object categories that are critical for understanding real-world indoor environments. This targeted dataset is designed to tackle specific challenges associated with indoor object detection, such as variable lighting and complex backgrounds.

To increase the robustness and diversity of our dataset, we incorporated the mosaic data augmentation technique. This technique constructs a single training image from four distinct images, enhancing the model's exposure to a variety of scenarios. It simulates complex interactions and occlusions between objects, which is vital for improving the model's generalization ability across different indoor settings.



Figure 2. Classes distribution across the dataset

Figure 2 illustrates the distribution of class instances within the dataset. Notably, class 13 (chair) is more prevalent, resulting in a skewed distribution. However, we chose not to balance this class distribution in order to evaluate the effectiveness of the mosaic data augmentation[4] and DFL loss[12]. These methods are intended to mitigate the impact of such imbalances on the model's performance during final evaluations.

In Figure 3 a training batch example is presented and the effect of the mosaic data augmentation technique can be clearly seen.

#### 3.2 Model Architecture

In this section, we detail the architecture of our object detection model, which is a hybrid system combining Convolutional Neural Networks (CNNs) and Transformer-based models. This design is strategically chosen to effectively handle the spatial hierarchies and contextual dependencies typical in indoor environments. The CNN component of our model serves as a high-performance feature extractor, essential for recognizing and delineating object features at various scales. In contrast, the Transformer component of the architecture takes these extracted features and integrates them across the entire image. This integration allows the model to dynamically focus on areas of interest, leveraging the local processing strengths of CNNs with the global perspective capabilities of Transformers. This



Figure 3. Training batch

synergy aims to provide a robust solution for the complexities of indoor object detection.

Our model's architecture is inspired by the YOLOv8n framework but is uniquely constructed from scratch using PyTorch. This approach gives us the flexibility to innovate without the constraints of the Ultralytics library.

A key modification in our model is the replacement of the conventional convolution layer at the input with a Focus layer. This layer concentrates spatial information from the input tensor into a channelrich tensor, effectively doubling the channel capacity while reducing spatial dimensions. This configuration aids the network in learning spatial hierarchies more efficiently. The ConvMod ( Convolutional Module ) is a conv2d layer followed by a batch normalization layer and a SiLU activation function.

Further enhancements include the integration of a custom module, the Spatial Pyramid Pooling Transformer (SPPT). This module combines a conventional spatial pyramid pooling (SPP) layer with a Transformer block inspired by the Vision Transformer (ViT)[6]. The introduction of the SPPT block has significantly reduced the model's floating-point operations (FLOPs), enhancing its computational efficiency and speed.

The SPPT module follows the backbone's final block and employs a standard Transformer layer to capture global information. The Transformer Encoder within this block features a Multi-Head Self-Attention Mechanism that updates and combines features from different subspaces for linear projection. This mechanism helps in capturing contextual information across the entire image and minimizing the loss of global features. The output through a multi-layer perceptron (MLP) ensures the non-linear enhancement of the selfattention mechanism's expressiveness.

Additionally, we draw inspiration from the YOLOv5 architecture's C3 module, which uses convolutional layers and bottleneck modules to reduce training parameters and computation. Our custom CSPlayer mimics these features but innovates by incorporating the Focus layer within the bottleneck structures rather than at the initial convolution layer.

This architectural design is depicted in Figure 4, illustrating both the individual components and their integration, highlighting how each contributes to efficient and accurate object detection within challenging indoor environments.

The loss structure consists of Binary Cross Entropy (BCE) for the

classification branch and combines Distribution Focal Loss (DFL) and Complete Intersection over Union (CIoU) loss for the regression branch.

DFL targets the class imbalance and improves accuracy in predicting bounding boxes, especially for objects with ambiguous boundaries by estimating the probability distribution of bounding box coordinates. CIoU loss enhances accuracy by considering the aspect ratios and overlap of the predicted and actual boxes.

The Binary Cross-Entropy loss for a binary classification task is defined as:

$$BCE = -\frac{1}{N} \sum_{i=1}^{N} \left[ y_i \cdot \log(p_i) + (1 - y_i) \cdot \log(1 - p_i) \right] \quad (1)$$

where N is the number of observations,  $y_i$  is the actual label (0 or 1), and  $p_i$  is the predicted probability of the observation being in class 1.

The Distribution Focal Loss, a variant that focuses more on difficult examples, is given by:

$$DFL = -\frac{1}{N} \sum_{i=1}^{N} \alpha_i (1 - p_i^{\gamma}) \log(p_i)$$
(2)

where  $p_i$  is the predicted probability,  $\alpha_i$  is a weighting factor for class imbalance,  $\gamma$  is a focusing parameter, and N is the number of observations.

The Complete Intersection over Union loss, which includes terms for overlap, distance, and aspect ratio, is defined as:

$$CIoU = 1 - IoU + \frac{\rho^2(b_{pred}, b_{true})}{c^2} + \alpha \cdot v$$
(3)

where:

- IoU is the Intersection over Union,
- ρ(b<sub>pred</sub>, b<sub>true</sub>) is the Euclidean distance between the center points
   of the predicted and actual bounding boxes,
- c is the diagonal length of the smallest enclosing box covering both bounding boxes,
- v is an aspect ratio consistency term,
- $\alpha$  is a trade-off parameter.

#### 3.3 Training Process

Each model was trained for 200 epochs, or until convergence, which was monitored through performance on a validation set. The training was conducted on an NVIDIA RTX 4090, utilizing a custom learning rate schedule to optimize convergence speed and model accuracy. The batch size was set at 16 to balance the trade-off between memory usage and the granularity of gradient updates. This training setup was chosen to ensure that the model learns effectively from the augmented dataset, adapting to the varied and complex scenarios presented during the training phase.

These hyperparameters are shown in Table 1

#### 3.4 Evaluation Metrics

The performance of our models was evaluated using several key metrics: accuracy, precision, recall, number of parameters and floatingpoint operations. Accuracy measures the proportion of correct predictions (both true positives and true negatives) among all evaluations. Precision and recall provide insights into the model's capability to classify indoor objects correctly without overfitting to frequent



Figure 4. The general architecture of our model

OS	Windows 11
CPU	I7 13700K
RAM	32Gb
GPU	RTX 4090 24Gb
Epochs	200
Batch size	16
Optimizer	SGD
Learning rate	Lambda LR

Table 1. Training environment and different hyper-parameters

labels or missing less common objects. Number of parameters and FlOPs are decisive for real-time applications, as they determine the feasibility of deploying these models in scenarios where rapid object detection is essential. These metrics collectively help in assessing the effectiveness and efficiency of our proposed model in real-world indoor settings.

Average Precision (AP) for a single class is calculated from the area under the precision-recall curve:

$$AP = \int_0^1 p(r) \, dr \tag{4}$$

Mean Average Precision (mAP) is the mean of the average precision scores for each class, widely used to evaluate object detection models:

$$\mathbf{mAP} = \frac{1}{N} \sum_{i=1}^{N} AP_i \tag{5}$$

where N is the number of classes and  $AP_i$  is the average precision for class *i*.

#### 4 Results and discussion

In our discussion of the results presented in the graph shown in Figure 5 and Table 2, we analyze the performance metrics across various models, focusing primarily on the mean Average Precision (mAP) and other key metrics such as recall and precision.

The graph shown in Figure5 illustrates the evolution of the mAP over training epochs, where it becomes evident that while all models exhibit convergence, their performance in terms of precision across different classes is quite similar. This similarity in mAP values, ranging between 0.375 and 0.410 for most models, underscores a general consistency in model performance. Notably, the RT-DETR model slightly leads with a mAP of 0.47, a marginal but significant upper hand which can be attributed to its considerably larger model size compared to its counterparts. Further reinforcing these findings, the results for recall and precision metrics also reveal minimal variability among the models, suggesting comparable effectiveness in identifying and accurately classifying objects within their respective datasets. This consistency is key for practical applications where predictability and reliability are valued.

Moreover, Table 2 offers a detailed comparison of the models in terms of computational efficiency and speed. Our model stands out as the fastest, evidenced by its fewer parameters and lower computational complexity, measured in GFLOPs. This efficiency translates

Table 2. Comparaison of different models

Model	Year	Size	Param(M)	FLOPs(B)	Precision	Recall	mAP50	mAP50/95	Inference time(ms)
Ours	2024	640	2.772	4.4	0.63	0.509	0.54	0.388	12.2
Yolov8-rtdetr	2024	640	9.545	16.9	0.62	0.506	0.595	0.403	134
RT-DETR	2023	640	32.04	103.6	0.677	0.609	0.628	0.467	184.5
Yolov8n (pretrained on COCO)	2023	640	3.2	8.7	0.625	0.55	0.582	0.414	13.4
Yolov8n (from scratch)	2023	640	3.2	8.7	0.614	0.538	0.575	0.407	19.9
Yolov7 tiny	2022	640	6.09	13.3	0.606	0.562	0.571	0.376	15.9
Yolov6n	2022	640	4.237	11.8	0.614	0.533	0.555	0.397	15.2
Yolov5n	2020	640	2.509	7.1	0.618	0.52	0.554	0.384	16.3



Figure 5. mAP50-95 evolution of all models across 200 epochs

directly into enhanced performance, particularly when considering the inference time tested on an RTX 2070 Mobile GPU. The data clearly show that our model not only maintains competitive accuracy but also excels in operational speed, making it particularly suited for real-time applications where rapid processing is paramount. These observations highlight the trade-offs between model complexity and speed, suggesting that while larger, more complex models like RT-DETR might achieve slightly higher accuracy, the increased computational demand may not always justify the marginal gains in performance. Our model's balance of speed and precision offers a compelling alternative, particularly for deployment in environments where both factors are critical. This balance will be a basis for further optimization and development aimed at refining our approach to indoor object detection, potentially extending its applicability to more demanding real-time environments.

In Figure 7, we present the confusion matrix for our model, providing a visual and quantitative analysis of its performance across various classes within our dataset. The matrix reveals areas where the model excels, accurately predicting certain classes, while also identifying classes where performance is less satisfactory. A primary factor contributing to these disparities is the uneven representation of classes within the training data. Classes that are underrepresented tend to be more challenging for the model to learn, resulting in a lower predictive accuracy.

Despite these challenges, the overall performance of our model is deemed commendable. The strengths demonstrated in wellrepresented classes indicate the model's capability to learn and generalize from the data provided.

Figure 6 offers a qualitative comparison between our model's detection outcomes and the ground truth for bounding box annotations,



Figure 6. Detection comparison between ground truth (left) and our model's detection (right).

across three randomly selected images from the test subset of our dataset. On the left side of the figure, we depict the ground truth bounding boxes, which serve as the benchmark for accurate object localization and classification. On the right, we display the corresponding detections made by our model, allowing for a direct visual assessment of its performance.

#### 5 Conclusion and future works

In conclusion, our work in real-time indoor object detection marks significant advancements in the field, particularly through the creation and utilization of a customized dataset derived from the expansive OpenImages dataset. Our model, a hybrid CNN-Transformer architecture, demonstrates the effectiveness of integrating these powerful technologies. We rigorously trained and compared various stateof-the-art real-time object detection models to our newly developed model. While the precision of these models was comparably high, our model stood out due to its lighter architecture, which substantially enhances processing speed. This feature is especially valuable



Figure 7. Confusion matrix of our model

in applications such as augmented reality (AR) and mixed reality (MR), where real-time processing on edge devices is pivotal. The efficiency and speed of our model not only meet these demanding requirements but also open avenues for further research and development across various real-time application scenarios, promising significant contributions to the practical deployment of different technologies in everyday use

## 5.1 Future works :

Despite the advancements achieved in our study, there remain several avenues for further enhancement and exploration in future work. Firstly, our model still relies on Non-Maximum Suppression (NMS) to refine detections, which can impact both speed and accuracy. To address this, we plan to explore end-to-end techniques like the one introduced by [22] and their DEYO mode, such technique could eliminate the need for NMS, thereby simplifying the detection process and potentially improving performance metrics. Additionally, integrating advanced tracking algorithms such as SORT[1] or DeepSORT [36] will enhance the model's applicability in real-time scenarios, particularly in video surveillance and interactive systems where we need to maintain object consistency . Further development and expansion of our dataset will also be a priority, with a focus on including a wider variety of indoor scenarios to better train models to handle diverse lighting conditions, occlusions, and complex object interactions. Lastly, we aim to revisit and refine the loss function used during training to better accommodate the unique challenges posed by indoor object detection. This will involve tailoring the loss function to enhance precision in bounding box predictions and to effectively address issues related to class imbalance. These improvements will help to solidify the foundation for more robust and accurate object detection systems, particularly for applications on edge devices.

## References

 A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft. Simple online and realtime tracking. In 2016 IEEE International Conference on Image Processing (ICIP). IEEE, Sept. 2016. doi: 10.1109/icip.2016.7533003. URL http://dx.doi.org/10.1109/ICIP.2016.7533003.

- [2] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao. Yolov4: Optimal speed and accuracy of object detection, 2020.
- [3] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-end object detection with transformers, 2020.
- [4] Y. Chen, P. Zhang, Z. Li, Y. Li, X. Zhang, L. Qi, J. Sun, and J. Jia. Dynamic scale training for object detection, 2021.
- [5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), volume 1, pages 886–893 vol. 1, 2005. doi: 10.1109/CVPR.2005.177.
- [6] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [7] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian. Centernet: Keypoint triplets for object detection, 2019.
- [8] Z. Ge, S. Liu, F. Wang, Z. Li, and J. Sun. Yolox: Exceeding yolo series in 2021, 2021.
- [9] R. Girshick. Fast r-cnn, 2015.
- [10] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation, 2014.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition, page 346–361. Springer International Publishing, 2014. ISBN 9783319105789. doi: 10.1007/978-3-319-10578-9\_23. URL http://dx.doi.org/10.1007/ 978-3-319-10578-9\_23.
- [12] M. S. Hossain, A. P. Paplinski, and J. M. Betts. Adaptive class weight based dual focal loss for improved semantic segmentation, 2020.
- [13] G. Jocher, A. Chaurasia, and J. Qiu. Ultralytics YOLO, Jan. 2023. URL https://github.com/ultralytics/ultralytics.
- [14] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Malloci, A. Kolesnikov, T. Duerig, and V. Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [15] S.-e. Laidoudi, M. Maidi, and S. Otmane. Real-time detection of lowtextured objects based on deep learning. In 2023 IEEE 25th International Workshop on Multimedia Signal Processing (MMSP), pages 1–6, 2023. doi: 10.1109/MMSP59012.2023.10337653.
- [16] S.-E. Laidoudi, M. Maidi, and S. Otmane. A cnn-based model for low-textured object detection and tracking. In 2023 IEEE 19th International Conference on Intelligent Computer Communication and Processing (ICCP), pages 221–227, 2023. doi: 10.1109/ICCP60212.2023. 10398729.
- [17] H. Law and J. Deng. Cornernet: Detecting objects as paired keypoints, 2019.
- [18] C. Li, L. Li, H. Jiang, K. Weng, Y. Geng, L. Li, Z. Ke, Q. Li, M. Cheng, W. Nie, Y. Li, B. Zhang, Y. Liang, L. Zhou, X. Xu, X. Chu, X. Wei, and X. Wei. Yolov6: A single-stage object detection framework for industrial applications, 2022.
- [19] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár. Microsoft coco: Common objects in context, 2015.
- [20] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia. Path aggregation network for instance segmentation, 2018.
- [21] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single Shot MultiBox Detector, page 21–37. Springer International Publishing, 2016. ISBN 9783319464480. doi: 10.1007/978-3-319-46448-0\_2. URL http://dx.doi.org/10.1007/ 978-3-319-46448-0\_2.
- [22] H. Ouyang. Deyo: Detr with yolo for end-to-end object detection, 2024.
- [23] J. Redmon and A. Farhadi. Yolo9000: Better, faster, stronger, 2016.
- [24] J. Redmon and A. Farhadi. Yolov3: An incremental improvement, 2018.
- [25] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection, 2016.
- [26] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks, 2016.
- [27] S. Shao, Z. Li, T. Zhang, C. Peng, G. Yu, X. Zhang, J. Li, and J. Sun. Objects365: A large-scale, high-quality dataset for object detection. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 8429–8438, 2019. doi: 10.1109/ICCV.2019.00852.
- [28] Z. Tian, C. Shen, H. Chen, and T. He. Fcos: Fully convolutional onestage object detection, 2019.
- [29] Ultralytics. YOLOv5: A state-of-the-art real-time object detection system. https://docs.ultralytics.com, 2021. Accessed: insert date here.
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2023.
- [31] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE Computer Society*

Conference on Computer Vision and Pattern Recognition. CVPR 2001, volume 1, pages I–I, 2001. doi: 10.1109/CVPR.2001.990517.

- [32] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors, 2022.
- [33] C.-Y. Wang, I.-H. Yeh, and H.-Y. M. Liao. Yolov9: Learning what you want to learn using programmable gradient information, 2024.
- [34] X. Wang and J. Song. Iciou: Improved loss based on complete intersection over union for bounding box regression. *IEEE Access*, 9:105686– 105695, 2021. doi: 10.1109/ACCESS.2021.3100414.
- [35] K. Weng, X. Chu, X. Xu, J. Huang, and X. Wei. Efficientrep:an efficient repvgg-style convnets with hardware-aware neural network design, 2023.
- [36] N. Wojke, A. Bewley, and D. Paulus. Simple online and realtime tracking with a deep association metric, 2017.
- [37] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization, 2018.
- [38] H. Zhang, Y. Wang, F. Dayoub, and N. Sünderhauf. Varifocalnet: An iou-aware dense object detector, 2021.
- [39] Y. Zhao, W. Lv, S. Xu, J. Wei, G. Wang, Q. Dang, Y. Liu, and J. Chen. Detrs beat yolos on real-time object detection, 2024.