Ahmad Alotaibi, 2026, 14:4
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Applications of YOLO in the Oil and Gas Industry

Ahmad Alotaibi[1], Masnour Asiri[2], Naif Qahtani[3], Abdullah Aldawsari[4]
King Fahd University of Petroleum and Minerals

Abstract- The You Only Look Once (YOLO) algorithm has become one of the most widely used deep learning frameworks for real-time object detection due to its high detection accuracy, computational efficiency, and ability to simultaneously perform object localization and classification. This paper reviews the applications of YOLO in industrial environments, with particular emphasis on battery leakage detection, facial detection for plant access control, and hazardous gas leakage detection. These applications demonstrate how YOLO enhances manufacturing quality, strengthens industrial security, improves workplace safety, and supports intelligent automation through rapid and reliable object detection. Despite its significant advantages, the practical implementation of YOLO presents several challenges. The model requires substantial computational resources and specialized hardware to achieve real-time performance, particularly when processing high-resolution images or multiple video streams. Furthermore, YOLO relies heavily on large, diverse, and accurately annotated datasets, while class imbalance in industrial datasets can reduce detection performance for rare but safety-critical objects. To overcome these limitations, several solutions are discussed, including lightweight YOLO architectures, model compression techniques, edge artificial intelligence (AI) hardware acceleration, data augmentation, synthetic data generation, balanced sampling methods, and advanced loss functions. Overall, the continuous evolution of YOLO architectures and optimization techniques has significantly improved its applicability across industrial sectors, making YOLO a promising and adaptable solution for smart manufacturing, industrial inspection, surveillance, and safety monitoring within Industry 4.0 environments.

Keywords: Routing Protocols, Convergence, Security, and Scalability.

I. INTRODUCTION

The rapid advancement of artificial intelligence (AI) and deep learning has transformed the field of computer vision, enabling machines to automatically detect, classify, and analyze objects with unprecedented accuracy and speed [1]. Among the various deep learning techniques, object detection has become a fundamental technology for numerous industrial applications, including quality inspection, surveillance, access control, autonomous systems, and safety monitoring [2].

Traditional computer vision methods often rely on handcrafted features and image processing techniques, which are sensitive to changes in lighting, background complexity, object orientation, and occlusion, resulting in limited performance in dynamic industrial environments [3]. In contrast, deep learning-based object detection models learn discriminative features directly from large datasets, significantly improving detection accuracy and robustness under diverse operating conditions [4].

The You Only Look Once (YOLO) algorithm has emerged as one of the most widely adopted real-time object detection frameworks because it performs object localization and classification simultaneously using a single convolutional neural network [5]. Unlike two-stage object detectors, which first generate region proposals and then classify objects, YOLO processes the entire image in a single forward pass, enabling high-speed inference while maintaining competitive detection accuracy [5], [6].

Since the introduction of the original YOLO model in 2016, several improved versions have been developed, incorporating advanced backbone networks, feature fusion techniques, attention mechanisms, and lightweight architectures that further enhance detection performance and computational efficiency [6]. These continuous improvements have enabled YOLO to be successfully applied across a wide range of industrial applications requiring real-time decision-making.

In industrial environments, YOLO has demonstrated remarkable effectiveness in tasks such as battery leakage detection, hazardous gas leakage monitoring, and facial detection for secure plant access control. These applications contribute to improved product quality, enhanced operational safety, increased security, and reduced dependence on manual inspection by enabling automated, intelligent monitoring systems [2], [7]. However, despite its significant advantages, the practical deployment of YOLO still faces challenges related to computational resource requirements, hardware limitations, dependency on large annotated datasets, and class imbalance during model training [6], [8]. Addressing these challenges through lightweight model architectures, data augmentation techniques, and hardware optimization is essential for maximizing the performance and scalability of YOLO-based systems. Consequently, YOLO continues to play a central role in advancing Industry 4.0 by providing an efficient, accurate, and adaptable solution for intelligent object detection in modern industrial environments.

II. APPLICATIONS OF YOLO IN OIL & GAS INDUSTRY

This section sheds light on the applications of YOLO in Oil and Gas Industry.

Battery Leakage

Battery leakage is one of the most critical defects in lithium-ion batteries because it can reduce battery performance, shorten operational lifespan, and increase the risk of thermal runaway, fire, and explosion, making early detection essential in modern battery manufacturing and quality control [9], [10]. Conventional inspection methods, including manual visual inspection and traditional image processing techniques, are often limited by inconsistent accuracy, operator fatigue, sensitivity to illumination changes, and the inability to perform reliable real-time inspection in high-speed production environments [11], [12].

To overcome these limitations, deep learning-based object detection algorithms have become increasingly popular, with the You Only Look Once (YOLO) family emerging as one of the most effective approaches for automated defect detection due to its capability of simultaneously performing object localization and classification in a single neural network inference [13]. Recent versions such as YOLOv8 and YOLOv11 incorporate improved backbone networks, feature pyramid structures, and attention mechanisms that significantly enhance the detection of small and low-contrast defects, including battery leakage, cracks, scratches, swelling, and surface contamination [9], [10].

These improvements enable YOLO models to achieve high detection accuracy while maintaining real-time inference speeds exceeding 100 frames per second, making them suitable for deployment in industrial production lines where batteries move continuously under machine vision systems [9], [13]. During model training, annotated datasets containing leakage images captured under different lighting conditions, camera perspectives, and manufacturing environments allow the network to learn robust feature representations and generalize effectively to unseen battery samples [9], [11].

Furthermore, data augmentation techniques such as Mosaic augmentation, MixUp, image flipping, scaling, and random erasing improve model robustness and reduce overfitting, particularly because battery leakage defects are relatively rare compared to normal battery samples [9], [13]. Integrating YOLO with industrial machine vision systems enables continuous online inspection, automatic rejection of defective batteries, reduced labor costs, improved product quality, and enhanced manufacturing safety [10], [12]. Consequently, YOLO has become one of the leading deep learning frameworks for battery leakage detection because it offers an excellent balance between detection accuracy, computational efficiency, and real-time industrial applicability [9], [11], [13].

Face Recognition

Face recognition has become one of the most widely adopted biometric technologies for controlling access to industrial facilities because it offers a contactless, fast, and secure method of authenticating personnel while reducing the risks associated with unauthorized entry [14], [15].

In modern industrial plants, access control systems must accurately identify employees and authorized visitors in real time despite challenges such as varying illumination, different viewing angles, protective equipment, and high pedestrian traffic [16]. Traditional face detection algorithms, including Haar Cascade classifiers and Histogram of Oriented Gradients (HOG)-based methods, often experience reduced accuracy under these challenging conditions and require multiple processing stages before recognition can be performed [17].

The You Only Look Once (YOLO) family of object detection algorithms has significantly improved face detection by performing object localization and classification simultaneously in a single forward pass through a convolutional neural network, enabling high-speed and high-accuracy detection suitable for real-time applications [18]. Recent versions such as YOLOv8 and YOLOv11 incorporate enhanced feature extraction networks, feature pyramid architectures, and attention mechanisms that improve the detection of small, partially occluded, and low-resolution faces while maintaining real-time inference speeds exceeding 100 frames per second on modern graphics processing units (GPUs) [18], [19]. In industrial access control applications, YOLO is typically integrated with facial recognition models such as FaceNet, ArcFace, or InsightFace, where YOLO first detects and crops facial regions before the recognition network verifies the person's identity against a secure employee database [15], [20].

This two-stage pipeline significantly improves both detection accuracy and recognition speed while minimizing computational overhead [15]. Furthermore, YOLO-based face detection systems can be trained using large annotated datasets containing faces captured under different lighting conditions, facial expressions, camera angles, and environmental settings, enabling robust performance in real-world industrial environments [16], [17]. Data augmentation techniques, including random rotation, brightness adjustment, scaling, and horizontal flipping, further enhance model generalization and reduce overfitting [5]. Consequently, integrating YOLO with modern facial recognition technologies provides an efficient and reliable solution for plant access control by enhancing security, reducing unauthorized access, automating attendance management, and supporting Industry 4.0 smart manufacturing environments through real-time intelligent surveillance and authentication [14], [15], [19].

Detecting Hazardous Gases

Detecting hazardous gas leakage is a critical requirement in industrial environments such as oil and gas facilities, chemical plants, refineries, and manufacturing industries because accidental gas releases can lead to fires, explosions, environmental pollution, and severe health hazards for workers [21], [22].

Conventional gas detection systems typically rely on fixed-point gas sensors that continuously measure gas concentrations at specific locations. Although these sensors provide accurate concentration measurements, they are limited by restricted spatial coverage, high installation costs, and their inability to visually identify the exact source of a leak [23].

Recent advances in computer vision and deep learning have enabled the development of vision-based gas leakage detection systems using the You Only Look Once (YOLO) family of object detection algorithms [4]. In these systems, optical gas imaging (OGI) cameras, thermal infrared cameras, or hyperspectral imaging devices capture gas plume images, while YOLO models automatically detect and localize the leakage regions in real time [22], [25].

Unlike traditional object detection methods that require multiple processing stages, YOLO performs object localization and classification simultaneously in a single forward pass, allowing high detection accuracy with low computational latency suitable for industrial monitoring applications [4]. Recent YOLO versions, including YOLOv8 and YOLOv11, incorporate improved feature extraction networks, feature pyramid structures, and attention mechanisms that significantly enhance the detection of small, low-contrast, and irregularly shaped gas plumes under varying environmental conditions such as wind, lighting, and background clutter [24], [25].

During model development, annotated datasets containing gas leakage images captured under different weather conditions, camera distances, and gas release rates are used to train the network to recognize diverse leakage patterns [22], [25]. Furthermore, data augmentation techniques such as random scaling, image rotation, brightness adjustment, Mosaic augmentation, and horizontal flipping improve model robustness and reduce overfitting, enabling reliable performance on previously unseen data [4].

Integrating YOLO with industrial monitoring systems allows automatic alarm generation, continuous remote surveillance, rapid localization of gas leaks, and immediate emergency response, thereby reducing accident risks, minimizing production downtime, and improving workplace safety [21], [23]. Consequently, YOLO-based vision systems have become a promising technology for hazardous gas leakage detection because they combine real-time performance, high detection accuracy, and efficient deployment within Industry 4.0 intelligent monitoring infrastructures [22], [24], [25].

Performance Comparison of Routing Protocols in terms of Convergence, Security, Scalability, Speed and Simplicity.

This section sheds light on the challenges of applying YOLO in the oil and gas industry.
Hardware and Computational Resource Intensiveness One of the major challenges associated with deploying the You Only Look Once (YOLO) object detection framework is its hardware and computational resource intensiveness, particularly in real-time industrial applications where edge devices and embedded systems are commonly used [26], [27]. Although recent YOLO versions, such as YOLOv8 and YOLOv11, have significantly improved detection accuracy and inference speed, these improvements often come at the cost of increased computational complexity, memory consumption, and graphics processing unit (GPU) requirements [26]. Training YOLO models typically requires high-performance GPUs with substantial memory capacity, while inference on high-resolution images demands considerable computational power to maintain real-time performance [28].

These requirements can limit the deployment of YOLO in industrial environments where edge devices, programmable logic controllers (PLCs), or low-power embedded computers have restricted processing capabilities and energy budgets [29]. Furthermore, industrial applications such as battery defect inspection, hazardous gas leakage monitoring, and facial recognition for access control often require continuous operation, making high energy consumption and hardware costs important considerations during system implementation [27], [29]. In addition, processing multiple camera streams simultaneously can significantly increase memory utilization and inference latency, thereby reducing the scalability of YOLO-based monitoring systems [30].

Several strategies have been proposed to address these limitations while maintaining acceptable detection performance. One effective solution is the use of lightweight YOLO variants, such as YOLO-Nano, YOLO-Tiny, or other compact architectures, which reduce the number of model parameters and floating-point operations without substantially sacrificing detection accuracy [26], [30]. Model compression techniques, including network pruning, weight quantization, and knowledge distillation, further decrease computational requirements by eliminating redundant parameters and reducing numerical precision while preserving most of the model's predictive capability [6].

Additionally, deploying YOLO models on specialized artificial intelligence accelerators, such as NVIDIA Jetson devices, Google Coral Tensor Processing Units (TPUs), or Intel OpenVINO-enabled hardware, can significantly improve inference speed while lowering power consumption [31], [32]. Optimizing the inference pipeline through image resizing, batch processing, and region-of-interest (ROI) detection also reduces computational overhead by limiting unnecessary processing [27]. Consequently, combining lightweight architectures, model optimization techniques, and dedicated AI hardware provides a practical solution for overcoming the hardware and resource challenges of YOLO, enabling efficient and reliable deployment in real-time industrial applications while maintaining high detection accuracy [25], [30], [32].

Strong Dependency on Large, High-quality Annotated Datasets

Another significant challenge in deploying the You Only Look Once (YOLO) object detection framework is its strong dependency on large, high-quality annotated datasets and the class imbalance commonly encountered in real-world applications [33], [34]. YOLO is a supervised deep learning model whose detection performance relies heavily on the quantity, diversity, and quality of the training data. In many industrial applications, such as battery leakage detection, hazardous gas monitoring, and manufacturing defect inspection, abnormal events occur far less frequently than normal operating conditions, resulting in highly imbalanced datasets [35].

Consequently, the model is exposed to substantially more examples of the majority class during training, causing it to become biased toward frequently occurring objects while exhibiting poor detection performance for rare but safety-critical defects [36]. This imbalance often leads to lower recall, increased false-negative rates, and reduced localization accuracy for minority classes, which may have severe consequences in safety-critical environments where missing a single defective object or hazardous event is unacceptable [34], [35]. Furthermore, collecting and manually annotating sufficient samples of rare events is both time-consuming and expensive because these defects naturally occur infrequently and often require expert labeling [36]. Variations in illumination, object orientation, background complexity, occlusion, and camera viewpoints further increase the amount of data required for effective model generalization, making dataset preparation one of the most resource-intensive stages of YOLO deployment [36], [37].

Several techniques have been proposed to address data dependency and class imbalance while maintaining high detection accuracy. Data augmentation methods, including Mosaic augmentation, MixUp, random cropping, image rotation, flipping, color jittering, and scaling, artificially increase dataset diversity and improve model robustness without requiring additional data collection [33], [38]. Synthetic data generation using generative adversarial networks (GANs) or simulation environments can further increase the number of minority-class samples, enabling the model to learn rare defect characteristics more effectively [37].

In addition, specialized loss functions such as Focal Loss reduce the influence of easily classified majority-class samples while assigning greater importance to difficult and underrepresented classes, thereby improving minority-class detection performance [5]. Oversampling minority classes, balanced batch sampling, and active learning strategies can also help create more representative training datasets by prioritizing informative and difficult samples during model training [33], [35]. Consequently, combining comprehensive data augmentation, synthetic data generation, balanced sampling techniques, and improved loss functions provides an effective solution for mitigating data dependency and class imbalance, enabling YOLO models to achieve more reliable and robust object detection performance in industrial and safety-critical applications [33], [36], [39].

III. CONCLUSION

The You Only Look Once (YOLO) framework has established itself as one of the most effective deep learning algorithms for real-time object detection in industrial applications. Its ability to perform object localization and classification simultaneously enables high detection accuracy while maintaining the speed required for continuous monitoring and automation. This makes YOLO particularly suitable for safety-critical and quality assurance tasks where rapid decision-making is essential.

In industrial environments, YOLO has demonstrated significant potential in a variety of applications, including battery leakage detection, facial detection for plant access control, and hazardous gas leakage monitoring. These applications contribute to improved product quality, enhanced workplace safety, strengthened security, and reduced reliance on manual inspection. By enabling automated, real-time monitoring, YOLO supports the transition toward intelligent manufacturing systems and Industry 4.0 technologies.

Despite its numerous advantages, YOLO also presents several challenges that must be addressed to maximize its effectiveness. The computational resources required for training and deployment can limit its implementation on low-power or edge devices. In addition, the model's performance depends heavily on the availability of large, diverse, and well-annotated datasets. Class imbalance, where defective or hazardous events occur much less frequently than normal conditions, may also reduce detection accuracy for critical objects if not properly managed.

Fortunately, these limitations can be mitigated through the adoption of lightweight YOLO architectures, model compression techniques, hardware acceleration, and optimized inference strategies that reduce computational demands while maintaining acceptable performance. Likewise, data augmentation, synthetic data generation, balanced sampling methods, and advanced training techniques can improve model robustness and enhance the detection of underrepresented classes.

Overall, YOLO provides a powerful and flexible solution for modern industrial vision systems by combining speed, accuracy, and adaptability. As deep learning technologies continue to evolve, future improvements in model efficiency, edge computing capabilities, and intelligent training methodologies are expected to further enhance YOLO's performance and broaden its adoption across industrial sectors. Consequently, YOLO is expected to remain a key enabling technology for smart manufacturing, automated inspection, industrial safety, and intelligent surveillance systems in the years ahead.

Acknowledgements

We would like to express our sincere thanks and gratitude to KFUPM management for guiding and supporting us to implement the performance test in out lab in order to enhance and optimize the company operations in the future. Moreover, special thanks for the planning and performance group for supporting us by funding the idea and providing the required paperwork to proceed with the performance test.

REFERENCES

1. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.
2. A. S. Mohammed, M. A. Khan, and R. Alhamid, "Applications of Deep Learning in Smart Manufacturing: A Review," IEEE Access, vol. 10, pp. 88412–88437, 2022.
3. R. Szeliski, Computer Vision: Algorithms and Applications, 2nd ed. Cham, Switzerland: Springer, 2022.
4. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
5. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788.
6. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
7. G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2024.
8. H. He and E. A. Garcia, "Learning from Imbalanced Data," IEEE Transactions on Knowledge and Data Engineering, vol. 21, no. 9, pp. 1263–1284, 2009.
9. H. Zhou, Y. Yu, K. Wang, and Y. Hu, "A YOLOv8-Based Approach for Real-Time Lithium-Ion Battery Electrode Defect Detection with High Accuracy," Electronics, vol. 13, no. 1, Art. no. 173, 2024.

10. D.-J. Kang, S.-Y. Son, and D.-H. Choi, "Battery Pouch Cell Defect Detection System Using YOLOv11 Lightweight Model," Journal of Electrochemical Science and Technology, 2025.
11. N. Choudhary et al., "Autonomous Visual Detection of Defects from Battery Electrode Production Using Deep Learning," Advanced Engineering Materials, vol. 24, no. 5, 2022.
12. L. Deng, D. Kang, and Q. Liu, "SLI_FRM: Multi-scale Real-Time Detection of Lithium Battery Swelling and Early Smoke Based on YOLOv8," Case Studies in Thermal Engineering, vol. 67, 2026.
13. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
14. A. K. Jain, A. Ross, and S. Prabhakar, "An Introduction to Biometric Recognition," IEEE Transactions on Circuits and Systems for Video Technology, vol. 14, no. 1, pp. 4–20, Jan. 2004.
15. J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 4690–4699.
16. O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep Face Recognition," in Proc. British Machine Vision Conference (BMVC), 2015.
17. P. Viola and M. Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2001, pp. I-511–I-518.
18. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
19. G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2024. Available: https://github.com/ultralytics/ultralytics
20. F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 815–823.
21. M. R. R. Khan, M. M. Rahman, and M. A. Hossain, "Industrial Gas Leakage Detection and Monitoring Systems: A Review," Journal of Loss Prevention in the Process Industries, vol. 74, Art. no. 104640, 2022.
22. A. P. Merkus and H. F. J. M. Van den Boom, "Optical Gas Imaging for Industrial Methane Leak Detection: A Review," Sensors, vol. 23, no. 5, Art. no. 2556, 2023.
23. A. C. Tolias, G. K. Papadopoulos, and E. N. Zois, "Gas Leak Detection Technologies in Industrial Environments: A Comprehensive Review," Process Safety and Environmental Protection, vol. 159, pp. 122–142, 2022.
24. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
25. G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2024. Available: https://github.com/ultralytics/ultralytics
26. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
27. G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2024. Available: https://github.com/ultralytics/ultralytics
28. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788.
29. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.
30. 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," arXiv preprint, arXiv:2207.02696, 2022.
31. S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," in Proc. International

Conference on Learning Representations (ICLR), 2016.
32. NVIDIA Corporation, "NVIDIA Jetson Platform for Edge AI and Robotics," NVIDIA Technical Documentation, 2024.
33. J. Terven and D. Cordova-Esparza, "A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS," arXiv preprint, arXiv:2304.00501, 2023.
34. G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," GitHub Repository, 2024. Available: https://github.com/ultralytics/ultralytics.
35. H. He and E. A. Garcia, "Learning from Imbalanced Data," IEEE Transactions on Knowledge and Data Engineering, vol. 21, no. 9, pp. 1263–1284, Sept. 2009.
36. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.
37. T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," in Proc. IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2980–2988.
38. 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," arXiv preprint, arXiv:2207.02696, 2022.
39. A. Goodfellow et al., "Generative Adversarial Nets," in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2014, pp. 2672–2680.