

YOLOv8 computer vision for automated offside detection in professional football validated through supervised learning

Received: 9 January 2026

Accepted: 6 May 2026

Published online: 11 May 2026

Cite this article as: Abdel-Fattah A., Belhaouari S.B., Ceylan H.İ. *et al.* YOLOv8 computer vision for automated offside detection in professional football validated through supervised learning. *Sci Rep* (2026). <https://doi.org/10.1038/s41598-026-52668-4>

Abdul-Rahman Abdel-Fattah, Samir Brahim Belhaouari, Halil İbrahim Ceylan, Serdar Bayrakdaroglu, Raul Ioan Muntean, Wissem Dhahbi, Ismail Dergaa & Abdelfatteh El Omri

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

YOLOv8 Computer Vision for Automated Offside Detection in Professional Football Validated Through Supervised Learning

Abdul-Rahman Abdel-Fattah ¹, Samir Brahim Belhaouari ¹, Halil İbrahim Ceylan², Serdar Bayrakdaroglu³, Raul Ioan Muntean⁴, Wissem Dhahbi ^{5,6}, Ismail Dergaa ^{7,8}, Abdelfatteh El Omri ^{9,10}

¹ Division of Information and Computing Technology, College of Science and Engineering, Hamad Ben Khalifa University, Doha, Qatar. abab46015@hbku.edu.qa , sbelhaouari@hbku.edu.qa

²Faculty of Sports Sciences, Atatürk University, Erzurum 25240, Türkiye. halil.ibrahimceylan60@gmail.com

³Faculty of Sports Sciences, Gümüşhane University, Gümüşhane 29100, Türkiye. bayrakdaroglu85@gmail.com

⁴Department of Physical Education and Sport, Faculty of Law and Social Sciences, University “1 Decembrie 1918” of Alba Iulia, Alba Iulia, Romania. muntean.raul@uab.ro

⁵ Research Unit “Sport Sciences, Health and Movement”, High Institute of Sports and Physical Education of Kef, University of Jendouba, Kef 8100, Tunisia.

⁶ Training Department, Police College, Qatar Police Academy, Doha 59J5+C7R, Qatar. wissem.dhahbi@gmail.com

⁷ High Institute of Sport and Physical Education of Ksar Said, University of Manouba, Manouba 2010, Tunisia

⁸ Physical Activity Research Unit, Sport and Health (UR18JS01), National Observatory of Sports, Tunis 1003, Tunisia. Phd.dergaa@gmail.com

⁹ Surgical Research Section, Department of Surgery, Hamad Medical Corporation, Doha, Qatar

¹⁰ QU-Health, Qatar University, Doha, Qatar. AElOmri@hamad.qa

***Correspondence:** halil.ibrahimceylan60@gmail.com or muntean.raul@uab.ro

ABSTRACT

Aim: Offside calls are integral to maintaining fairness and competitive integrity in football, yet current officiating processes are constrained by limitations in accuracy despite technological advancements. This investigation developed and validated an automated offside detection system using YOLOv8 object detection architecture to address documented limitations of Video Assistant Referee (VAR) systems in offside decision accuracy. **Methods:** Supervised learning validation employed 400 annotated instances from professional match footage. The YOLOv8 medium variant was configured for four-class detection, with a focus on extreme player positions, reducing computational complexity by approximately 70% relative to comprehensive player tracking while accepting accuracy trade-offs in marginal scenarios. The system was optimized for computational efficiency through highly accurate tracking of player positions. Data were partitioned into 40% (n=160) for training and 60% (n=240) for testing, with 5-fold cross-validation. **Results:** The automated system achieved robust performance metrics across 240 test instances: overall accuracy of 83.0%, precision of 85.0%, recall of 87.0%, and F1-score of 86.0%. Confusion matrix analysis revealed true positives (n=130), true negatives (n=70), false positives (n=22), and false negatives (n=18). Individual case analysis demonstrated 80% agreement with ground truth classifications. System performance significantly exceeded the random classification baseline ($p < 0.001$). The strategic extreme player tracking approach achieved substantial computational efficiency while maintaining decision accuracy across diverse match scenarios. **Conclusion:** A YOLOv8-based offside detection system demonstrates proof-of-concept feasibility for decision-support applications in football officiating. Performance (83% accuracy, 95% CI: 79-87%) on 240 test instances remains below professional benchmarks (>93%), with single-camera constraints and limited dataset size restricting generalizability to proof-of-concept validation rather than operational deployment.

KEYWORDS: Artificial intelligence; machine learning; neural networks; offside detection; real-time object detection; sports officiating technology; video assistant referee; YOLOv8

1. INTRODUCTION

Competitive football officiating requires 137 observable decisions per match at the international level, with offside calls accounting for 15-20% of these decisions and directly affecting scoring opportunities [1]. Among these critical decisions, offside infractions are among the most challenging and controversial aspects of football officiating, requiring split-second judgments in fast-paced, dynamic environments where human perceptual limitations frequently compromise decision accuracy.

The complexity of offside detection stems from multiple interrelated factors that challenge experienced officials. Assistant referees in elite European competitions (UEFA Champions League, English Premier League, 2015-2019 seasons) achieve 94% offside accuracy (95% CI: 92-96%, $n=1,247$ decisions coded via video analysis with inter-rater reliability $\kappa=0.91$) [2]. This benchmark remains consistent with recent evaluations reporting 93-95% accuracy pre-VAR implementation, with error concentration in marginal violations (<30 cm player separation from offside threshold). Impact quantification across 760 top-tier league matches (2017-2020) indicates that erroneous offside calls affect 8-12 goals per season, with documented instances that alter final league standings (2-3 position shifts) and European competition qualification in five seasons across major leagues. Most of these errors stem from the limits of human perception. Assistant referees are required to track several players at once, follow the ball's path, and judge the exact moment of contact—all while staying in the right position. Research shows that crowded situations, in which many players cluster, make these judgments even more difficult, and that players who are not directly involved can still influence decisions through their positioning or gestures [3]. External factors also play a role: crowd noise, for example, has been shown to reduce foul-detection rates against home teams by 15.5% relative to silent conditions, highlighting how auditory distractions can impair focus and decision-making [4]. In addition, the pressure of high-stakes matches adds further cognitive strain, often leading to systematic biases in refereeing decisions.

Video Assistant Referee (VAR) technology, integrated into Laws of the Game in 2018, encompasses distinct architectural categories [5]. Conventional VAR employs multi-angle video review with human adjudication across all decision types (goals, penalties, red cards, mistaken identity), achieving 93-96% overall accuracy in professional competitions [6]. Semi-automated offside technology (SAOT), deployed in the FIFA World Cup 2022 and the UEFA Champions League 2023-present, uses 12-29 dedicated tracking cameras (minimum 50 fps) with limb-detection algorithms (29 body points per player, 3 cm spatial tolerance) to achieve $>98\%$ offside-specific accuracy. Fully automated systems eliminating human oversight remain absent from professional competition due to rule complexity (player involvement assessment) and liability constraints. A meta-

analysis of 20 competitions (9,871 matches across all VAR decision categories) identified persistent error mechanisms: communication protocol failures, compressed decision timeframes, and camera angle limitations [7]. Offside decisions constituted 18-22% of VAR interventions across competitions. A well-known example occurred in September 2023 during Tottenham Hotspur's match against Liverpool, when hurried communication between officials led to a valid goal being disallowed [8]. Furthermore, VAR's reliance on camera angles can create blind spots, making it less effective for close offside calls, especially when players obstruct the view or the exact moment of ball contact cannot be determined with certainty. Operational challenges extend beyond classification accuracy to decision latency (conventional VAR: 60-180 seconds per review), transparency (opaque adjudication processes), and public acceptance (52% fan approval in European surveys, 2021-2023). The proposed YOLOv8-based approach addresses a distinct niche: single-camera broadcast integration without specialized hardware, trading SAOT-level precision for accessibility in resource-constrained competitive contexts.

Advances in computer vision have created new possibilities for automating sports officiating through real-time object detection and tracking. In particular, deep learning models such as the You Only Look Once (YOLO) family have shown strong performance in identifying and classifying multiple objects in complex visual settings [9]. The evolution from YOLOv1 to YOLOv8 has led to steady improvements in detection accuracy, processing speed, and efficiency, making real-time analysis in sports increasingly practical [10]. Research on automated offside detection has explored contour-based silhouette mapping, background subtraction with Kalman filtering, and ensemble-classifier approaches [11]. However, these methods exhibit critical limitations: contour-based techniques achieve only 76-81% accuracy on standardized datasets and fail under player-occlusion conditions; ensemble approaches require computationally prohibitive multi-model inference, which is unsuitable for real-time processing; and most validations use synthetic datasets or controlled laboratory conditions rather than broadcast match footage. Dataset constraints remain severe, with typical studies employing fewer than 200 annotated instances, which is insufficient for robust generalization in deep learning. Architecturally, contour-based methods rely on handcrafted feature extraction and threshold-based spatial logic, rendering them brittle under variable illumination and player occlusion; ensemble classifiers achieve marginal accuracy gains at the cost of multi-model inference latency exceeding 200 ms, which precludes real-time deployment. By contrast, YOLOv8's anchor-free convolutional architecture performs end-to-end multi-class detection (players, ball, officials) within a single inference pass of 19-21 ms, directly addressing the occlusion sensitivity and latency constraints of predecessor methods. Unlike SAOT, which requires

proprietary arrays of 12-29 calibrated cameras per stadium, the proposed system operates on standard 1080p broadcast feeds, eliminating hardware barriers in resource-constrained competitive contexts at the accepted cost of reduced spatial precision. These methodological gaps motivated the adoption of YOLOv8's single-shot detection architecture, which balances accuracy and computational efficiency while enabling training on limited real-world data.

Recent YOLOv8 deployments in football contexts include camera operator segmentation for broadcast analysis [16] and multi-person tracking for tactical evaluation, demonstrating the feasibility of real-time detection (45-60 fps on consumer GPUs). However, these applications omit the implementation of the offside rule and lack validation against referee ground truth. Methodological approaches span player segmentation (instance masks for spatial localization), multi-object tracking (temporal trajectory estimation), and field geometry projection (homography-based coordinate transformation). Limitations include depth ambiguity in monocular vision, occlusion sensitivity in crowded scenarios, and the absence of an assessment of player involvement required under FIFA Law 11.

YOLOv8 is the latest generation of single-shot object detection models, introducing key upgrades such as an anchor-free detection system, improved convolutional layers, and refined training strategies that yield improved accuracy and speed trade-offs compared with earlier versions [12]. It can handle 1080p video at around 50 frames per second while preserving high detection accuracy, making it well-suited for real-time sports applications [13]. Studies have further confirmed its reliability in challenging scenarios; for example, Hermens et al. [14] reported that YOLOv8 achieved near-perfect detection accuracy even when trained on relatively small datasets of 100-350 images.

Although VAR and computer vision technologies have advanced considerably, there is still no fully automated offside-detection system that delivers the accuracy and reliability required at the professional level. Studies show that human officials perceive ball contact with an average delay of about 132 milliseconds, which can introduce systematic bias in close offside calls [14]. The research gap comprises three dimensions: (i) absence of validated single-camera offside detection at broadcast frame rates (30-60 fps) without specialized hardware, addressing accessibility barriers in sub-elite competitions; (ii) lack of performance quantification on operational broadcast footage under authentic match conditions rather than controlled laboratory datasets; (iii) undefined accuracy-cost trade-off boundaries for contexts where SAOT infrastructure (\$2-4 million per stadium) is prohibitive. The primary objective is to validate computational efficiency for resource-constrained implementation while not exceeding the precision of existing SAOT capabilities. Existing solutions require either multi-camera infrastructure with proprietary hardware (semi-automated offside technology systems) or sacrifice

accuracy for speed (single-model approaches). This gap constrains access to officiating support technology across competitive levels and impedes the systematic evaluation of automated detection in resource-constrained environments. This study addresses this issue by adapting YOLOv8's object detection capabilities for football analysis. The approach emphasizes detecting extreme player positions and tracking ball movement, reducing computational demands while preserving decision accuracy. The main contributions are threefold: (i) a streamlined algorithm that targets key player positions instead of tracking the entire field, (ii) integration of temporal analysis to capture ball dynamics and passing direction, and (iii) a thorough evaluation using real match footage to demonstrate its potential for professional use. The system is positioned as a decision-support prototype for evaluation rather than a replacement for existing officiating infrastructure.

2. METHODS

2.1 Study Design

This study evaluates a deep learning-based offside detection system using supervised learning and held-out test validation on the YOLOv8 architecture. Annotated professional match footage provided training and testing data, with validation protocols following established computer vision benchmarking standards adapted for sports officiating contexts.

2.2 Dataset Collection and Characteristics

The dataset consisted of 400 high-resolution images and video clips from professional football matches, captured under various playing conditions. Broadcast footage was used to ensure coverage across different stadiums, lighting environments, and camera angles typically found in competitive settings. Only footage with clear visibility of players, the ball, and field markings was included, with a minimum resolution of 1080p to provide sufficient detail for accurate detection and classification. Eligible samples contained well-defined offside situations, including unobstructed views of key players and ball position during active play sequences such as passes and positional shifts [15]. Footage was excluded if it exhibited excessive motion blur, significant occlusion, poor lighting, or originated from non-competitive contexts (Table 1). The final sample size of 400 annotated instances was informed by standard computer vision research practices and pilot testing. Sample size determination followed pragmatic constraints of annotation resource availability rather than formal power calculations, which are not standard practice for deep learning classifier development, where performance bounds depend on model capacity, data complexity, and training convergence rather than predetermined effect size detection [16]. This sample size, while enabling initial proof-of-concept validation, imposes a constraint on claims of generalization in deep learning. Dataset limitations restrict

performance estimates to approximate bounds rather than precise population-level accuracy projections.

****Please Insert Table 1 About Here****

2.3 Data Preprocessing and Annotation Protocol

Data preprocessing followed standard computer vision practices to maintain consistency across the dataset. All images were resized to 640×640 pixels to match YOLOv8 requirements while retaining essential visual detail. Color normalization was applied by converting to RGB format and scaling pixel values to [0,1], ensuring stable training across varied lighting conditions and camera setups [17]. An annotation framework defined four object categories: players from each team, the ball, and match officials, using bounding boxes with precise coordinates. To ensure reliability, two independent annotators with sports science backgrounds labeled all instances, achieving Cohen's kappa $\kappa = 0.87$ (95% CI: 0.83-0.91). Discrepancies ($n = 52$, 13% of the dataset) were resolved through structured consensus discussions among annotators, and the remaining disagreements ($n = 8$, 2% of the dataset) were adjudicated by a certified UEFA B-licensed referee with 12 years of professional match analysis experience, whose decisions established the final ground-truth classifications. To increase dataset diversity and improve generalization, several augmentation techniques were applied, including random horizontal flips ($p = 0.5$), rotations of up to $\pm 15^\circ$, brightness shifts of $\pm 20\%$, and contrast changes of $\pm 15\%$. These transformations introduced natural variation while preserving spatial relationships needed for offside analysis. Mosaic augmentation was used during training but removed in the final 10 epochs to reduce the risk of overfitting.

2.4 YOLOv8 Model Implementation and Training

The system was built using the medium YOLOv8 (YOLOv8m) model, chosen for its balance between detection accuracy and computational efficiency in real-time scenarios [18]. The architecture included anchor-free detection, improved convolutional blocks, and optimized feature extraction layers, with custom output heads designed for four-class detection and weighted to account for class imbalance [19]. Training parameters were optimized through a grid search, with the final configuration set to a learning rate (η) of 0.01, batch size of 16, and 300 epochs, using early stopping when validation loss plateaued. The Adam optimizer with weight decay ($\lambda = 0.0005$) was used to reduce overfitting. The data were split using stratified sampling: 40% ($n=160$) for training and 60% ($n=240$) for testing, with 5-fold cross-validation applied to the training partition for model optimization. Transfer learning was employed by initializing the model with YOLOv8 weights pre-trained on COCO, followed by fine-tuning for football-specific detection. After initial convergence, all network layers were

unfrozen and retrained with a lower learning rate ($\eta = 0.001$) to refine performance [20].

2.5 Team Classification and Player Identification

Team classification methodology employed advanced color segmentation techniques combined with machine learning algorithms to distinguish between opposing teams. The approach utilized HSV color space analysis to extract dominant jersey colors, followed by clustering algorithms (K-means, $k=2$) to separate teams based on chromatic characteristics, with validation procedures including manual verification across diverse lighting conditions and jersey color combinations. Player identification protocols focused on extracting the spatial coordinates of detected players within the field coordinate system, using geometric transformation techniques to convert pixel coordinates to normalized field positions via perspective projection. Kalman filtering was implemented to smooth player position trajectories and handle temporary occlusions during match play.

2.6 Automated Offside Detection Algorithm

The offside detection algorithm prioritized computational efficiency by using extreme player tracking (the leftmost and rightmost positions per team) rather than comprehensive field coverage. This approach assumes the second-to-last defender occupies an extreme horizontal position, reducing computational load by approximately 70% relative to full-field tracking but introducing potential spatial errors when defensive lines compress centrally or camera perspectives create depth ambiguity. This strategic reduction in tracking complexity significantly improved processing efficiency while maintaining decision accuracy for offside determination. Ball detection and trajectory analysis incorporated temporal information to determine pass direction and timing through motion vector calculations, establishing ball velocity and direction, with pass detection thresholds calibrated based on empirical analysis of ball movement patterns in professional match footage. The core offside detection algorithm implemented official FIFA offside rules through computational logic, evaluating player positions relative to the second-last opponent at the moment of ball contact and incorporating temporal analysis to ensure accurate timing assessment, with decision logic including exception handling for situations where offside rules do not apply and confidence scoring for decision reliability assessment.

Recognized limitations of extreme player tracking include: (i) failure to detect offside violations when the second-to-last defender occupies a central rather than extreme position, (ii) systematic spatial errors under diagonal camera angles that distort horizontal positioning, (iii) reduced accuracy in compressed defensive formations where multiple defenders cluster near the offside threshold, and (iv) potential misclassification when goalkeeper positioning creates non-standard defensive geometry. These

constraints represent deliberate trade-offs between the feasibility of real-time processing and the spatial completeness, with acknowledged accuracy costs in geometrically complex matching scenarios. No direct performance comparison with full-field player tracking was conducted due to resource constraints, preventing empirical quantification of the accuracy trade-offs between extreme position heuristics and comprehensive spatial coverage.

****Please Insert Figure 1 About Here****

The algorithm employs a streamlined approach that focuses on extreme player positions (leftmost and rightmost per team) rather than comprehensive player tracking, thereby significantly reducing computational complexity while maintaining decision accuracy. Parallel processing of player identification and ball trajectory analysis enables real-time performance suitable for live match implementation (Figure 1).

2.7 Performance Evaluation and Validation

Model performance was evaluated using standard computer vision metrics, such as precision, recall, F1-score, and mean Average Precision (mAP), as well as football-specific measures, including offside decision accuracy, false positive rate, and temporal consistency across video sequences. To ensure robustness, a five-fold cross-validation ($k=5$) approach was used, and additional test datasets from different leagues, stadiums, and playing conditions were included to assess generalizability [21]. Ground-truth labels were created by certified referees and match analysts, and inter-observer agreement was measured using Fleiss' kappa to provide a benchmark for human decision accuracy.

2.8 Statistical Analysis

Statistical analysis was conducted in Python 3.9 using SciPy 1.11.0, NumPy 1.24.0, and scikit-learn 1.3.0. For performance comparisons, paired t-tests were used for normally distributed data, whereas Wilcoxon signed-rank tests were used for non-normally distributed data. Effect sizes were reported using Cohen's d where appropriate. Ninety-five percent confidence intervals were estimated using bias-corrected accelerated (BCa) bootstrap resampling (1,000 iterations); no Bayesian inference or Markov Chain Monte Carlo procedures were employed. BCa resampling was selected for its greater robustness to distributional skewness compared with standard percentile bootstrap methods. Classifier discrimination was assessed through precision-recall trade-off analysis. Five-fold cross-validation on the training partition informed hyperparameter selection (learning rate, batch size, and augmentation parameters), with final performance evaluated on the held-out test set to prevent overfitting and ensure unbiased accuracy estimation. Statistical significance was set at $\alpha = 0.05$, and

Bonferroni corrections were applied to adjust for multiple comparisons. The primary performance metric (overall accuracy: 83.0%) was computed exclusively on the held-out test set (n=240), which was withheld from all training and hyperparameter optimization procedures. The 5-fold cross-validation applied to the training partition (n=160) served solely for hyperparameter selection; fold-level performance metrics are reported in Table 3.

3. RESULTS

3.1 Overall System Performance and Statistical Validation

The trained model achieved 83.0% accuracy (95% CI: 79.1-86.9%) on 240 test instances, with a precision 85.0% (95% CI: 81.3-88.7%), a recall 87.0% (95% CI: 83.5-90.5%), and an F1-score 86.0% (95% CI: 82.7-89.3%). Five-fold cross-validation on the training partition (n=160) yielded a mean cross-validated accuracy of $83.3 \pm 1.0\%$ across folds (Table 3), confirming convergence stability prior to held-out test evaluation. Precision reached 85.0%, meaning that about 85 out of every 100 flagged offsides were correct, while recall was 87.0%, reflecting the system's ability to reliably detect true violations. The resulting F1-score of 86.0% highlights a balanced trade-off between sensitivity and specificity, confirming the system's potential for use in automated officiating.

3.2 Confusion Matrix Analysis and Classification Performance

Confusion matrix analysis provided insight into the system's classification performance (Table 2). The largest group was true positives (n = 130), showing the model's strong ability to detect genuine offside violations. True negatives (n = 70) confirmed reliable recognition of legal play, preventing unnecessary stoppages. False positives (n = 22) occurred primarily in marginal cases in which players were positioned very close to the offside line, a scenario that often challenges even experienced referees. The relatively low number of such errors reflects a conservative decision-making bias that favors accuracy in uncertain situations. False negatives (n = 18) were largely attributable to temporary occlusion or rapid movements that made precise timing difficult to determine, underscoring the inherent limitations of single-camera analysis. Class-specific analysis derived from the confusion matrix reveals asymmetric performance across decision categories. For the offside class: precision 85.5%, recall 87.8%, F1-score 86.6%. For the not-offside class: precision 79.5%, recall 76.1%, F1-score 77.7%. The macro-average F1-score was 82.2%, and the weighted F1-score was 83.3%, consistent with the reported overall accuracy. The 8.9-percentage-point F1 differential between classes reflects the system's stronger capacity to identify genuine violations than to confirm legal play, attributable to the class imbalance in the test set (offside: n=148; not-offside: n=92).

****Please Insert Table 2 About Here****

****Please Insert Table 3 About Here****

3.3 Visual Performance Analysis and Model Training Validation

Training and validation results showed clear convergence without signs of overfitting, with both accuracy curves progressing in parallel (Figure 2). This indicates that the model learned generalizable features rather than memorizing training data. On real match footage, the system consistently detected all object categories with high confidence under normal conditions. The main challenge arose in occlusion scenarios, such as when the ball was partially obscured by players or when player positioning blocked key elements, which reduced detection confidence and occasionally led to errors (Figure 3).

****Please Insert Figure 2 About Here****

****Please Insert Figure 3 About Here****

3.4 Practical Application Demonstration and Real-World Validation

The system's practical value was shown by testing it on real offside situations from professional matches. It successfully detected attacking players beyond the defensive line at the moment of ball contact, accurately identified teams, tracked the ball, and analyzed the spatial positions required for offside decisions (Figure 4). It also correctly identified situations without violations, even when player positioning could have created visual uncertainty for referees, demonstrating its reliability across both positive and negative decisions (Figure 5).

****Please Insert Figure 4 About Here****

****Please Insert Figure 5 About Here****

3.5 Error Analysis and Technical Limitations

Error analysis showed that false positives ($n = 22$) primarily occurred in marginal cases in which players were positioned very close to the offside line, making precise spatial measurement especially difficult. False negatives ($n = 18$) were primarily due to visual occlusion, in which defenders temporarily blocked the attackers' view, with additional errors arising from difficult camera angles or lighting conditions. Importantly, errors were evenly distributed across teams and field positions, indicating that the limitations arose from measurement challenges rather than algorithmic bias. The extreme player-tracking approach likely contributed to errors in matches featuring compressed defensive lines (estimated 15-20% of professional match scenarios) or non-standard camera perspectives, in which the assumption that offside thresholds align with horizontally extreme players may be violated.

3.6 Processing Efficiency and Implementation Feasibility

Processing efficiency: NVIDIA RTX 3080 hardware (10 GB VRAM, consumer-grade GPU) achieved 48-52 fps on 1080p input, exceeding 30 fps broadcast integration threshold. Per-frame inference time: 19-21 ms (mean: 20 ms, SD: 1.2 ms), with player count variation (18-22 detected instances) producing <5% latency fluctuation. Operational deployment latency (frame acquisition through the broadcast infrastructure, preprocessing, inference, and output transmission to the referee communication system) remains unquantified; offline validation cannot verify real-time stability under live-match conditions with network transmission delays and concurrent processing demands. The efficiency gains enabled deployment on standard broadcasting equipment without requiring specialized high-performance computing infrastructure, reducing implementation barriers and costs for widespread adoption.

3.7 System Reliability and Performance Context

The overall system performance of 83% accuracy, 85% precision, 87% recall, and 86% F1-score demonstrates proof-of-concept feasibility for automated detection, though it remains below the accuracy of professional assistant referees (~94%) and VAR-assisted decision-making (93-98% in elite competitions). Balanced performance across precision and recall metrics indicates appropriate trade-offs between conservative and aggressive decision-making approaches, which are essential for the practical implementation of referee decision-support tools. The performance consistency demonstrated across diverse match scenarios supports the system's potential for professional implementation while acknowledging technical challenges inherent in automated visual analysis of dynamic sports scenarios.

4. DISCUSSION

This study demonstrates that YOLOv8-based automated offside detection achieves moderate accuracy suitable for decision-support prototyping in football officiating contexts. The system achieved an overall accuracy of 83.0%, precision of 85.0%, recall of 87.0%, and F1-score of 86.0%. These results represent a notable step forward in automated officiating, particularly given the efficiency gains from focusing on extreme player positions rather than tracking the entire field. Confusion matrix results further highlight this strength: true positives ($n = 130$) were the largest group, reflecting the model's ability to identify genuine offside violations, whereas the relatively low number of false positives ($n = 22$) indicates a conservative bias that favors accuracy in borderline cases.

The asymmetric class-specific performance carries distinct operational implications. False positives ($n=22$; 9.2% of test instances) result in wrongful disallowance of legal attacking

sequences, directly affecting scoring opportunities and constituting officiating errors equivalent to incorrect flag raises. False negatives ($n=18$; 7.5% of test instances) represent undetected violations that confer an illegitimate positional advantage to attacking players. Under the officiating principle of conservative decision-making, the higher FP cost relative to FN cost suggests that deployment thresholds should be calibrated toward higher precision at the acceptable cost of reduced recall. At the current performance level (not-offside F1: 77.7%), autonomous deployment in professional contexts is not operationally acceptable given the established human assistant referee benchmark of 93–94% [2]; however, the system meets a threshold appropriate for supplementary flagging in sub-elite contexts where no technological support is available.

The system's performance compares favorably with recent research on automated sports officiating and directly addresses key shortcomings of current VAR use. Studies show that VAR remains prone to human error, communication lapses, and time delays, which can undermine decision quality [22]. Harwood's motion-capture work, for example, found that referees perceive ball contact with an average delay of 132 milliseconds, thereby introducing systematic bias in marginal offside calls [14]. By using frame-by-frame analysis, the proposed system removes this subjective timing factor, thereby reducing a major source of error in VAR decisions. Compared with earlier computer vision approaches, which typically achieved 76–81% accuracy [11], the YOLOv8-based model reached 83% accuracy while also offering better computational efficiency.

The system's strong recall (87%) is largely due to YOLOv8's advanced feature extraction and its focus on tracking extreme player positions. The model's anchor-free detection and improved convolutional blocks enhance localization accuracy, even in fast-paced game scenarios with frequent occlusions [23]. The relatively low false positive rate (9.2%) reflects a conservative decision-making tendency, consistent with officiating principles that favor allowing play to continue in uncertain situations rather than risking an incorrect stoppage. This mirrors patterns observed in psychological studies of referees, who tend to avoid false positives that could unfairly disrupt scoring opportunities [4]. The balance between precision (85%) and recall (87%) indicates that the detection threshold was effectively tuned, avoiding the typical extremes of overly cautious or overly aggressive classification.

Relative to prior automated approaches, the YOLOv8-based system surpasses contour-based silhouette methods (76–81% [11]) by addressing their principal failure modes through learned convolutional feature representations rather than handcrafted segmentation, and outperforms ensemble classifiers operationally by reducing inference latency from >200 ms to 19–21 ms per frame. Compared with SAOT, the single-camera approach accepts a spatial precision trade-off (estimated depth error 10–25 cm versus SAOT's <3

cm through multi-camera triangulation) in exchange for accessibility in contexts where dedicated camera infrastructure is financially prohibitive. The system therefore occupies a distinct architectural niche: single-pass real-time detection without proprietary hardware, targeting competitive levels where VAR and SAOT deployment are economically or logistically unfeasible. Unlike multi-camera semi-automated offside technology (SAOT) systems deployed in elite competitions, which achieve >98% spatial accuracy through triangulated depth estimation and limb-tracking algorithms, the single-camera approach prioritizes accessibility and computational efficiency over precision. The performance gap reflects fundamental trade-offs between system complexity, implementation cost, and accuracy requirements across competitive levels.

The system's performance characteristics suggest potential application as a supplementary decision-support tool in non-critical officiating contexts, referee training programs, and competitive environments where VAR infrastructure is financially or logistically unavailable, thereby meeting key operational demands in sports technology adoption. Its 83% accuracy offers reliable referee assistance while preserving human officials' ultimate authority, consistent with technology-assisted officiating principles [24]. Computational efficiency through extreme player tracking enables deployment on standard broadcasting equipment, reducing implementation barriers relative to comprehensive systems and enhancing accessibility across competitive levels. Consistent performance across diverse scenarios and a balanced error distribution, indicating no systematic bias, supports fair, competitive applications crucial for governing bodies and stakeholders.

Hawk-Eye's semi-automated offside technology, deployed in FIFA World Cup and UEFA Champions League competitions, employs 12-29 high-frame-rate cameras (50 fps minimum) with limb-tracking algorithms that analyze 29 body points per player, achieving reported accuracy >99% for spatial positioning within 3 cm tolerance. This system serves as the operational-accuracy benchmark for professional offside detection but requires dedicated camera infrastructure, proprietary calibration software, and specialized personnel, with estimated implementation costs exceeding \$3 million per stadium. The current YOLOv8-based approach trades spatial precision for accessibility, utilizing standard broadcast feeds without additional hardware, thereby enabling feasibility assessment in contexts where SAOT deployment is economically prohibitive. Performance differences reflect this infrastructure trade-off: Hawk-Eye's multi-camera triangulation resolves depth ambiguity and occlusion constraints inherent to monocular vision systems, while the single-camera approach accepts higher error rates (17% vs. <1%) in exchange for reduced implementation barriers. The systems address distinct market segments rather than competing directly: SAOT for elite professional competitions requiring maximal accuracy, single-

camera detection for training environments, lower-tier leagues, and broadcast analysis, where cost constraints dominate accuracy requirements. The proliferation of AI-driven sensor technologies in sports necessitates standardized validation methodologies, algorithmic transparency, and performance benchmarking across competitive contexts [25]. Such frameworks would enable systematic comparison between emerging detection systems and established officiating technologies.

System positioning: The approach serves as a low-cost, single-camera offside-detection prototype for contexts lacking SAOT infrastructure, rather than as a VAR replacement or an accuracy enhancement for elite competitions. Target applications comprise (i) referee training programs providing offline decision analysis with performance feedback, (ii) sub-elite competitive levels (regional leagues, youth competitions) without VAR access, and (iii) broadcast analysis for viewer education. The system does not address VAR operational deficiencies (communication protocols, multi-camera review, player involvement assessment), but it offers accessibility in resource-limited contexts where established technologies are economically prohibitive. Operational deployment would require resolution of real-time processing constraints not addressed in offline validation. Decision latency from frame capture to output classification must remain below 200 ms to inform referee decisions before play resumes, necessitating optimized inference pipelines and protocols for integrating broadcast feeds. Frame synchronization between ball contact detection and player position assessment introduces temporal uncertainty (± 33 ms at 30 fps broadcast rates), potentially causing misclassification when players move 0.5-1.0 m between frames during sprint sequences. Communication protocols for conveying automated outputs to match officials without disrupting match flow, integration with existing referee communication systems, and procedures for adjudicating disputed decisions remain unspecified, posing implementation barriers beyond algorithmic performance validation.

The interpretation and generalizability of these findings are subject to several limitations. The dataset of 240 test instances, while enabling initial validation, is limited compared to the vast number of real-world offside events, potentially affecting generalizability to less frequent scenarios. The single-camera analysis, although computationally efficient, inherently suffers from perspective-based limitations that can lead to occlusions, as evidenced by 18 false negatives during such events, suggesting that multi-camera systems could enhance reliability [26]. The binary classification approach (offside/not offside) omits the nuanced assessment of player involvement that is crucial to modern interpretation of the offside rule under FIFA regulations, thereby restricting the system's applicability to scenarios that require subjective judgment beyond spatial analysis.

Integrating the current system with multi-camera setups offers an immediate avenue for improved performance, particularly by mitigating occlusion-related false negatives through advanced sensor fusion. Developing capabilities for player-involvement assessment, potentially using machine learning to analyze movement, gaze, and spatial relationships, would address a key limitation by enabling comprehensive enforcement of the offside rule [27]. Real-time implementation studies and field testing in professional matches are needed to validate performance under operational conditions and inform integration with existing infrastructure.

4.1 Limitations

This study's interpretation requires acknowledgment of several methodological constraints. Dataset constraints impose interpretation boundaries: (1) Sample size ($n=400$, 240 test instances) limits generalization relative to deep learning standards requiring 2,000-10,000 instances for robust performance estimation across diverse scenarios. (2) Selection bias: Inclusion criteria (clear player visibility, minimal occlusion, optimal lighting) excluded 30-40% of broadcast footage representing operationally challenging conditions (adverse weather, suboptimal camera angles, player clustering). This bias likely inflates reported accuracy by 5-8 percentage points relative to unfiltered operational deployment, in which marginal-visibility scenarios predominate among controversial offside calls.

Second, single-camera analysis imposes fundamental spatial resolution constraints. Monocular depth estimation from perspective projection introduces systematic spatial uncertainty. Professional SAOT systems achieve <3 cm positional accuracy through multi-camera triangulation, while single-camera approaches exhibit estimated depth errors of 10-25 cm at typical broadcast distances (40-70 m from touchline), sufficient to cause misclassification in marginal offside situations (player separation <30 cm from offside line). Occlusion-related failures ($n=18$, 7.5% of test instances) further constrain the reliability of monocular detection. Multi-camera fusion architectures employed in professional VAR and semi-automated offside technology (SAOT) achieve superior spatial accuracy through triangulated depth estimation, a capability absent from the current implementation.

Third, extreme player tracking prioritizes computational efficiency over spatial completeness, thereby introducing systematic errors in compressed defensive formations (estimated to occur in 15-20% of professional scenarios) and in non-standard camera perspectives. The approach fails when the second-to-last defender occupies a central rather than an extreme horizontal position, thereby constraining its applicability to matches with stereotypical defensive geometry.

Fourth, the binary spatial classification omits FIFA Law 11 requirements: (i) Player involvement assessment (interfering with

play, gaining advantage, interfering with opponent) requires subjective judgment beyond positional analysis. (ii) Body part eligibility specification (arms excluded, other body parts with which the ball may be legally played included) demands limb-level segmentation absent from bounding box detection. (iii) Goalkeeper positional dynamics (second-to-last defender determination when goalkeeper advances) require role-specific tracking logic not implemented. (iv) Deflection scenarios (ball contact by defending player) necessitate causal analysis linking ball trajectory changes to player actions. These omissions preclude autonomous, rule-compliant decision-making, thereby limiting the system's utility to preliminary spatial flagging that requires human adjudication.

Fifth, generalizability to heterogeneous competitive contexts (different leagues, stadium configurations, broadcast standards, tactical systems) remains unvalidated. The system's performance under varying camera heights, zoom levels, frame rates, and match tempos requires empirical verification before deployment recommendations can extend beyond the sampled conditions.

Dataset restrictions impose ecological validity constraints: (i) Competitive context: UEFA Champions League matches exclusively (2021-2023 seasons), representing elite intensity (sprint distance 28-32 km/match, transition frequency 15-18 per match) with VAR/SAOT infrastructure present. (ii) Camera specifications: Sideline elevation 42-58 m, vertical angle 35-45°, fixed zoom, 1080p resolution, 30-50 fps. (iii) Tactical diversity: Predominantly 4-3-3 and 4-2-3-1 formations with high defensive lines (38-42 m from goal); low-block defenses (<25 m) and alternative systems (5-3-2, 3-5-2) are underrepresented. (iv) Venue sampling: Three stadiums only; generalizability to variable pitch dimensions, camera positioning, and lighting conditions unvalidated. (5) Processing latency: 185-220 ms measured offline; real-time broadcast integration latency untested. (6) Operational integration: Referee communication protocols, user interface design, and IFAB Law 11 compliance procedures are unaddressed.

The study lacks empirical validation of the impact of decision-making. Performance metrics reflect algorithmic classification accuracy against expert-annotated ground truth but do not assess whether system outputs improve referee decision quality, reduce error rates in operational contexts, or enhance officiating efficiency. Evaluation of automation-bias risks, communication-protocol effectiveness, and integration with existing referee workflows remains necessary before operational deployment claims can be substantiated. Finally, the study lacks ecological validation through real-time implementation testing, assessment of referee interactions, or integration with existing officiating workflows. Decision latency, communication protocols, and automation bias risks remain unaddressed, limiting conclusions regarding operational feasibility in professional match environments.

4.2 Practical Recommendations

Based on the research findings and performance analysis, professional football leagues should consider phased deployment strategies that begin with decision support rather than replacing human officials, allowing gradual integration that builds confidence and systematically addresses operational challenges. The system's 83% accuracy, while insufficient for autonomous professional implementation, may provide supplementary information in training contexts or lower-stakes competitive environments where human oversight maintains final decision authority, aligning with successful technology adoption patterns in other sports applications. Referee training integration requires empirical validation, absent from the current study design. Proposed application: post-match case analysis providing automated classifications for discussion without real-time influence on decisions. Critical research needs include assessing automation bias (overreliance on system outputs despite errors), measuring the impact of decision quality (pre- and post-intervention comparisons), and identifying thresholds for output confidence levels that warrant presentation to officials. Implementation without these validations risks degrading human judgment by inappropriately deferring to algorithmic classifications. Realistic implementation would require phased validation: (i) offline post-match analysis to establish reliability across diverse competitive contexts, (ii) controlled field trials in non-elite competitions with parallel human officiating to quantify decision concordance rates, (iii) referee training integration to assess usability and automation bias risks, and (iv) graduated deployment contingent on demonstrated decision quality improvements and regulatory approval. Technical infrastructure would require integrating broadcast feeds, specifying real-time processing hardware, and defining communication protocols to deliver outputs to match officials. These steps remain unvalidated in the current study, precluding immediate recommendations for professional implementation. A prospective single-camera deployment architecture would require: (i) a consumer-grade GPU (minimum 8 GB VRAM) connected to a standard broadcast capture card accepting 1080p input at 30-50 fps via HDMI or SDI; (ii) camera positioning at sideline elevation 42-58 m with a vertical angle of 35-45° to minimize perspective distortion at the offside threshold; (iii) end-to-end processing latency below 200 ms, achieved through local GPU inference (19-21 ms per frame) rather than cloud-based routing; and (iv) output delivery to assistant referees via existing wireless earpiece infrastructure, with a standardized protocol distinguishing high-confidence automated flags (confidence score >0.85) from advisory alerts requiring human adjudication. These specifications constitute prospective design requirements and have not been validated under live-match conditions.

Empirically-grounded development priorities: (i) Occlusion mitigation. Error analysis indicates 72% of false negatives ($n=13/18$) occurred when attackers were obscured by defenders at identical

depth (± 2 m), suggesting second sideline-oriented camera (perpendicular baseline, 90° offset from primary view) could resolve depth ambiguity through stereoscopic triangulation; (ii) multi-camera prototype evaluation using synchronized broadcast feeds to quantify depth estimation improvements, requiring triangulation algorithm development (estimated 6-month engineering timeline); (iii) Player involvement classification. FIFA Law 11 requires assessment beyond spatial positioning (interfering with play, gaining advantage, and interfering with opponents). Implementation pathway: positional data integration with expert-coded involvement labels (binary classification: involved/not involved) using logistic regression or shallow neural network, trained on minimum 500 annotated offside scenarios with certified referee ground truth; (iv) player involvement classification module integrating spatial analysis with gaze tracking and movement pattern recognition, representing substantial algorithmic complexity increase; (v) prospective validation through controlled field trials in semi-professional competitions to assess operational reliability and referee acceptance, contingent on regulatory approval and liability frameworks. Implementation feasibility decreases with system complexity; single-camera refinement offers the most pragmatic near-term development pathway. Compliance with FIFA's semi-automated offside technology standards, which specify <3 cm spatial accuracy, < 0.5 -second decision latency, and a comprehensive player-involvement assessment would require an architectural redesign that incorporates limb-tracking algorithms, multi-camera sensor fusion, and a real-time communication infrastructure. These specifications exceed the capabilities of a single camera, positioning the current system as exploratory research rather than a pathway to a FIFA-compliant implementation.

5. CONCLUSION

This study demonstrates proof of concept for automated offside detection using extreme player tracking on single-camera broadcast footage, achieving computational efficiency (50 fps) suitable for resource-constrained competitive environments while maintaining moderate classification accuracy (83%). The system achieved moderate classification performance (83% accuracy) with computational efficiency suitable for single-camera broadcast integration. Performance levels indicate suitability for non-critical support applications, referee training, and resource-constrained environments, rather than for replacing established VAR infrastructure. While limitations remain, the evaluation confirms the feasibility of supplementary decision support in non-autonomous contexts in which final authority remains with human officials. By adapting modern computer vision techniques to the complexities of officiating, this work contributes to the growing field of sports technology. YOLOv8's anchor-free detection architecture and

optimized convolutional blocks enable real-time processing on standard broadcast hardware without specialized multi-camera infrastructure, thereby differentiating this approach from SAOT systems that require proprietary sensor arrays and from ensemble classifier methods with prohibitive inference latency. This computational accessibility enables the feasibility assessment of offside detection in lower-tier competitions, referee training programs, and broadcast analysis applications, where deployment of VAR infrastructure is financially or logistically constrained. The findings support prospective, phased implementation contingent on live-match validation, with final decision authority maintained by human officials throughout any deployment context.

Conflict of Interest

The authors declare that they have no commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

Data curation, A.A.F and S.B.B; Formal analysis, A.A.F and S.B.B; Funding acquisition, R.I.M.; Investigation, W.D., A.A.F S.B., and S.B.B; Methodology, A.E.O, H.I.C., W.D., A.A.F., I.D., and S.B.B; Project administration, A.E.O, A.A.F., H.I.C., and S.B.B; Software, A.A.F and S.B.B; Validation, A.E.O and W.D.; Writing – original draft, A.A.F., H.I.C., S.B., R.I.M., I.D., and S.B.B; Writing – review & editing, A.A.F., S.B.B., H.I.C., S.B., R.I.M., I.D., W.D and A.E.O. All authors have read and agreed to the published version of the manuscript.

Funding

The authors declare that no financial support was received for the research, authorship, and/or publication of this article.

Acknowledgments

The authors thank all the experts of the Qatar Football Federation for their participation in this work.

Data Availability Statement

The raw data supporting the conclusions of this article will be made available by the corresponding authors on request.

REFERENCES

1. Spitz J, Wagemans J, Memmert D, Williams AM, Helsen WF: **Video assistant referees (VAR): The impact of technology on decision making in association football referees.** *Journal of Sports Sciences* 2021, **39**(2):147-153.

2. Oudejans RR, Verheijen R, Bakker FC, Gerrits JC, Steinbrückner M, Beek PJ: **Errors in judging 'offside' in football.** *Nature* 2000, **404**(6773):33.
3. Unkelbach C, Memmert D: **Crowd noise as a cue in referee decisions contributes to the home advantage.** *J Sport Exerc Psychol* 2010, **32**(4):483-498.
4. Nevill AM, Balmer NJ, Mark Williams A: **The influence of crowd noise and experience upon refereeing decisions in football.** *Psychology of Sport and Exercise* 2002, **3**(4):261-272.
5. Carlos L-P, Ezequiel R, Anton K: **How does Video Assistant Referee (VAR) modify the game in elite soccer?** *International Journal of Performance Analysis in Sport* 2019, **19**(4):646-653.
6. Nagle T, Sammon D, Pope A: **Exploring the socio-technical dynamics of VAR implementation and use.** *Journal of Decision Systems* 2024, **33**(sup1):47-62.
7. Rogerson M, Knight D, Scherer R, Jones B, McManus C, Waterworth S, Murray K, Hope E: **Meta-analysis of the effects of VAR on goals scored and home advantage in football.** *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, **0**(0):17543371241242914.
8. Alhassan I, Sammon D, Daly M, Wibisono A, Kasraian L, Nagle T, Heavin C, Dennehy D, Zamani E, Qaffas A: **The use of open, axial and selective coding techniques: A literature analysis of IS research.** 2023.
9. Naik BT, Hashmi MF, Bokde ND: **A comprehensive review of computer vision in sports: Open issues, future trends and research directions.** *Applied Sciences* 2022, **12**(9):4429.
10. Terven J, Córdova-Esparza D-M, Romero-González J-A: **A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS.** *Machine Learning and Knowledge Extraction* 2023, **5**(4):1680-1716.
11. Panse N, Mahabaleshwarkar A: **A dataset & methodology for computer vision based offside detection in soccer.** In: *Proceedings of the 3rd International Workshop on Multimedia Content Analysis in Sports: 2020*; 2020: 19-26.
12. Jocher G, Chaurasia A, Qiu J: **YOLO by Ultralytics.** In.; 2023.
13. Maaroo MKA, Bouhleb MS: **Real-Time Object Detection Using YOLO-8 Model: A Drone-Based Approach.** *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2025, **16**(1):190-204.
14. Harwood H: **Investigating Video Assistant Referees (VARs): Does contextual information influence VAR decision-accuracy and decision-severity in football.** University of Essex; 2024.
15. Postupaiev S, Damaševičius R, Maskeliūnas R: **Real-Time Camera Operator Segmentation with YOLOv8 in Football Video Broadcasts.** *AI* 2024, **5**(2):842-872.
16. Hermens F: **Automatic object detection for behavioural research using YOLOv8.** *Behavior Research Methods* 2024, **56**(7):7307-7330.
17. Li C, Zhu Y, Zheng M: **A multi-objective dynamic detection model in autonomous driving based on an improved YOLOv8.** *Alexandria Engineering Journal* 2025, **122**:453-464.

18. Fedorov VA: **Railway Infrastructure Detection Based on YOLOv8 with NPU Acceleration.** *Doklady Mathematics* 2024, **110**(1):S42-S48.
19. Ayachi R, Said Y, Afif M, Alshammari A, Hleili M, Ben Abdelali A: **Assessing YOLO models for real-time object detection in urban environments for advanced driver-assistance systems (ADAS).** *Alexandria Engineering Journal* 2025, **123**:530-549.
20. Li A, Li D, Wang A: **A Two-Stage YOLOv5s-U-Net Framework for Defect Localization and Segmentation in Overhead Transmission Lines.** *Sensors* 2025, **25**(9):2903.
21. Rabee A, Anwar Z, AbdelMoety A, Abdelsallam A, Ali M: **Comparative analysis of automated foul detection in football using deep learning architectures.** *Sci Rep* 2025, **15**(1):14236.
22. Dowsett B: **AI is helping referee games in major sports leagues, but limitations remain.** *Scientific American* <https://www.scientificamerican.com/article/ai-is-helping-referee-games-in-major-sports-leagues-but-limitations-remain> 2024.
23. Fakhrurroja H, Fashihullisan AA, Bangkit H, Pramesti D, Ismail N, Mahardiono NA: **A Vision-Based System: Detecting Traffic Law Violation Case Study of Red-Light Running Using Pre-Trained YOLOv8 Model and OpenCV.** In: *2024 IEEE International Conference on Smart Mechatronics (ICSMech): 2024*: IEEE; 2024: 1-6.
24. Kolbinger O, Lames M: **Scientific approaches to technological officiating aids in game sports.** *Current Issues in Sport Science* 2017, **2**:1-10.
25. Dhahbi W, Chamari K: **The Algorithmic Athlete: A Call to Standardize Assessment of Sensor Technologies and Artificial Intelligence.** *Int J Sports Physiol Perform* 2026:1-2.
26. Amosa TI, Sebastian P, Izhar LI, Ibrahim O, Ayinla LS, Bahashwan AA, Bala A, Samaila YA: **Multi-camera multi-object tracking: A review of current trends and future advances.** *Neurocomputing* 2023, **552**:126558.
27. Elstak I, Paul S, and McLean S: **Artificial intelligence applications in the football codes: A systematic review.** *Journal of Sports Sciences* 2024, **42**(13):1184-1199.

TABLES

Table 1. Dataset characteristics and distribution

Characteristic	Category	Count	Percentage
Total Samples	Images/Video Segments	400	100%
	Clear Offside	180	45%
Match Scenarios	No Offside	160	40%
	Marginal Cases	60	15%
Playing Conditions	Daylight	240	60%
	Evening/Artificial Light	160	40%
Camera Angles	Side View	280	70%
	Elevated View	120	30%
Team Combinations	High Contrast Jerseys	320	80%
	Similar Color Jerseys	80	20%
Data Partition	Training Set	160	40%
	Validation Set	240	60%

Note: Comprehensive breakdown of dataset characteristics showing distribution across key variables affecting automated offside detection performance. Model validation was performed using 5-fold cross-validation on the training partition rather than a held-out validation set, thereby maximizing test-set allocation for robust performance assessment.

Table 2. Confusion matrix analysis for binary offside classification performance

	Predicted Not Offside	Predicted Offside
Actual Not Offside	70 (TN)	22 (FP)
Actual Offside	18 (FN)	130 (TP)

TN = True Negative, TP = True Positive, FP = False Positive, FN = False Negative. Analysis **Note:** based on 240 test instances across diverse match scenarios.

Caption: Comprehensive confusion matrix summarizing automated offside detection system performance across 240 test instances. True Negative (TN) = 70 represents correct identification of valid player positions; True Positive (TP) = 130 indicates accurate detection of genuine offside violations; False Positive (FP) = 22 denotes incorrect flagging of valid positions as offside; False Negative (FN) = 18 represents missed detection of actual violations. Matrix elements demonstrate the system's strength in positive detection (TP is the largest category) while maintaining an acceptable false-positive rate. Overall accuracy = 83.0%, precision = 85.0%, recall = 87.0%, F1-score = 86.0%. Statistical analysis confirms performance significantly exceeds random classification baseline ($p < 0.001$).

Class-specific metrics: Offside class: precision 85.5%, recall 87.8%, F1-score 86.6%. Not-Offside class: precision 79.5%, recall 76.1%, F1-score 77.7%. Macro-average F1: 82.2%; weighted F1: 83.3%.

Table 3. Per-fold cross-validation performance metrics on training partition (n=160)

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Fold 1	82.1%	84.0%	86.5%	85.2%
Fold 2	84.2%	86.1%	88.0%	87.0%
Fold 3	83.0%	85.0%	87.4%	86.2%
Fold 4	82.7%	84.6%	86.9%	85.7%
Fold 5	84.5%	86.3%	88.2%	87.2%
Mean $\pm$ SD	83.3 $\pm$ 1.0%	85.2 $\pm$ 1.0%	87.4 $\pm$ 0.7%	86.3 $\pm$ 0.8%

Note: Cross-validation was applied exclusively to the training partition (n=160) for hyperparameter optimization. Final performance estimates are reported on the independent held-out test set (n=240; Table 2). No Bayesian inference was employed; confidence intervals were estimated via BCa bootstrap resampling (1,000 iterations).




Figure 1: Algorithmic workflow for automated offside detection using YOLOv8 architecture.




Figure 2: Model training and validation performance curves across training epochs. **Caption:** Training and validation accuracy progression throughout the model development process, demonstrating convergence characteristics and generalization capability. The parallel trajectories of the training and validation curves indicate successful learning without overfitting, with both metrics stabilizing at optimal performance levels. Training accuracy approaches 87%, while validation accuracy converges to 83%, with a narrow performance

gap suggesting appropriate model complexity and robust generalization to unseen data. The smooth convergence pattern confirms optimal hyperparameter selection and the effectiveness of the training protocol for adapting the YOLOv8 architecture to football offside detection tasks.

ARTICLE IN PRESS



Figure 3: Real-time object detection performance on professional match footage.

Caption: Demonstration of YOLOv8 system performance on actual match imagery showing detection accuracy for primary object classes: Team 1 players, Team 2 players, and ball identification. Bounding boxes indicate successful object detection with confidence scores displayed for each classification. The image illustrates typical detection performance under standard viewing conditions and highlights the challenges posed by player occlusion. Visual analysis reveals consistently high detection confidence across all object categories, with particular attention to situations where ball positioning between players increases detection complexity, which may affect subsequent offside analysis accuracy.




Figure 4: Successful automated offside detection in a professional match context.

Caption: Representative example of accurate offside violation detection showing attacking player positioned beyond the defensive line. This practical demonstration confirms that laboratory performance validation

effectively translates to real-world conditions, with processing completed within an acceptable timeframe for near-real-time implementation.

ARTICLE IN PRESS



Figure 5: Correct negative classification in a complex player positioning scenario

Caption: Example of accurate negative offside determination where the system correctly identified that no violation occurred despite potentially ambiguous player positioning. Multiple players in proximity to the offside threshold create visual complexity that might challenge human officials, yet the automated system-maintained

classification accuracy. This negative case validation demonstrates the system's discriminative capability and reliability across the full decision spectrum, supporting its potential as a decision support tool for professional officiating applications.

ARTICLE IN PRESS

YOLOv8 Computer Vision for Automated Offside Detection in Professional Football Validated Through Supervised Learning

Abdul-Rahman Abdel-Fattah ¹, Samir Brahim Belhaouari ¹, Halil İbrahim Ceylan², Serdar Bayrakdaroglu³, Raul Ioan Muntean⁴, Wissem Dhahbi ^{5,6}, Ismail Dergaa ^{7,8}, Abdelfatteh El Omri ^{9,10}

¹ Division of Information and Computing Technology, College of Science and Engineering, Hamad Ben Khalifa University, Doha, Qatar. abab46015@hbku.edu.qa , sbelhaouari@hbku.edu.qa

²Faculty of Sports Sciences, Atatürk University, Erzurum 25240, Türkiye. halil.ibrahimceylan60@gmail.com

³Faculty of Sports Sciences, Gümüşhane University, Gümüşhane 29100, Türkiye. bayrakdaroglu85@gmail.com

⁴Department of Physical Education and Sport, Faculty of Law and Social Sciences, University “1 Decembrie 1918” of Alba Iulia, Alba Iulia, Romania. muntean.raul@uab.ro

⁵ Research Unit “Sport Sciences, Health and Movement”, High Institute of Sports and Physical Education of Kef, University of Jendouba, Kef 8100, Tunisia.

⁶ Training Department, Police College, Qatar Police Academy, Doha 59J5+C7R, Qatar. wissem.dhahbi@gmail.com

⁷ High Institute of Sport and Physical Education of Ksar Said, University of Manouba, Manouba 2010, Tunisia

⁸ Physical Activity Research Unit, Sport and Health (UR18JS01), National Observatory of Sports, Tunis 1003, Tunisia. Phd.dergaa@gmail.com

⁹ Surgical Research Section, Department of Surgery, Hamad Medical Corporation, Doha, Qatar

¹⁰ QU-Health, Qatar University, Doha, Qatar. AEIomri@hamad.qa

***Correspondence:** halil.ibrahimceylan60@gmail.com or muntean.raul@uab.ro

ABSTRACT

Aim: Offside calls are integral to maintaining fairness and competitive integrity in football, yet current officiating processes are constrained by limitations in accuracy despite technological advancements. This investigation developed and validated an automated offside detection system using YOLOv8 object detection architecture to address documented limitations of Video Assistant Referee (VAR) systems in offside decision accuracy. **Methods:** Supervised learning validation employed 400 annotated instances from professional match footage. The YOLOv8 medium variant was configured for four-class detection, with a focus on extreme player positions, reducing computational complexity by approximately 70% relative to comprehensive player tracking while accepting accuracy trade-offs in marginal scenarios. The system was optimized for computational efficiency through highly accurate tracking of player positions. Data were partitioned into 40% (n=160) for training and 60% (n=240) for testing, with 5-fold cross-validation. **Results:** The automated system achieved robust performance metrics across 240 test instances: overall accuracy of 83.0%, precision of 85.0%, recall of 87.0%, and F1-score of 86.0%. Confusion matrix analysis revealed true positives (n=130), true negatives (n=70), false positives (n=22), and false negatives (n=18). Individual case analysis demonstrated 80% agreement with ground truth classifications. System performance significantly exceeded the random classification baseline ($p < 0.001$). The strategic extreme player tracking approach achieved substantial computational efficiency while maintaining decision accuracy across diverse match scenarios. **Conclusion:** A YOLOv8-based offside detection system demonstrates proof-of-concept feasibility for decision-support applications in football officiating. Performance (83% accuracy, 95% CI: 79-87%) on 240 test instances remains below professional benchmarks (>93%), with single-camera constraints and limited dataset size restricting generalizability to proof-of-concept validation rather than operational deployment.

KEYWORDS: Artificial intelligence; machine learning; neural networks; offside detection; real-time object detection; sports officiating technology; video assistant referee; YOLOv8

1. INTRODUCTION

Competitive football officiating requires 137 observable decisions per match at the international level, with offside calls accounting for 15-20% of these decisions and directly affecting scoring opportunities [1]. Among these critical decisions, offside infractions are among the most challenging and controversial aspects of football officiating, requiring split-second judgments in fast-paced, dynamic environments where human perceptual limitations frequently compromise decision accuracy.

The complexity of offside detection stems from multiple interrelated factors that challenge experienced officials. Assistant referees in elite European competitions (UEFA Champions League, English Premier League, 2015-2019 seasons) achieve 94% offside accuracy (95% CI: 92-96%, $n=1,247$ decisions coded via video analysis with inter-rater reliability $\kappa=0.91$) [2]. This benchmark remains consistent with recent evaluations reporting 93-95% accuracy pre-VAR implementation, with error concentration in marginal violations (<30 cm player separation from offside threshold). Impact quantification across 760 top-tier league matches (2017-2020) indicates that erroneous offside calls affect 8-12 goals per season, with documented instances that alter final league standings (2-3 position shifts) and European competition qualification in five seasons across major leagues. Most of these errors stem from the limits of human perception. Assistant referees are required to track several players at once, follow the ball's path, and judge the exact moment of contact—all while staying in the right position. Research shows that crowded situations, in which many players cluster, make these judgments even more difficult, and that players who are not directly involved can still influence decisions through their positioning or gestures [3]. External factors also play a role: crowd noise, for example, has been shown to reduce foul-detection rates against home teams by 15.5% relative to silent conditions, highlighting how auditory distractions can impair focus and decision-making [4]. In addition, the pressure of high-stakes matches adds further cognitive strain, often leading to systematic biases in refereeing decisions.

Video Assistant Referee (VAR) technology, integrated into Laws of the Game in 2018, encompasses distinct architectural categories [5]. Conventional VAR employs multi-angle video review with human adjudication across all decision types (goals, penalties, red cards, mistaken identity), achieving 93-96% overall accuracy in professional competitions [6]. Semi-automated offside technology (SAOT), deployed in the FIFA World Cup 2022 and the UEFA Champions League 2023-present, uses 12-29 dedicated tracking cameras (minimum 50 fps) with limb-detection algorithms (29 body points per player, 3 cm spatial tolerance) to achieve $>98\%$ offside-specific accuracy. Fully automated systems eliminating human oversight remain absent from professional competition due to rule complexity (player involvement assessment) and liability constraints. A meta-

analysis of 20 competitions (9,871 matches across all VAR decision categories) identified persistent error mechanisms: communication protocol failures, compressed decision timeframes, and camera angle limitations [7]. Offside decisions constituted 18-22% of VAR interventions across competitions. A well-known example occurred in September 2023 during Tottenham Hotspur's match against Liverpool, when hurried communication between officials led to a valid goal being disallowed [8]. Furthermore, VAR's reliance on camera angles can create blind spots, making it less effective for close offside calls, especially when players obstruct the view or the exact moment of ball contact cannot be determined with certainty. Operational challenges extend beyond classification accuracy to decision latency (conventional VAR: 60-180 seconds per review), transparency (opaque adjudication processes), and public acceptance (52% fan approval in European surveys, 2021-2023). The proposed YOLOv8-based approach addresses a distinct niche: single-camera broadcast integration without specialized hardware, trading SAOT-level precision for accessibility in resource-constrained competitive contexts.

Advances in computer vision have created new possibilities for automating sports officiating through real-time object detection and tracking. In particular, deep learning models such as the You Only Look Once (YOLO) family have shown strong performance in identifying and classifying multiple objects in complex visual settings [9]. The evolution from YOLOv1 to YOLOv8 has led to steady improvements in detection accuracy, processing speed, and efficiency, making real-time analysis in sports increasingly practical [10]. Research on automated offside detection has explored contour-based silhouette mapping, background subtraction with Kalman filtering, and ensemble-classifier approaches [11]. However, these methods exhibit critical limitations: contour-based techniques achieve only 76-81% accuracy on standardized datasets and fail under player-occlusion conditions; ensemble approaches require computationally prohibitive multi-model inference, which is unsuitable for real-time processing; and most validations use synthetic datasets or controlled laboratory conditions rather than broadcast match footage. Dataset constraints remain severe, with typical studies employing fewer than 200 annotated instances, which is insufficient for robust generalization in deep learning. Architecturally, contour-based methods rely on handcrafted feature extraction and threshold-based spatial logic, rendering them brittle under variable illumination and player occlusion; ensemble classifiers achieve marginal accuracy gains at the cost of multi-model inference latency exceeding 200 ms, which precludes real-time deployment. By contrast, YOLOv8's anchor-free convolutional architecture performs end-to-end multi-class detection (players, ball, officials) within a single inference pass of 19-21 ms, directly addressing the occlusion sensitivity and latency constraints of predecessor methods. Unlike SAOT, which requires

proprietary arrays of 12-29 calibrated cameras per stadium, the proposed system operates on standard 1080p broadcast feeds, eliminating hardware barriers in resource-constrained competitive contexts at the accepted cost of reduced spatial precision. These methodological gaps motivated the adoption of YOLOv8's single-shot detection architecture, which balances accuracy and computational efficiency while enabling training on limited real-world data.

Recent YOLOv8 deployments in football contexts include camera operator segmentation for broadcast analysis [16] and multi-person tracking for tactical evaluation, demonstrating the feasibility of real-time detection (45-60 fps on consumer GPUs). However, these applications omit the implementation of the offside rule and lack validation against referee ground truth. Methodological approaches span player segmentation (instance masks for spatial localization), multi-object tracking (temporal trajectory estimation), and field geometry projection (homography-based coordinate transformation). Limitations include depth ambiguity in monocular vision, occlusion sensitivity in crowded scenarios, and the absence of an assessment of player involvement required under FIFA Law 11.

YOLOv8 is the latest generation of single-shot object detection models, introducing key upgrades such as an anchor-free detection system, improved convolutional layers, and refined training strategies that yield improved accuracy and speed trade-offs compared with earlier versions [12]. It can handle 1080p video at around 50 frames per second while preserving high detection accuracy, making it well-suited for real-time sports applications [13]. Studies have further confirmed its reliability in challenging scenarios; for example, Hermens et al. [14] reported that YOLOv8 achieved near-perfect detection accuracy even when trained on relatively small datasets of 100-350 images.

Although VAR and computer vision technologies have advanced considerably, there is still no fully automated offside-detection system that delivers the accuracy and reliability required at the professional level. Studies show that human officials perceive ball contact with an average delay of about 132 milliseconds, which can introduce systematic bias in close offside calls [14]. The research gap comprises three dimensions: (i) absence of validated single-camera offside detection at broadcast frame rates (30-60 fps) without specialized hardware, addressing accessibility barriers in sub-elite competitions; (ii) lack of performance quantification on operational broadcast footage under authentic match conditions rather than controlled laboratory datasets; (iii) undefined accuracy-cost trade-off boundaries for contexts where SAOT infrastructure (\$2-4 million per stadium) is prohibitive. The primary objective is to validate computational efficiency for resource-constrained implementation while not exceeding the precision of existing SAOT capabilities. Existing solutions require either multi-camera infrastructure with proprietary hardware (semi-automated offside technology systems) or sacrifice

accuracy for speed (single-model approaches). This gap constrains access to officiating support technology across competitive levels and impedes the systematic evaluation of automated detection in resource-constrained environments. This study addresses this issue by adapting YOLOv8's object detection capabilities for football analysis. The approach emphasizes detecting extreme player positions and tracking ball movement, reducing computational demands while preserving decision accuracy. The main contributions are threefold: (i) a streamlined algorithm that targets key player positions instead of tracking the entire field, (ii) integration of temporal analysis to capture ball dynamics and passing direction, and (iii) a thorough evaluation using real match footage to demonstrate its potential for professional use. The system is positioned as a decision-support prototype for evaluation rather than a replacement for existing officiating infrastructure.

2. METHODS

2.1 Study Design

This study evaluates a deep learning-based offside detection system using supervised learning and held-out test validation on the YOLOv8 architecture. Annotated professional match footage provided training and testing data, with validation protocols following established computer vision benchmarking standards adapted for sports officiating contexts.

2.2 Dataset Collection and Characteristics

The dataset consisted of 400 high-resolution images and video clips from professional football matches, captured under various playing conditions. Broadcast footage was used to ensure coverage across different stadiums, lighting environments, and camera angles typically found in competitive settings. Only footage with clear visibility of players, the ball, and field markings was included, with a minimum resolution of 1080p to provide sufficient detail for accurate detection and classification. Eligible samples contained well-defined offside situations, including unobstructed views of key players and ball position during active play sequences such as passes and positional shifts [15]. Footage was excluded if it exhibited excessive motion blur, significant occlusion, poor lighting, or originated from non-competitive contexts (Table 1). The final sample size of 400 annotated instances was informed by standard computer vision research practices and pilot testing. Sample size determination followed pragmatic constraints of annotation resource availability rather than formal power calculations, which are not standard practice for deep learning classifier development, where performance bounds depend on model capacity, data complexity, and training convergence rather than predetermined effect size detection [16]. This sample size, while enabling initial proof-of-concept validation, imposes a constraint on claims of generalization in deep learning. Dataset limitations restrict

performance estimates to approximate bounds rather than precise population-level accuracy projections.

****Please Insert Table 1 About Here****

2.3 Data Preprocessing and Annotation Protocol

Data preprocessing followed standard computer vision practices to maintain consistency across the dataset. All images were resized to 640×640 pixels to match YOLOv8 requirements while retaining essential visual detail. Color normalization was applied by converting to RGB format and scaling pixel values to [0,1], ensuring stable training across varied lighting conditions and camera setups [17]. An annotation framework defined four object categories: players from each team, the ball, and match officials, using bounding boxes with precise coordinates. To ensure reliability, two independent annotators with sports science backgrounds labeled all instances, achieving Cohen's kappa $\kappa = 0.87$ (95% CI: 0.83-0.91). Discrepancies ($n = 52$, 13% of the dataset) were resolved through structured consensus discussions among annotators, and the remaining disagreements ($n = 8$, 2% of the dataset) were adjudicated by a certified UEFA B-licensed referee with 12 years of professional match analysis experience, whose decisions established the final ground-truth classifications. To increase dataset diversity and improve generalization, several augmentation techniques were applied, including random horizontal flips ($p = 0.5$), rotations of up to $\pm 15^\circ$, brightness shifts of $\pm 20\%$, and contrast changes of $\pm 15\%$. These transformations introduced natural variation while preserving spatial relationships needed for offside analysis. Mosaic augmentation was used during training but removed in the final 10 epochs to reduce the risk of overfitting.

2.4 YOLOv8 Model Implementation and Training

The system was built using the medium YOLOv8 (YOLOv8m) model, chosen for its balance between detection accuracy and computational efficiency in real-time scenarios [18]. The architecture included anchor-free detection, improved convolutional blocks, and optimized feature extraction layers, with custom output heads designed for four-class detection and weighted to account for class imbalance [19]. Training parameters were optimized through a grid search, with the final configuration set to a learning rate (η) of 0.01, batch size of 16, and 300 epochs, using early stopping when validation loss plateaued. The Adam optimizer with weight decay ($\lambda = 0.0005$) was used to reduce overfitting. The data were split using stratified sampling: 40% ($n=160$) for training and 60% ($n=240$) for testing, with 5-fold cross-validation applied to the training partition for model optimization. Transfer learning was employed by initializing the model with YOLOv8 weights pre-trained on COCO, followed by fine-tuning for football-specific detection. After initial convergence, all network layers were

unfrozen and retrained with a lower learning rate ($\eta = 0.001$) to refine performance [20].

2.5 Team Classification and Player Identification

Team classification methodology employed advanced color segmentation techniques combined with machine learning algorithms to distinguish between opposing teams. The approach utilized HSV color space analysis to extract dominant jersey colors, followed by clustering algorithms (K-means, $k=2$) to separate teams based on chromatic characteristics, with validation procedures including manual verification across diverse lighting conditions and jersey color combinations. Player identification protocols focused on extracting the spatial coordinates of detected players within the field coordinate system, using geometric transformation techniques to convert pixel coordinates to normalized field positions via perspective projection. Kalman filtering was implemented to smooth player position trajectories and handle temporary occlusions during match play.

2.6 Automated Offside Detection Algorithm

The offside detection algorithm prioritized computational efficiency by using extreme player tracking (the leftmost and rightmost positions per team) rather than comprehensive field coverage. This approach assumes the second-to-last defender occupies an extreme horizontal position, reducing computational load by approximately 70% relative to full-field tracking but introducing potential spatial errors when defensive lines compress centrally or camera perspectives create depth ambiguity. This strategic reduction in tracking complexity significantly improved processing efficiency while maintaining decision accuracy for offside determination. Ball detection and trajectory analysis incorporated temporal information to determine pass direction and timing through motion vector calculations, establishing ball velocity and direction, with pass detection thresholds calibrated based on empirical analysis of ball movement patterns in professional match footage. The core offside detection algorithm implemented official FIFA offside rules through computational logic, evaluating player positions relative to the second-last opponent at the moment of ball contact and incorporating temporal analysis to ensure accurate timing assessment, with decision logic including exception handling for situations where offside rules do not apply and confidence scoring for decision reliability assessment.

Recognized limitations of extreme player tracking include: (i) failure to detect offside violations when the second-to-last defender occupies a central rather than extreme position, (ii) systematic spatial errors under diagonal camera angles that distort horizontal positioning, (iii) reduced accuracy in compressed defensive formations where multiple defenders cluster near the offside threshold, and (iv) potential misclassification when goalkeeper positioning creates non-standard defensive geometry. These

constraints represent deliberate trade-offs between the feasibility of real-time processing and the spatial completeness, with acknowledged accuracy costs in geometrically complex matching scenarios. No direct performance comparison with full-field player tracking was conducted due to resource constraints, preventing empirical quantification of the accuracy trade-offs between extreme position heuristics and comprehensive spatial coverage.

****Please Insert Figure 1 About Here****

The algorithm employs a streamlined approach that focuses on extreme player positions (leftmost and rightmost per team) rather than comprehensive player tracking, thereby significantly reducing computational complexity while maintaining decision accuracy. Parallel processing of player identification and ball trajectory analysis enables real-time performance suitable for live match implementation (Figure 1).

2.7 Performance Evaluation and Validation

Model performance was evaluated using standard computer vision metrics, such as precision, recall, F1-score, and mean Average Precision (mAP), as well as football-specific measures, including offside decision accuracy, false positive rate, and temporal consistency across video sequences. To ensure robustness, a five-fold cross-validation ($k=5$) approach was used, and additional test datasets from different leagues, stadiums, and playing conditions were included to assess generalizability [21]. Ground-truth labels were created by certified referees and match analysts, and inter-observer agreement was measured using Fleiss' kappa to provide a benchmark for human decision accuracy.

2.8 Statistical Analysis

Statistical analysis was conducted in Python 3.9 using SciPy 1.11.0, NumPy 1.24.0, and scikit-learn 1.3.0. For performance comparisons, paired t-tests were used for normally distributed data, whereas Wilcoxon signed-rank tests were used for non-normally distributed data. Effect sizes were reported using Cohen's d where appropriate. Ninety-five percent confidence intervals were estimated using bias-corrected accelerated (BCa) bootstrap resampling (1,000 iterations); no Bayesian inference or Markov Chain Monte Carlo procedures were employed. BCa resampling was selected for its greater robustness to distributional skewness compared with standard percentile bootstrap methods. Classifier discrimination was assessed through precision-recall trade-off analysis. Five-fold cross-validation on the training partition informed hyperparameter selection (learning rate, batch size, and augmentation parameters), with final performance evaluated on the held-out test set to prevent overfitting and ensure unbiased accuracy estimation. Statistical significance was set at $\alpha = 0.05$, and

Bonferroni corrections were applied to adjust for multiple comparisons. The primary performance metric (overall accuracy: 83.0%) was computed exclusively on the held-out test set (n=240), which was withheld from all training and hyperparameter optimization procedures. The 5-fold cross-validation applied to the training partition (n=160) served solely for hyperparameter selection; fold-level performance metrics are reported in Table 3.

3. RESULTS

3.1 Overall System Performance and Statistical Validation

The trained model achieved 83.0% accuracy (95% CI: 79.1-86.9%) on 240 test instances, with a precision 85.0% (95% CI: 81.3-88.7%), a recall 87.0% (95% CI: 83.5-90.5%), and an F1-score 86.0% (95% CI: 82.7-89.3%). Five-fold cross-validation on the training partition (n=160) yielded a mean cross-validated accuracy of $83.3 \pm 1.0\%$ across folds (Table 3), confirming convergence stability prior to held-out test evaluation. Precision reached 85.0%, meaning that about 85 out of every 100 flagged offsides were correct, while recall was 87.0%, reflecting the system's ability to reliably detect true violations. The resulting F1-score of 86.0% highlights a balanced trade-off between sensitivity and specificity, confirming the system's potential for use in automated officiating.

3.2 Confusion Matrix Analysis and Classification Performance

Confusion matrix analysis provided insight into the system's classification performance (Table 2). The largest group was true positives (n = 130), showing the model's strong ability to detect genuine offside violations. True negatives (n = 70) confirmed reliable recognition of legal play, preventing unnecessary stoppages. False positives (n = 22) occurred primarily in marginal cases in which players were positioned very close to the offside line, a scenario that often challenges even experienced referees. The relatively low number of such errors reflects a conservative decision-making bias that favors accuracy in uncertain situations. False negatives (n = 18) were largely attributable to temporary occlusion or rapid movements that made precise timing difficult to determine, underscoring the inherent limitations of single-camera analysis. Class-specific analysis derived from the confusion matrix reveals asymmetric performance across decision categories. For the offside class: precision 85.5%, recall 87.8%, F1-score 86.6%. For the not-offside class: precision 79.5%, recall 76.1%, F1-score 77.7%. The macro-average F1-score was 82.2%, and the weighted F1-score was 83.3%, consistent with the reported overall accuracy. The 8.9-percentage-point F1 differential between classes reflects the system's stronger capacity to identify genuine violations than to confirm legal play, attributable to the class imbalance in the test set (offside: n=148; not-offside: n=92).

****Please Insert Table 2 About Here****

****Please Insert Table 3 About Here****

3.3 Visual Performance Analysis and Model Training Validation

Training and validation results showed clear convergence without signs of overfitting, with both accuracy curves progressing in parallel (Figure 2). This indicates that the model learned generalizable features rather than memorizing training data. On real match footage, the system consistently detected all object categories with high confidence under normal conditions. The main challenge arose in occlusion scenarios, such as when the ball was partially obscured by players or when player positioning blocked key elements, which reduced detection confidence and occasionally led to errors (Figure 3).

****Please Insert Figure 2 About Here****

****Please Insert Figure 3 About Here****

3.4 Practical Application Demonstration and Real-World Validation

The system's practical value was shown by testing it on real offside situations from professional matches. It successfully detected attacking players beyond the defensive line at the moment of ball contact, accurately identified teams, tracked the ball, and analyzed the spatial positions required for offside decisions (Figure 4). It also correctly identified situations without violations, even when player positioning could have created visual uncertainty for referees, demonstrating its reliability across both positive and negative decisions (Figure 5).

****Please Insert Figure 4 About Here****

****Please Insert Figure 5 About Here****

3.5 Error Analysis and Technical Limitations

Error analysis showed that false positives ($n = 22$) primarily occurred in marginal cases in which players were positioned very close to the offside line, making precise spatial measurement especially difficult. False negatives ($n = 18$) were primarily due to visual occlusion, in which defenders temporarily blocked the attackers' view, with additional errors arising from difficult camera angles or lighting conditions. Importantly, errors were evenly distributed across teams and field positions, indicating that the limitations arose from measurement challenges rather than algorithmic bias. The extreme player-tracking approach likely contributed to errors in matches featuring compressed defensive lines (estimated 15-20% of professional match scenarios) or non-standard camera perspectives, in which the assumption that offside thresholds align with horizontally extreme players may be violated.

3.6 Processing Efficiency and Implementation Feasibility

Processing efficiency: NVIDIA RTX 3080 hardware (10 GB VRAM, consumer-grade GPU) achieved 48-52 fps on 1080p input, exceeding 30 fps broadcast integration threshold. Per-frame inference time: 19-21 ms (mean: 20 ms, SD: 1.2 ms), with player count variation (18-22 detected instances) producing <5% latency fluctuation. Operational deployment latency (frame acquisition through the broadcast infrastructure, preprocessing, inference, and output transmission to the referee communication system) remains unquantified; offline validation cannot verify real-time stability under live-match conditions with network transmission delays and concurrent processing demands. The efficiency gains enabled deployment on standard broadcasting equipment without requiring specialized high-performance computing infrastructure, reducing implementation barriers and costs for widespread adoption.

3.7 System Reliability and Performance Context

The overall system performance of 83% accuracy, 85% precision, 87% recall, and 86% F1-score demonstrates proof-of-concept feasibility for automated detection, though it remains below the accuracy of professional assistant referees (~94%) and VAR-assisted decision-making (93-98% in elite competitions). Balanced performance across precision and recall metrics indicates appropriate trade-offs between conservative and aggressive decision-making approaches, which are essential for the practical implementation of referee decision-support tools. The performance consistency demonstrated across diverse match scenarios supports the system's potential for professional implementation while acknowledging technical challenges inherent in automated visual analysis of dynamic sports scenarios.

4. DISCUSSION

This study demonstrates that YOLOv8-based automated offside detection achieves moderate accuracy suitable for decision-support prototyping in football officiating contexts. The system achieved an overall accuracy of 83.0%, precision of 85.0%, recall of 87.0%, and F1-score of 86.0%. These results represent a notable step forward in automated officiating, particularly given the efficiency gains from focusing on extreme player positions rather than tracking the entire field. Confusion matrix results further highlight this strength: true positives ($n = 130$) were the largest group, reflecting the model's ability to identify genuine offside violations, whereas the relatively low number of false positives ($n = 22$) indicates a conservative bias that favors accuracy in borderline cases.

The asymmetric class-specific performance carries distinct operational implications. False positives ($n=22$; 9.2% of test instances) result in wrongful disallowance of legal attacking

sequences, directly affecting scoring opportunities and constituting officiating errors equivalent to incorrect flag raises. False negatives ($n=18$; 7.5% of test instances) represent undetected violations that confer an illegitimate positional advantage to attacking players. Under the officiating principle of conservative decision-making, the higher FP cost relative to FN cost suggests that deployment thresholds should be calibrated toward higher precision at the acceptable cost of reduced recall. At the current performance level (not-offside F1: 77.7%), autonomous deployment in professional contexts is not operationally acceptable given the established human assistant referee benchmark of 93–94% [2]; however, the system meets a threshold appropriate for supplementary flagging in sub-elite contexts where no technological support is available.

The system's performance compares favorably with recent research on automated sports officiating and directly addresses key shortcomings of current VAR use. Studies show that VAR remains prone to human error, communication lapses, and time delays, which can undermine decision quality [22]. Harwood's motion-capture work, for example, found that referees perceive ball contact with an average delay of 132 milliseconds, thereby introducing systematic bias in marginal offside calls [14]. By using frame-by-frame analysis, the proposed system removes this subjective timing factor, thereby reducing a major source of error in VAR decisions. Compared with earlier computer vision approaches, which typically achieved 76–81% accuracy [11], the YOLOv8-based model reached 83% accuracy while also offering better computational efficiency.

The system's strong recall (87%) is largely due to YOLOv8's advanced feature extraction and its focus on tracking extreme player positions. The model's anchor-free detection and improved convolutional blocks enhance localization accuracy, even in fast-paced game scenarios with frequent occlusions [23]. The relatively low false positive rate (9.2%) reflects a conservative decision-making tendency, consistent with officiating principles that favor allowing play to continue in uncertain situations rather than risking an incorrect stoppage. This mirrors patterns observed in psychological studies of referees, who tend to avoid false positives that could unfairly disrupt scoring opportunities [4]. The balance between precision (85%) and recall (87%) indicates that the detection threshold was effectively tuned, avoiding the typical extremes of overly cautious or overly aggressive classification.

Relative to prior automated approaches, the YOLOv8-based system surpasses contour-based silhouette methods (76–81% [11]) by addressing their principal failure modes through learned convolutional feature representations rather than handcrafted segmentation, and outperforms ensemble classifiers operationally by reducing inference latency from >200 ms to 19–21 ms per frame. Compared with SAOT, the single-camera approach accepts a spatial precision trade-off (estimated depth error 10–25 cm versus SAOT's <3

cm through multi-camera triangulation) in exchange for accessibility in contexts where dedicated camera infrastructure is financially prohibitive. The system therefore occupies a distinct architectural niche: single-pass real-time detection without proprietary hardware, targeting competitive levels where VAR and SAOT deployment are economically or logistically unfeasible. Unlike multi-camera semi-automated offside technology (SAOT) systems deployed in elite competitions, which achieve >98% spatial accuracy through triangulated depth estimation and limb-tracking algorithms, the single-camera approach prioritizes accessibility and computational efficiency over precision. The performance gap reflects fundamental trade-offs between system complexity, implementation cost, and accuracy requirements across competitive levels.

The system's performance characteristics suggest potential application as a supplementary decision-support tool in non-critical officiating contexts, referee training programs, and competitive environments where VAR infrastructure is financially or logistically unavailable, thereby meeting key operational demands in sports technology adoption. Its 83% accuracy offers reliable referee assistance while preserving human officials' ultimate authority, consistent with technology-assisted officiating principles [24]. Computational efficiency through extreme player tracking enables deployment on standard broadcasting equipment, reducing implementation barriers relative to comprehensive systems and enhancing accessibility across competitive levels. Consistent performance across diverse scenarios and a balanced error distribution, indicating no systematic bias, supports fair, competitive applications crucial for governing bodies and stakeholders.

Hawk-Eye's semi-automated offside technology, deployed in FIFA World Cup and UEFA Champions League competitions, employs 12-29 high-frame-rate cameras (50 fps minimum) with limb-tracking algorithms that analyze 29 body points per player, achieving reported accuracy >99% for spatial positioning within 3 cm tolerance. This system serves as the operational-accuracy benchmark for professional offside detection but requires dedicated camera infrastructure, proprietary calibration software, and specialized personnel, with estimated implementation costs exceeding \$3 million per stadium. The current YOLOv8-based approach trades spatial precision for accessibility, utilizing standard broadcast feeds without additional hardware, thereby enabling feasibility assessment in contexts where SAOT deployment is economically prohibitive. Performance differences reflect this infrastructure trade-off: Hawk-Eye's multi-camera triangulation resolves depth ambiguity and occlusion constraints inherent to monocular vision systems, while the single-camera approach accepts higher error rates (17% vs. <1%) in exchange for reduced implementation barriers. The systems address distinct market segments rather than competing directly: SAOT for elite professional competitions requiring maximal accuracy, single-

camera detection for training environments, lower-tier leagues, and broadcast analysis, where cost constraints dominate accuracy requirements. The proliferation of AI-driven sensor technologies in sports necessitates standardized validation methodologies, algorithmic transparency, and performance benchmarking across competitive contexts [25]. Such frameworks would enable systematic comparison between emerging detection systems and established officiating technologies.

System positioning: The approach serves as a low-cost, single-camera offside-detection prototype for contexts lacking SAOT infrastructure, rather than as a VAR replacement or an accuracy enhancement for elite competitions. Target applications comprise (i) referee training programs providing offline decision analysis with performance feedback, (ii) sub-elite competitive levels (regional leagues, youth competitions) without VAR access, and (iii) broadcast analysis for viewer education. The system does not address VAR operational deficiencies (communication protocols, multi-camera review, player involvement assessment), but it offers accessibility in resource-limited contexts where established technologies are economically prohibitive. Operational deployment would require resolution of real-time processing constraints not addressed in offline validation. Decision latency from frame capture to output classification must remain below 200 ms to inform referee decisions before play resumes, necessitating optimized inference pipelines and protocols for integrating broadcast feeds. Frame synchronization between ball contact detection and player position assessment introduces temporal uncertainty (± 33 ms at 30 fps broadcast rates), potentially causing misclassification when players move 0.5-1.0 m between frames during sprint sequences. Communication protocols for conveying automated outputs to match officials without disrupting match flow, integration with existing referee communication systems, and procedures for adjudicating disputed decisions remain unspecified, posing implementation barriers beyond algorithmic performance validation.

The interpretation and generalizability of these findings are subject to several limitations. The dataset of 240 test instances, while enabling initial validation, is limited compared to the vast number of real-world offside events, potentially affecting generalizability to less frequent scenarios. The single-camera analysis, although computationally efficient, inherently suffers from perspective-based limitations that can lead to occlusions, as evidenced by 18 false negatives during such events, suggesting that multi-camera systems could enhance reliability [26]. The binary classification approach (offside/not offside) omits the nuanced assessment of player involvement that is crucial to modern interpretation of the offside rule under FIFA regulations, thereby restricting the system's applicability to scenarios that require subjective judgment beyond spatial analysis.

Integrating the current system with multi-camera setups offers an immediate avenue for improved performance, particularly by mitigating occlusion-related false negatives through advanced sensor fusion. Developing capabilities for player-involvement assessment, potentially using machine learning to analyze movement, gaze, and spatial relationships, would address a key limitation by enabling comprehensive enforcement of the offside rule [27]. Real-time implementation studies and field testing in professional matches are needed to validate performance under operational conditions and inform integration with existing infrastructure.

4.1 Limitations

This study's interpretation requires acknowledgment of several methodological constraints. Dataset constraints impose interpretation boundaries: (1) Sample size ($n=400$, 240 test instances) limits generalization relative to deep learning standards requiring 2,000-10,000 instances for robust performance estimation across diverse scenarios. (2) Selection bias: Inclusion criteria (clear player visibility, minimal occlusion, optimal lighting) excluded 30-40% of broadcast footage representing operationally challenging conditions (adverse weather, suboptimal camera angles, player clustering). This bias likely inflates reported accuracy by 5-8 percentage points relative to unfiltered operational deployment, in which marginal-visibility scenarios predominate among controversial offside calls.

Second, single-camera analysis imposes fundamental spatial resolution constraints. Monocular depth estimation from perspective projection introduces systematic spatial uncertainty. Professional SAOT systems achieve <3 cm positional accuracy through multi-camera triangulation, while single-camera approaches exhibit estimated depth errors of 10-25 cm at typical broadcast distances (40-70 m from touchline), sufficient to cause misclassification in marginal offside situations (player separation <30 cm from offside line). Occlusion-related failures ($n=18$, 7.5% of test instances) further constrain the reliability of monocular detection. Multi-camera fusion architectures employed in professional VAR and semi-automated offside technology (SAOT) achieve superior spatial accuracy through triangulated depth estimation, a capability absent from the current implementation.

Third, extreme player tracking prioritizes computational efficiency over spatial completeness, thereby introducing systematic errors in compressed defensive formations (estimated to occur in 15-20% of professional scenarios) and in non-standard camera perspectives. The approach fails when the second-to-last defender occupies a central rather than an extreme horizontal position, thereby constraining its applicability to matches with stereotypical defensive geometry.

Fourth, the binary spatial classification omits FIFA Law 11 requirements: (i) Player involvement assessment (interfering with

play, gaining advantage, interfering with opponent) requires subjective judgment beyond positional analysis. (ii) Body part eligibility specification (arms excluded, other body parts with which the ball may be legally played included) demands limb-level segmentation absent from bounding box detection. (iii) Goalkeeper positional dynamics (second-to-last defender determination when goalkeeper advances) require role-specific tracking logic not implemented. (iv) Deflection scenarios (ball contact by defending player) necessitate causal analysis linking ball trajectory changes to player actions. These omissions preclude autonomous, rule-compliant decision-making, thereby limiting the system's utility to preliminary spatial flagging that requires human adjudication.

Fifth, generalizability to heterogeneous competitive contexts (different leagues, stadium configurations, broadcast standards, tactical systems) remains unvalidated. The system's performance under varying camera heights, zoom levels, frame rates, and match tempos requires empirical verification before deployment recommendations can extend beyond the sampled conditions.

Dataset restrictions impose ecological validity constraints: (i) Competitive context: UEFA Champions League matches exclusively (2021-2023 seasons), representing elite intensity (sprint distance 28-32 km/match, transition frequency 15-18 per match) with VAR/SAOT infrastructure present. (ii) Camera specifications: Sideline elevation 42-58 m, vertical angle 35-45°, fixed zoom, 1080p resolution, 30-50 fps. (iii) Tactical diversity: Predominantly 4-3-3 and 4-2-3-1 formations with high defensive lines (38-42 m from goal); low-block defenses (<25 m) and alternative systems (5-3-2, 3-5-2) are underrepresented. (iv) Venue sampling: Three stadiums only; generalizability to variable pitch dimensions, camera positioning, and lighting conditions unvalidated. (5) Processing latency: 185-220 ms measured offline; real-time broadcast integration latency untested. (6) Operational integration: Referee communication protocols, user interface design, and IFAB Law 11 compliance procedures are unaddressed.

The study lacks empirical validation of the impact of decision-making. Performance metrics reflect algorithmic classification accuracy against expert-annotated ground truth but do not assess whether system outputs improve referee decision quality, reduce error rates in operational contexts, or enhance officiating efficiency. Evaluation of automation-bias risks, communication-protocol effectiveness, and integration with existing referee workflows remains necessary before operational deployment claims can be substantiated. Finally, the study lacks ecological validation through real-time implementation testing, assessment of referee interactions, or integration with existing officiating workflows. Decision latency, communication protocols, and automation bias risks remain unaddressed, limiting conclusions regarding operational feasibility in professional match environments.

4.2 Practical Recommendations

Based on the research findings and performance analysis, professional football leagues should consider phased deployment strategies that begin with decision support rather than replacing human officials, allowing gradual integration that builds confidence and systematically addresses operational challenges. The system's 83% accuracy, while insufficient for autonomous professional implementation, may provide supplementary information in training contexts or lower-stakes competitive environments where human oversight maintains final decision authority, aligning with successful technology adoption patterns in other sports applications. Referee training integration requires empirical validation, absent from the current study design. Proposed application: post-match case analysis providing automated classifications for discussion without real-time influence on decisions. Critical research needs include assessing automation bias (overreliance on system outputs despite errors), measuring the impact of decision quality (pre- and post-intervention comparisons), and identifying thresholds for output confidence levels that warrant presentation to officials. Implementation without these validations risks degrading human judgment by inappropriately deferring to algorithmic classifications. Realistic implementation would require phased validation: (i) offline post-match analysis to establish reliability across diverse competitive contexts, (ii) controlled field trials in non-elite competitions with parallel human officiating to quantify decision concordance rates, (iii) referee training integration to assess usability and automation bias risks, and (iv) graduated deployment contingent on demonstrated decision quality improvements and regulatory approval. Technical infrastructure would require integrating broadcast feeds, specifying real-time processing hardware, and defining communication protocols to deliver outputs to match officials. These steps remain unvalidated in the current study, precluding immediate recommendations for professional implementation. A prospective single-camera deployment architecture would require: (i) a consumer-grade GPU (minimum 8 GB VRAM) connected to a standard broadcast capture card accepting 1080p input at 30–50 fps via HDMI or SDI; (ii) camera positioning at sideline elevation 42–58 m with a vertical angle of 35–45° to minimize perspective distortion at the offside threshold; (iii) end-to-end processing latency below 200 ms, achieved through local GPU inference (19–21 ms per frame) rather than cloud-based routing; and (iv) output delivery to assistant referees via existing wireless earpiece infrastructure, with a standardized protocol distinguishing high-confidence automated flags (confidence score >0.85) from advisory alerts requiring human adjudication. These specifications constitute prospective design requirements and have not been validated under live-match conditions.

Empirically-grounded development priorities: (i) Occlusion mitigation. Error analysis indicates 72% of false negatives ($n=13/18$) occurred when attackers were obscured by defenders at identical

depth (± 2 m), suggesting second sideline-oriented camera (perpendicular baseline, 90° offset from primary view) could resolve depth ambiguity through stereoscopic triangulation; (ii) multi-camera prototype evaluation using synchronized broadcast feeds to quantify depth estimation improvements, requiring triangulation algorithm development (estimated 6-month engineering timeline); (iii) Player involvement classification. FIFA Law 11 requires assessment beyond spatial positioning (interfering with play, gaining advantage, and interfering with opponents). Implementation pathway: positional data integration with expert-coded involvement labels (binary classification: involved/not involved) using logistic regression or shallow neural network, trained on minimum 500 annotated offside scenarios with certified referee ground truth; (iv) player involvement classification module integrating spatial analysis with gaze tracking and movement pattern recognition, representing substantial algorithmic complexity increase; (v) prospective validation through controlled field trials in semi-professional competitions to assess operational reliability and referee acceptance, contingent on regulatory approval and liability frameworks. Implementation feasibility decreases with system complexity; single-camera refinement offers the most pragmatic near-term development pathway. Compliance with FIFA's semi-automated offside technology standards, which specify < 3 cm spatial accuracy, < 0.5 -second decision latency, and a comprehensive player-involvement assessment would require an architectural redesign that incorporates limb-tracking algorithms, multi-camera sensor fusion, and a real-time communication infrastructure. These specifications exceed the capabilities of a single camera, positioning the current system as exploratory research rather than a pathway to a FIFA-compliant implementation.

5. CONCLUSION

This study demonstrates proof of concept for automated offside detection using extreme player tracking on single-camera broadcast footage, achieving computational efficiency (50 fps) suitable for resource-constrained competitive environments while maintaining moderate classification accuracy (83%). The system achieved moderate classification performance (83% accuracy) with computational efficiency suitable for single-camera broadcast integration. Performance levels indicate suitability for non-critical support applications, referee training, and resource-constrained environments, rather than for replacing established VAR infrastructure. While limitations remain, the evaluation confirms the feasibility of supplementary decision support in non-autonomous contexts in which final authority remains with human officials. By adapting modern computer vision techniques to the complexities of officiating, this work contributes to the growing field of sports technology. YOLOv8's anchor-free detection architecture and

optimized convolutional blocks enable real-time processing on standard broadcast hardware without specialized multi-camera infrastructure, thereby differentiating this approach from SAOT systems that require proprietary sensor arrays and from ensemble classifier methods with prohibitive inference latency. This computational accessibility enables the feasibility assessment of offside detection in lower-tier competitions, referee training programs, and broadcast analysis applications, where deployment of VAR infrastructure is financially or logistically constrained. The findings support prospective, phased implementation contingent on live-match validation, with final decision authority maintained by human officials throughout any deployment context.

Conflict of Interest

The authors declare that they have no commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

Data curation, A.A.F and S.B.B; Formal analysis, A.A.F and S.B.B; Funding acquisition, R.I.M.; Investigation, W.D., A.A.F S.B., and S.B.B; Methodology, A.E.O, H.I.C., W.D., A.A.F., I.D., and S.B.B; Project administration, A.E.O, A.A.F., H.I.C., and S.B.B; Software, A.A.F and S.B.B; Validation, A.E.O and W.D.; Writing – original draft, A.A.F., H.I.C., S.B., R.I.M., I.D., and S.B.B; Writing – review & editing, A.A.F., S.B.B., H.I.C., S.B., R.I.M., I.D., W.D and A.E.O. All authors have read and agreed to the published version of the manuscript.

Funding

The authors declare that no financial support was received for the research, authorship, and/or publication of this article.

Acknowledgments

The authors thank all the experts of the Qatar Football Federation for their participation in this work.

Data Availability Statement

The raw data supporting the conclusions of this article will be made available by the corresponding authors on request.

REFERENCES

1. Spitz J, Wagemans J, Memmert D, Williams AM, Helsen WF: **Video assistant referees (VAR): The impact of technology on decision making in association football referees.** *Journal of Sports Sciences* 2021, **39**(2):147-153.

2. Oudejans RR, Verheijen R, Bakker FC, Gerrits JC, Steinbrückner M, Beek PJ: **Errors in judging 'offside' in football.** *Nature* 2000, **404**(6773):33.
3. Unkelbach C, Memmert D: **Crowd noise as a cue in referee decisions contributes to the home advantage.** *J Sport Exerc Psychol* 2010, **32**(4):483-498.
4. Nevill AM, Balmer NJ, Mark Williams A: **The influence of crowd noise and experience upon refereeing decisions in football.** *Psychology of Sport and Exercise* 2002, **3**(4):261-272.
5. Carlos L-P, Ezequiel R, Anton K: **How does Video Assistant Referee (VAR) modify the game in elite soccer?** *International Journal of Performance Analysis in Sport* 2019, **19**(4):646-653.
6. Nagle T, Sammon D, Pope A: **Exploring the socio-technical dynamics of VAR implementation and use.** *Journal of Decision Systems* 2024, **33**(sup1):47-62.
7. Rogerson M, Knight D, Scherer R, Jones B, McManus C, Waterworth S, Murray K, Hope E: **Meta-analysis of the effects of VAR on goals scored and home advantage in football.** *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, **0**(0):17543371241242914.
8. Alhassan I, Sammon D, Daly M, Wibisono A, Kasraian L, Nagle T, Heavin C, Dennehy D, Zamani E, Qaffas A: **The use of open, axial and selective coding techniques: A literature analysis of IS research.** 2023.
9. Naik BT, Hashmi MF, Bokde ND: **A comprehensive review of computer vision in sports: Open issues, future trends and research directions.** *Applied Sciences* 2022, **12**(9):4429.
10. Terven J, Córdova-Esparza D-M, Romero-González J-A: **A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS.** *Machine Learning and Knowledge Extraction* 2023, **5**(4):1680-1716.
11. Panse N, Mahabaleshwarkar A: **A dataset & methodology for computer vision based offside detection in soccer.** In: *Proceedings of the 3rd International Workshop on Multimedia Content Analysis in Sports: 2020*; 2020: 19-26.
12. Jocher G, Chaurasia A, Qiu J: **YOLO by Ultralytics.** In.; 2023.
13. Maaroo MKA, Bouhleb MS: **Real-Time Object Detection Using YOLO-8 Model: A Drone-Based Approach.** *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2025, **16**(1):190-204.
14. Harwood H: **Investigating Video Assistant Referees (VARs): Does contextual information influence VAR decision-accuracy and decision-severity in football.** University of Essex; 2024.
15. Postupaiev S, Damaševičius R, Maskeliūnas R: **Real-Time Camera Operator Segmentation with YOLOv8 in Football Video Broadcasts.** *AI* 2024, **5**(2):842-872.
16. Hermens F: **Automatic object detection for behavioural research using YOLOv8.** *Behavior Research Methods* 2024, **56**(7):7307-7330.
17. Li C, Zhu Y, Zheng M: **A multi-objective dynamic detection model in autonomous driving based on an improved YOLOv8.** *Alexandria Engineering Journal* 2025, **122**:453-464.

18. Fedorov VA: **Railway Infrastructure Detection Based on YOLOv8 with NPU Acceleration.** *Doklady Mathematics* 2024, **110**(1):S42-S48.
19. Ayachi R, Said Y, Afif M, Alshammari A, Hleili M, Ben Abdelali A: **Assessing YOLO models for real-time object detection in urban environments for advanced driver-assistance systems (ADAS).** *Alexandria Engineering Journal* 2025, **123**:530-549.
20. Li A, Li D, Wang A: **A Two-Stage YOLOv5s-U-Net Framework for Defect Localization and Segmentation in Overhead Transmission Lines.** *Sensors* 2025, **25**(9):2903.
21. Rabee A, Anwar Z, AbdelMoety A, Abdelsallam A, Ali M: **Comparative analysis of automated foul detection in football using deep learning architectures.** *Sci Rep* 2025, **15**(1):14236.
22. Dowsett B: **AI is helping referee games in major sports leagues, but limitations remain.** *Scientific American* <https://www.scientificamerican.com/article/ai-is-helping-referee-games-in-major-sports-leagues-but-limitations-remain> 2024.
23. Fakhurroja H, Fashihullisan AA, Bangkit H, Pramesti D, Ismail N, Mahardiono NA: **A Vision-Based System: Detecting Traffic Law Violation Case Study of Red-Light Running Using Pre-Trained YOLOv8 Model and OpenCV.** In: *2024 IEEE International Conference on Smart Mechatronics (ICSMech): 2024*: IEEE; 2024: 1-6.
24. Kolbinger O, Lames M: **Scientific approaches to technological officiating aids in game sports.** *Current Issues in Sport Science* 2017, **2**:1-10.
25. Dhahbi W, Chamari K: **The Algorithmic Athlete: A Call to Standardize Assessment of Sensor Technologies and Artificial Intelligence.** *Int J Sports Physiol Perform* 2026:1-2.
26. Amosa TI, Sebastian P, Izhar LI, Ibrahim O, Ayinla LS, Bahashwan AA, Bala A, Samaila YA: **Multi-camera multi-object tracking: A review of current trends and future advances.** *Neurocomputing* 2023, **552**:126558.
27. Elstak I, Paul S, and McLean S: **Artificial intelligence applications in the football codes: A systematic review.** *Journal of Sports Sciences* 2024, **42**(13):1184-1199.

TABLES

Table 1. Dataset characteristics and distribution

Characteristic	Category	Count	Percentage
Total Samples	Images/Video Segments	400	100%
	Clear Offside	180	45%
Match Scenarios	No Offside	160	40%
	Marginal Cases	60	15%
Playing Conditions	Daylight	240	60%
	Evening/Artificial Light	160	40%
Camera Angles	Side View	280	70%
	Elevated View	120	30%
Team Combinations	High Contrast Jerseys	320	80%
	Similar Color Jerseys	80	20%
Data Partition	Training Set	160	40%
	Validation Set	240	60%

Note: Comprehensive breakdown of dataset characteristics showing distribution across key variables affecting automated offside detection performance. Model validation was performed using 5-fold cross-validation on the training partition rather than a held-out validation set, thereby maximizing test-set allocation for robust performance assessment.

Table 2. Confusion matrix analysis for binary offside classification performance

	Predicted Not Offside	Predicted Offside
Actual Not Offside	70 (TN)	22 (FP)
Actual Offside	18 (FN)	130 (TP)

TN = True Negative, TP = True Positive, FP = False Positive, FN = False Negative. Analysis **Note:** based on 240 test instances across diverse match scenarios.

Caption: Comprehensive confusion matrix summarizing automated offside detection system performance across 240 test instances. True Negative (TN) = 70 represents correct identification of valid player positions; True Positive (TP) = 130 indicates accurate detection of genuine offside violations; False Positive (FP) = 22 denotes incorrect flagging of valid positions as offside; False Negative (FN) = 18 represents missed detection of actual violations. Matrix elements demonstrate the system's strength in positive detection (TP is the largest category) while maintaining an acceptable false-positive rate. Overall accuracy = 83.0%, precision = 85.0%, recall = 87.0%, F1-score = 86.0%. Statistical analysis confirms performance significantly exceeds random classification baseline ($p < 0.001$).

Class-specific metrics: Offside class: precision 85.5%, recall 87.8%, F1-score 86.6%. Not-Offside class: precision 79.5%, recall 76.1%, F1-score 77.7%. Macro-average F1: 82.2%; weighted F1: 83.3%.

Table 3. Per-fold cross-validation performance metrics on training partition (n=160)

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Fold 1	82.1%	84.0%	86.5%	85.2%
Fold 2	84.2%	86.1%	88.0%	87.0%
Fold 3	83.0%	85.0%	87.4%	86.2%
Fold 4	82.7%	84.6%	86.9%	85.7%
Fold 5	84.5%	86.3%	88.2%	87.2%
Mean $\pm$ SD	83.3 $\pm$ 1.0%	85.2 $\pm$ 1.0%	87.4 $\pm$ 0.7%	86.3 $\pm$ 0.8%

Note: Cross-validation was applied exclusively to the training partition (n=160) for hyperparameter optimization. Final performance estimates are reported on the independent held-out test set (n=240; Table 2). No Bayesian inference was employed; confidence intervals were estimated via BCa bootstrap resampling (1,000 iterations).




Figure 1: Algorithmic workflow for automated offside detection using YOLOv8 architecture.




Figure 2: Model training and validation performance curves across training epochs. **Caption:** Training and validation accuracy progression throughout the model development process, demonstrating convergence characteristics and generalization capability. The parallel trajectories of the training and validation curves indicate successful learning without overfitting, with both metrics stabilizing at optimal performance levels. Training accuracy approaches 87%, while validation accuracy converges to 83%, with a narrow performance

gap suggesting appropriate model complexity and robust generalization to unseen data. The smooth convergence pattern confirms optimal hyperparameter selection and the effectiveness of the training protocol for adapting the YOLOv8 architecture to football offside detection tasks.

ARTICLE IN PRESS



Figure 3: Real-time object detection performance on professional match footage.

Caption: Demonstration of YOLOv8 system performance on actual match imagery showing detection accuracy for primary object classes: Team 1 players, Team 2 players, and ball identification. Bounding boxes indicate successful object detection with confidence scores displayed for each classification. The image illustrates typical detection performance under standard viewing conditions and highlights the challenges posed by player occlusion. Visual analysis reveals consistently high detection confidence across all object categories, with particular attention to situations where ball positioning between players increases detection complexity, which may affect subsequent offside analysis accuracy.




Figure 4: Successful automated offside detection in a professional match context.

Caption: Representative example of accurate offside violation detection showing attacking player positioned beyond the defensive line. This practical demonstration confirms that laboratory performance validation

effectively translates to real-world conditions, with processing completed within an acceptable timeframe for near-real-time implementation.

ARTICLE IN PRESS



Figure 5: Correct negative classification in a complex player positioning scenario

Caption: Example of accurate negative offside determination where the system correctly identified that no violation occurred despite potentially ambiguous player positioning. Multiple players in proximity to the offside threshold create visual complexity that might challenge human officials, yet the automated system-maintained

classification accuracy. This negative case validation demonstrates the system's discriminative capability and reliability across the full decision spectrum, supporting its potential as a decision support tool for professional officiating applications.

ARTICLE IN PRESS