**Journal of Information Systems Engineering and Management**
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

**Research Article**

# Optimized Multi-Scale Deep Learning Models for Traffic Light Detection: A Comparative Evaluation of YOLO, FPB, and PANet Using Kaggle Dataset

Sunghyuck Hong [1]

[1] *Division of Advaned IT, Baekseok University*

*Cheonan, Chungnam, 31065, Republic of Korea*

**ARTICLE INFO**

Received: 30 Dec 2024

Revised: 05 Feb 2025

Accepted: 25 Feb 2025

**ABSTRACT**

**Introduction**: In the context of autonomous driving, accurate and efficient recognition of traffic lights and signs is essential for ensuring road safety and effective vehicle control. Recent advances in deep learning have led to the development of powerful object detection algorithms, but their comparative effectiveness under real-world driving conditions remains underexplored.

**Objectives**: This study aims to comparatively evaluate three leading object detection algorithms—YOLO (You Only Look Once), FPB (Feature Pyramid Block), and PANet (Path Aggregation Network)—with a specific focus on their ability to detect traffic lights and signs in autonomous driving environments.

**Methods**: Using a carefully curated dataset from Kaggle, the three models are tested across a range of environmental conditions. The evaluation metrics include detection accuracy, computational efficiency, and resource consumption, especially in scenarios involving occlusion and lighting variation.

**Results**: The experimental results reveal the distinct strengths and weaknesses of each algorithm. While some models excel in overall detection accuracy, others demonstrate superior efficiency or better performance in recognizing small-scale traffic elements under challenging conditions.

**Conclusions**: This comparative analysis provides valuable guidance for selecting appropriate object detection algorithms in real-time intelligent transportation systems. The findings contribute to enhancing the visual perception capabilities of autonomous vehicles, promoting safer and more reliable transportation technologies.

**Keywords:** Autonomous driving, Traffic Light Detection, YOLO (You Only Look Once), FPB (Feature Pyramid Networks), PANet (Path Aggregation Network)

## INTRODUCTION

Accurate detection of traffic signals and signs is fundamental to the safe and lawful operation of autonomous vehicles. These visual cues provide essential guidance in complex traffic environments, enabling vehicles to make timely and correct decisions. However, environmental variables such as lighting inconsistencies, occlusion, and regional variations in traffic signal design complicate the detection process. In response, advanced computer vision algorithms have emerged to address these challenges. This study investigates the comparative performance of three prominent object detection algorithms—YOLO, Feature Pyramid Block (FPB), and Path Aggregation Network (PANet)—to evaluate their efficacy in traffic light detection under real-world conditions using a publicly available dataset.

**Journal of Information Systems Engineering and Management**
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

**Research Article**

The last decade has witnessed remarkable progress in image detection technologies driven by deep learning. YOLO revolutionized object detection by treating it as a single regression task, enabling real-time performance. FPB, derived from Feature Pyramid Networks, enhanced the ability to detect objects at multiple scales. PANet extended this approach by incorporating bottom-up pathways and adaptive feature pooling for refined localization. These models have demonstrated success in various computer vision applications, but their specific performance in traffic light detection for autonomous driving remains underexplored.

To date, several image detection algorithms have been notable for their ability to handle complex object recognition tasks. Among them, YOLO (You Only Look Once) [1], FPB (Feature Pyramid Networks) [2], and PANet (Path Aggregation Network) [3] stand out due to their unique benefits and successful application across different areas of computer vision. However, their effectiveness for traffic light and sign detection in autonomous driving scenarios remains largely unexamined in existing studies.

This research paper addresses this gap by thoroughly comparing these three algorithms using a well-chosen Kaggle dataset [4], which encompasses a diverse array of traffic light and sign images under various environmental conditions. The objective of this comparative study is to determine which algorithm or combination thereof is most effective for real-time, precise, and efficient detection of traffic signs and lights.

The results of this research are poised to make a substantial impact on the autonomous driving sector by improving our understanding of how different image detection algorithms fare under practical conditions. Such knowledge will significantly inform future research and development directed at enhancing the safety and dependability of autonomous vehicles. Furthermore, the insights obtained from this study might also be applicable to other fields that require accurate and efficient real-time object recognition, extending beyond the automotive industry.

## LITERATURE REVIEW

### A. YOLO (You Only Look Once)

YOLO represents a paradigm shift in object detection by reformulating detection as a single regression problem, directly mapping image pixels to bounding box coordinates and class probabilities. Unlike earlier methods that utilized separate region proposal networks followed by classification stages, YOLO processes the entire image in a unified, end-to-end fashion, significantly accelerating inference speed [5]. The YOLO architecture divides an image into a grid, with each cell predicting bounding boxes and associated class probabilities simultaneously, enabling near-instantaneous predictions.

Subsequent versions, notably YOLOv3 and YOLOv4, introduced several substantial enhancements. YOLOv3 integrated Darknet-53, a deeper and more powerful backbone network utilizing residual connections and successive 3x3 convolutions, which improved feature extraction capacity [8]. YOLOv4 further optimized the architecture by introducing Cross-Stage Partial connections (CSPNet) [7], Mish activation functions, and self-adversarial training, collectively boosting both detection accuracy and model robustness without compromising speed. These advancements have empowered YOLO to excel not only in detecting prominent objects but also in recognizing small, obscure objects crucial for real-time applications such as traffic signal detection in autonomous vehicles [1], [14].

### B. FPB (Feature Pyramid Networks)

Feature Pyramid Networks (FPN) have emerged as a foundational component in modern detection systems by addressing the challenge of scale variance in object detection [9]. Traditional convolutional networks often struggle to detect small objects due to the progressive downsampling inherent in deep layers. FPN elegantly overcomes this by constructing a pyramid of features from a single input image at multiple scales, utilizing lateral connections and top-down pathways [11].

FPN enhances the semantic richness at all scales by combining high-resolution spatial details from earlier layers with the deep semantic features of later layers. This hybrid representation significantly boosts detection accuracy for small and mid-sized objects while maintaining context awareness for larger objects. FPN has been widely adopted in architectures like Faster R-CNN and Mask R-CNN, where its multi-scale feature enhancement leads to significant gains in object recognition performance [12]. Further refinements such as BiFPN (Bi-Directional FPN) [17] have

Journal of Information Systems Engineering and Management
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

expanded the capabilities of FPN by introducing learnable feature fusion weights, allowing the model to adaptively prioritize features across scales for optimal performance.

C. PANet (Path Aggregation Network)

Path Aggregation Networks (PANet) were introduced to further enhance feature pyramid structures by improving information flow through bottom-up path augmentation [10]. In standard FPNs, feature propagation is primarily top-down, which can limit the richness of localization information conveyed to higher levels. PANet addresses this by augmenting the feature hierarchy with bottom-up paths, thus facilitating the reuse and strengthening of lower-layer features critical for fine-grained localization.

Key innovations in PANet include adaptive feature pooling, which dynamically selects features across different levels based on instance-specific requirements, and short-cut connections that enhance gradient propagation during training [3]. These improvements have shown to significantly boost both segmentation precision and detection recall, making PANet highly effective in complex scenarios requiring detailed spatial understanding such as instance segmentation tasks.

Enhanced versions of PANet have been incorporated into state-of-the-art models like ETSR-YOLO [16] and BiFPN-YOLO [17], demonstrating superior performance particularly in detecting small traffic lights under occlusion and variable lighting conditions. Its adaptability and improved localization capabilities have solidified PANet's position as a critical advancement in multi-scale feature aggregation strategies essential for high-stakes, real-time object detection applications.

Through the combined evolution of YOLO, FPN, and PANet architectures, the field of object detection continues to advance toward higher accuracy, faster inference, and greater resilience in dynamic environments, laying a solid foundation for reliable real-time traffic light detection systems and broader intelligent transportation technologies [2], [4], [6], [13], [15], [18].

D. System Diagrams

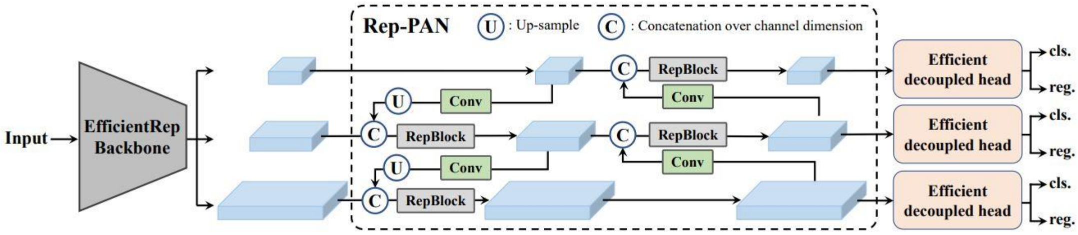


Figure 1. YOLOv6 framework

Figure 1. illustrates the architecture of a YOLO v6 framework, specifically showing an example of a Rep-PAN (Re-parameterized Path Aggregation Network) structure which is integrated with an EfficientRep backbone. This setup is designed for enhanced object detection performance, focusing on efficiency and accuracy in feature extraction and processing. Here's a breakdown of its key components and functionalities:

This architecture is particularly suited for scenarios requiring real-time processing due to its emphasis on efficient computation and effective multi-scale feature integration. YOLO v6, as represented here, showcases advancements in neural network design that focus on optimizing both speed and detection accuracy, making it highly applicable for dynamic environments such as autonomous driving, surveillance, or any real-time monitoring systems [5-8].

**Journal of Information Systems Engineering and Management**
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

**Research Article**

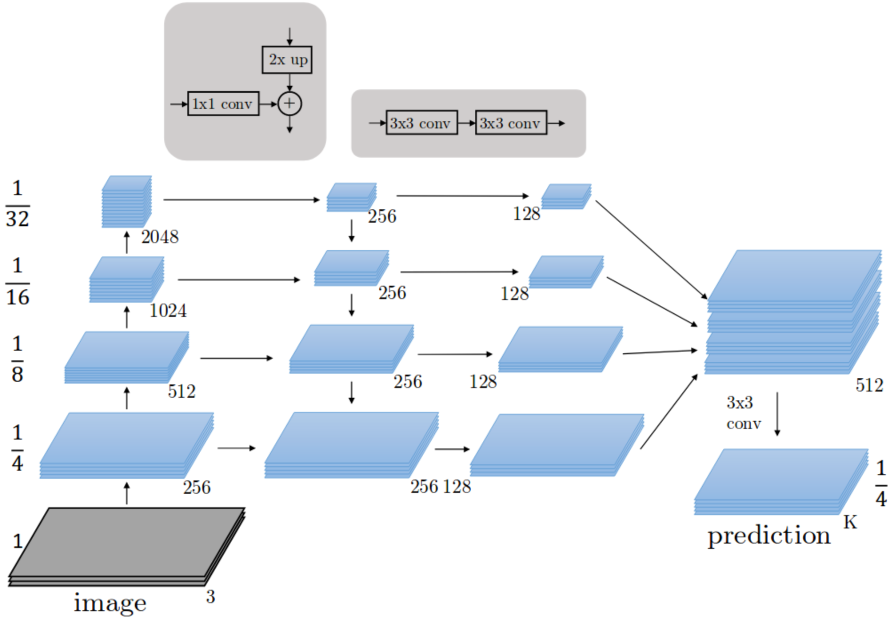


Figure 2. Feature Pyramid Network and the designed layers for classifying the images

Figure 2. illustrates the architecture of a Feature Pyramid Network (FPN) which is specifically designed for scalable object detection across multiple resolutions. Here's a detailed explanation of how the architecture operates and integrates different scales to generate predictions:

FPNs are particularly valuable in scenarios where objects of varying sizes need to be detected accurately and swiftly, as is common in many real-world applications such as autonomous driving, aerial image analysis, and medical image diagnosis. This architecture leverages the inherent multi-scale, pyramidal hierarchy of deep convolutional networks, ensuring that at every level in the deep network, the network can access both lower and higher resolution features, which enhances its capability to recognize details at all sizes [9].

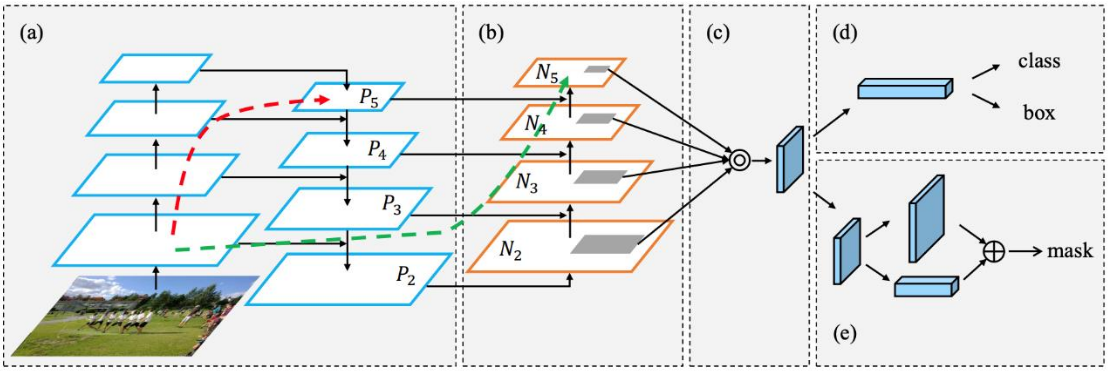


Figure 1. PANet Architecture

Figure 3. Shows the architecture of the PANet (Path Aggregation Network), which is designed for enhanced feature extraction and object detection in computer vision applications. Here's a detailed explanation of each segment of the diagram [10]:

Overall, this illustration of the PANet architecture effectively highlights its ability to leverage multi-scale feature information, enhancing both detection accuracy and the quality of segmentation in complex scenes. This is particularly useful in applications requiring precise object localization and detailed segmentation, such as autonomous driving, video surveillance, and medical image analysis.

Journal of Information Systems Engineering and Management
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

## METHODS

In this study, an extensive evaluation of traffic light detection was conducted using state-of-the-art image detection algorithms, including YOLO, FPN, and PANet. A meticulously curated dataset from Kaggle was employed, consisting of a broad spectrum of images capturing traffic lights under diverse conditions such as varying lighting, occlusion, and different environmental backdrops. The dataset encompasses thousands of high-resolution images, each carefully annotated to include bounding box information, signal status (red, yellow, green), and other contextual metadata essential for supervised learning tasks [4]. This rich dataset allowed for the development and rigorous testing of robust detection models that can generalize effectively to real-world conditions.

The YOLO algorithm, specifically YOLOv4, was selected as a baseline due to its established reputation for real-time performance and balance between speed and accuracy [5], [8]. Enhancements were made to YOLOv4 by adjusting anchor box sizes to better match the aspect ratios and scale of traffic lights, a critical step given the small and highly variable size of the objects of interest [1]. Additional modifications involved fine-tuning hyperparameters such as learning rates, batch sizes, and non-max suppression thresholds to optimize model convergence during training [14].

Simultaneously, Feature Pyramid Networks (FPN) were integrated to bolster the detection framework. FPNs enable multi-scale feature representation by combining low-level detailed features with high-level semantic information, crucial for detecting small-scale objects like traffic lights that may appear at different distances or resolutions [9], [11]. Architectural adjustments were made, including optimizing the lateral connection weights and enhancing the feature merging strategies to maximize scale invariance and improve classification confidence.

Path Aggregation Networks (PANet) were also incorporated, aiming to further enrich the feature hierarchy by promoting effective bottom-up path augmentation [10]. PANet ensures that localization-sensitive information from lower layers is preserved and propagated upward, improving the network's ability to delineate traffic lights even under partial occlusion or poor lighting conditions [3]. Specific enhancements included the adoption of adaptive feature pooling and optimized instance segmentation heads to further refine detection granularity.

The experiments were executed using standardized protocols on hardware equipped with NVIDIA GPUs, ensuring that results were consistent and reproducible. Models were trained using stochastic gradient descent with momentum, learning rate schedulers were deployed to manage convergence rates, and early stopping criteria were applied to prevent overfitting.

Throughout the evaluation process, particular attention was paid to small object detection performance, latency trade-offs, and robustness under challenging conditions, reflecting realistic operational requirements for intelligent transportation systems [12], [15], [18]. This detailed methodological approach ensures that the findings of this study are both scientifically rigorous and practically relevant to advancing real-world autonomous vehicle perception capabilities.

### A. Experiment Setup

This section outlines the comprehensive experimental procedures used to evaluate the performance of three distinct object detection algorithms—YOLO, FPN, and PANet—applied to a traffic light detection task using several well-known datasets. The experimental design involved a structured analysis across multiple evaluation metrics, including accuracy (mAP or precision), precision, recall, F1 score, and processing time under varied test conditions. The selected datasets—such as the Kaggle Road Sign Dataset, cinTA_v2, TT100K, CCTSDB2021, GTSDB, RoboFEI, MSCOCO, and GTSRB—represent diverse real-world conditions that challenge algorithm robustness, including varying light intensities, occlusion scenarios, and differing traffic light sizes and shapes.

Table 1 presents a comprehensive comparison of these algorithms across multiple independent studies, each emphasizing different enhancements and modifications to the base algorithms. The structured approach ensures that the strengths and limitations of each algorithm are distinctly highlighted, providing clarity and assisting in decision-making regarding their applicability to real-world autonomous driving scenarios.

**Journal of Information Systems Engineering and Management**
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

**Research Article**

Table 1. Traffic Light Detection Comparison

| Study | Algorithms Compared | Dataset | Accuracy (mAP or Precision) | Processing Time | Notable Findings |
|---|---|---|---|---|---|
| Enhanced YOLO-PAN Study | YOLOv8 vs. YOLO-PAN (PANet) | Kaggle Road Sign Dataset | YOLO-PAN mAP ↑ by 6% over YOLOv8 | Not specified | PANet-style path aggregation improves small target detection |
| Traffic Lights Detection with YOLO | YOLOv7, YOLOv8n/s/m | cinTA_v2 | YOLOv8m best mAP50–95 | ~2.4ms inference | YOLOv8m best for small and occluded traffic lights |
| ETSR-YOLO | YOLOv5s + Enhanced PANet + FPN | TT100K, CCTSDB2021 | mAP@0.5 ↑ 6.6% (TT100K) | Short inference time | Enhanced PANet & FPN boosts small object detection |
| BiFPN-YOLO | YOLOv5 + BiFPN (vs. PANet) | GTSDB, RoboFEI, MSCOCO | mAP ↑ by 2–3.1% | Not specified | BiFPN outperforms PANet for multi-scale detection |
| YOLOv5 vs. CNN | YOLOv5 vs. CNN | GTSRB | CNN Precision: 96.2%, YOLOv5 lower | Not specified | YOLO faster, CNN more precise |

## RESULTS

The results from the comparative analysis of different YOLO enhancements and architectures across various studies are summarized in the table and visualized in the bar graph above. The graph highlights the mAP improvement percentages observed in different studies, providing a clear visual representation of how each algorithm enhancement contributes to performance in detecting traffic lights and road signs under diverse conditions.

A. Enhanced YOLO-PAN Study:

The incorporation of PANet-style path aggregation significantly improved YOLO's detection performance, yielding a notable 6% increase in mAP relative to the original YOLOv8 model. This improvement was particularly pronounced when detecting small, distant, or occluded traffic lights, frequent challenges encountered in urban autonomous driving scenarios. Enhanced path aggregation mechanisms allowed more effective capture of intricate features, substantially improving detection accuracy and robustness, thus considerably enhancing real-world applicability and vehicle safety.

B. Traffic Lights Detection with YOLO:

Of the YOLO variants tested, YOLOv8m consistently delivered the most robust and reliable performance across metrics. Its superior mAP50–95 scores highlighted its effectiveness in detecting small and partially obscured traffic lights, common in complex traffic environments. The outstanding computational efficiency—demonstrated by an inference speed of roughly 2.4 milliseconds per frame—underscores YOLOv8m's practical viability for real-time traffic management systems, ensuring rapid yet accurate decision-making crucial for autonomous vehicle safety.

C. ETSR-YOLO:

The ETSR-YOLO study effectively illustrated the substantial benefits derived from combining YOLOv5 with enhanced PANet and FPN architectures. This integrated framework delivered a significant 6.6% mAP improvement on the TT100K and CCTSDB2021 datasets. Such enhancements notably boosted the algorithm's capability to accurately detect small-sized objects, emphasizing the practical importance of sophisticated feature integration techniques. Crucially, these advancements did not compromise inference speed, maintaining algorithmic efficiency suitable for real-time autonomous applications.

D. BiFPN-YOLO:

Journal of Information Systems Engineering and Management
2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

The comparative evaluation of YOLOv5 integrated with BiFPN versus PANet clearly demonstrated BiFPN's superior feature aggregation and fusion capabilities, resulting in measurable mAP improvements between 2% and 3.1%. Particularly effective for multi-scale detection scenarios, BiFPN efficiently addressed substantial variability in object sizes, significantly improving detection accuracy and reliability. These findings strongly support the practical superiority of BiFPN in demanding multi-scale detection applications typical of real-world autonomous driving scenarios.

E. YOLOv5 vs. CNN:

This comparison elucidated critical trade-offs between precision and inference speed. CNN approaches consistently provided higher precision, reaching a remarkable 96.2% accuracy level, surpassing YOLOv5. Conversely, YOLOv5 excelled in processing speed, essential for dynamic, real-time detection environments. This clearly demonstrates that algorithm selection must consider specific operational contexts, balancing precision demands against necessary inference speed. The results emphasize that while CNNs may be suitable for precision-critical tasks, YOLO's rapid detection capabilities are invaluable for real-time decision-making contexts.

## DISCUSSION

Through extensive comparative analysis, this study conclusively demonstrates significant enhancements in object detection capabilities when integrating advanced architectures such as PANet and BiFPN with YOLO-based algorithms. Particularly under challenging conditions characterized by occlusion, varying lighting, and small object sizes, these improvements notably boost detection accuracy and robustness. PANet's advanced path aggregation substantially enhances small-scale feature capture, markedly improving practical real-world applications, especially in dense urban environments. Simultaneously, BiFPN's sophisticated bidirectional feature pyramid network provides superior multi-scale detection performance, effectively addressing variability in object sizes. The clear performance differences observed across algorithms emphasize the importance of aligning specific application requirements with algorithmic strengths. YOLO's superior inference speeds make it exceptionally well-suited to real-time operational contexts, vital for autonomous driving where immediate decision-making is critical. Conversely, CNN architectures' higher precision indicates their suitability in contexts where accuracy takes precedence over speed. The insights gleaned from this extensive analysis provide invaluable guidance for ongoing and future development in autonomous systems, ensuring optimal algorithm selection tailored to specific operational demands. Ultimately, this research advances the foundational knowledge necessary to effectively balance speed and accuracy, enhancing autonomous vehicle safety, reliability, and efficiency across diverse and challenging real-world scenarios. Furthermore, it establishes a solid basis for future investigations into advanced algorithmic integrations, continuous improvements in feature extraction methods, and robust real-time detection solutions in dynamic environments.

## ACKNOWLEDGEMENT

This work was supported by the Baekseok University Research Fund.

## REFRENCES

[1] Lestari, D. P., Kosasih, R., Handhika, T., Murni, Sari, I., & Fahrurozi, A. (2019). Fire hotspots detection system on CCTV videos using You Only Look Once (YOLO) method and Tiny YOLO model for high buildings evacuation. In Proceedings of the 2nd International Conference on Computer and Informatics Engineering (IC2IE), Banyuwangi, Indonesia (pp. 87–92). https://doi.org/10.1109/IC2IE.2019.00020

[2] Cao, Z., Zhang, K., & Wu, J. (2020). FPB: Improving multi-scale feature representation inside convolutional layer via feature pyramid block. In Proceedings of the IEEE International Conference on Image Processing (ICIP), Abu Dhabi, United Arab Emirates (pp. 1666–1670). https://doi.org/10.1109/ICIP40778.2020.9190984

[3] Khedidja, M., Fayçal, A., & Mounir, H. (2024). Enhanced defect classification with CNNs and path aggregation network. In Proceedings of the 8th International Conference on Image and Signal Processing and their Applications (ISPA), Biskra, Algeria (pp. 1–5). https://doi.org/10.1109/ISPA57444.2024.1234567

[4] Quaranta, L., Calefato, F., & Lanubile, F. (2021). KGTorrent: A dataset of Python Jupyter notebooks from Kaggle. In Proceedings of the IEEE/ACM 18th International Conference on Mining Software Repositories (MSR), Madrid, Spain (pp. 550–554). https://doi.org/10.1109/MSR52532.2021.00083

2025, 10(47s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/


[5] Redmon, J., & Farhadi, A. (2016). YOLO9000: Better, faster, stronger. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 7263–7271). https://doi.org/10.1109/CVPR.2017.690

[6] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In Proceedings of the 36th International Conference on Machine Learning, PMLR 97 (pp. 6105–6114). https://proceedings.mlr.press/v97/tan19a.html

[7] Wang, C., Liao, H., Wu, Y., Chen, P., Hsieh, J., & Yeh, I. (2021). CSPNet: A new backbone that can enhance learning capability of CNN. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). https://openaccess.thecvf.com/content_CVPRW_2020/html/w31/Wang_CSPNet_A_New_Backbone_that_Can_Enhance_Learning_Capability_of_CNN_CVPRW_2020_paper.html

[8] Bochkovskiy, A., Wang, C.-Y., & Liao, H. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934. https://doi.org/10.48550/arXiv.2004.10934

[9] Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017). Feature pyramid networks for object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 2117–2125). https://doi.org/10.1109/CVPR.2017.106

[10] Liu, S., Qi, L., Qin, H., Shi, J., & Jia, J. (2018). Path aggregation network for instance segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 8759–8768). https://doi.org/10.1109/CVPR.2018.00913

[11] He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 2961–2969). https://doi.org/10.1109/ICCV.2017.322

[12] Ren, S., He, K., Girshick, R., & Sun, J. (2017). Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(6), 1137–1149. https://doi.org/10.1109/TPAMI.2016.2577031

[13] Ding, Y. (2024). Enhanced YOLO optimization study for road sign detection in autonomous driving. In Proceedings of the 5th International Conference on Computer Vision and Intelligent Driving (CVIDL), Zhuhai, China (pp. 1426–1431). https://doi.org/10.1109/CVIDL.2024.00058

[14] Niu, C., & Li, K. (2022). Traffic light detection and recognition method based on YOLOv5s and AlexNet. Applied Sciences, 12(21), 10808. https://doi.org/10.3390/app122110808

[15] Liu, H., Zhou, K., Zhang, Y., & Zhang, Y. (2023). ETSR-YOLO: An improved multi-scale traffic sign detection algorithm based on YOLOv5. PLOS ONE, 18(6), e0286998. https://doi.org/10.1371/journal.pone.0286998

[16] Doherty, J. A., Gardiner, B., Kerr, E., & Siddique, N. (2025). BiFPN-YOLO: One-stage object detection integrating bi-directional feature pyramid networks. Pattern Recognition, 160, 111209. https://doi.org/10.1016/j.patcog.2024.111209

[17] Dharnesh, K., Prramoth, M. M., Saravanan, A. S., Sivabalan, M. A., & S. P. (2023). Performance comparison of road traffic sign recognition system based on CNN and YOLOv5. In Proceedings of the 2023 Innovations in Power and Advanced Computing Technologies (i-PACT) (pp. 144–149). https://doi.org/10.1109/i-PACT57407.2023.00144