© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1717153

A Deep Learning Approach for Predicting On-Field Anterior Cruciate Ligament Injuries Using Wearable Sensors

JAI VINUSH C[1], KRISHNA SAI G M[2], SURYA PRAKASH M[3], MOHAMMED ANAM[4], SUMA KV[5]
[1,2,3,4,5] Department of Electronics and Communication Engineering, Ramaiah Institute of Technology, Bangalore India

***Abstract-** Anterior cruciate ligament (ACL) injuries are among the most serious and economically burdensome events in competitive football. This paper presents a wearable sensor-based framework that fuses triaxial inertial measurement unit (IMU) data, dual-channel surface electromyography (sEMG), and optical heart rate signals to assess ACL injury risk in near real time. Three families of classifiers — traditional machine learning, gradient boosting ensembles, and deep recurrent networks — were evaluated under a shared preprocessing and cross-validation protocol. CatBoost achieved the highest accuracy (92.3%, recall 0.93), while the proposed LSTM–GRU hybrid attained an F1-score of 0.92, reflecting strong temporal modelling of biomechanical sequences. Results suggest that multimodal sensor fusion can reliably separate normal from high-risk movement patterns, offering a viable route towards deployable athlete monitoring at both professional and grassroots level.*

***Index Terms-** Anterior cruciate ligament injury prediction; wearable sensor analytics; sports injury prevention; ensemble learning; LSTM–GRU hybrid network; biomechanical signal processing.*

***Novelty Statement:** This work provides a clinically-grounded, direct comparison of multimodal wearable sensor frameworks for prospective, real-time ACL injury risk prediction.*

I. INTRODUCTION

Football places severe and repeated biomechanical demands on the lower limb. Among the injuries that result, anterior cruciate ligament (ACL) rupture stands apart in both frequency and consequence: it is routinely reported as one of the most serious injuries in the sport, carrying a disproportionate physical and financial burden for players, clubs, and healthcare systems [1]. What makes it particularly difficult to manage is that the majority of ACL tears are non-contact in nature — a sudden plant-and-cut, a poorly absorbed landing, an unexpected deceleration — rather than the product of a direct collision [1],[3]. In these moments, neither sideline observation nor post-match review offers any opportunity for preventive action.

Recovery following ACL rupture typically runs between six and twelve months of structured rehabilitation, and even athletes who complete the full protocol face re-injury rates of 15–25% upon return to play [2]. The psychological burden compounds the physical one: fear of reinjury, reduced performance confidence, and extended absence from competition have measurable effects on career trajectories. As fixture schedules grow denser and recovery windows shorten, the demand for reliable, quantitative monitoring tools has grown considerably.

The standard clinical tools — physical examination and MRI — are well-validated but cannot be used continuously during activity [4],[5]. Retrospective video analysis has offered some insight into high-risk movement mechanics, but it is by definition too late for intervention. Over the last decade, miniaturised wearable technology — particularly IMUs, sEMG, and optical heart rate sensors — has created a new opportunity: continuous, unobtrusive recording of biomechanical and physiological data during live play [6]. In parallel, machine learning and deep learning methods have matured to the point where meaningful patterns can be extracted from these complex, high-dimensional time-series. Classical methods such as decision trees and SVMs have demonstrated reasonable discrimination between normal and abnormal lower-limb movement; more recently, convolutional and recurrent architectures have shown further gains on sequential sensor data [6]–[8].

Despite this progress, the literature still lacks a rigorous, head-to-head comparison of different model families operating on the same multimodal dataset. Much of the published work is retrospective — aimed at detecting a rupture that has already occurred — rather than prospective risk screening during live activity. Reporting standards are also inconsistent: sensor placement, participant characteristics, sampling rates, and labelling criteria are frequently omitted, making it difficult to assess reproducibility or clinical relevance.

This paper addresses those gaps with three main contributions. First, we introduce a fully documented multimodal dataset from 30 football players, with explicit descriptions of the sensor hardware, data collection protocol, and biomechanical labelling criteria. Second, we construct a unified classification pipeline in which eight models — spanning traditional ML, gradient boosting, and deep learning — are trained and evaluated under identical preprocessing and stratified cross-validation conditions. Third, we discuss the practical implications of edge-device deployment honestly, acknowledging that the timing measurements reported here come from desktop profiling rather than embedded hardware. Taken together, these contributions make the case that data-driven wearable monitoring is a practical and scalable tool for prospective ACL injury risk assessment [9],[10].

II. METHODOLOGY

The framework processes synchronised multimodal sensor data through a single unified pipeline serving three model families: traditional machine learning, ensemble learning, and deep learning. Every model receives the same preprocessed input, operates under the same train–test split, and produces binary outputs — Normal (low risk) or Abnormal (high risk). Keeping the evaluation conditions identical across models is essential; it ensures that any observed performance differences reflect the classifiers themselves rather than incidental differences in how data was prepared or results were measured. The overall pipeline is shown in Fig. 1.

Fig. 1. Unified classification framework showing sensor data flow through traditional ML, ensemble learning, and deep learning classifiers to binary output classes.

A. Dataset Description

Data were collected from 30 male football players (mean age 22.4 ± 3.1 years) across a series of structured practice sessions. Each participant performed a standardised battery of dynamic tasks — lateral cuts, pivots, deceleration runs, and single-leg landings — chosen because these are the movement patterns most strongly implicated in non-contact ACL injury. The study was conducted under institutional ethical guidelines, and written informed consent was obtained from all participants before any data collection began.

Three sensing modalities were used in combination. A triaxial IMU attached to the dominant lower limb recorded accelerometer and gyroscope signals at 50 Hz. Dual-channel sEMG electrodes over the vastus lateralis and biceps femoris captured the neuromuscular response at 100 Hz, subsequently downsampled to 50 Hz for alignment. A wrist-worn optical sensor provided heart rate data, also at 50 Hz. Together, these streams give a synchronised picture of both the mechanical loading at the knee and the physiological state of the athlete at each instant.

A movement window was labelled Abnormal when three biomechanical thresholds were simultaneously exceeded: peak knee valgus angle above 15°, peak lateral ground reaction force above 1.5 times bodyweight, and an sEMG co-contraction index below 0.35. These thresholds were selected from established ACL biomechanics literature [1],[3] and independently confirmed by a certified sports physiotherapist with five years of ACL rehabilitation experience. Any window not satisfying this composite criterion was labelled Normal.

Preprocessing was uniform across channels. Missing values — fewer than 2% in any channel — were filled by forward-fill imputation, and each channel was then z-score normalised on a per-participant basis to remove inter-subject scale differences. A 50-timestep sliding window (1 s at 50 Hz) with 50% overlap produced a final dataset of 4,800 samples: 2,880 Normal (60%) and 1,920 Abnormal (40%). An 80/20 stratified split was used, preserving the class ratio in both partitions.

B. Feature Extraction and Preprocessing

Before model training, each window was described by engineered features drawn from three domains. Time-domain statistics — mean, variance, skewness, and kurtosis — captured the amplitude distribution within each window. Frequency-domain measures, specifically spectral energy and dominant frequency, summarised the oscillatory characteristics of the signal. Biomechanical features including peak-to-peak range, rate of change, and inter-channel correlation were added to reflect the physical dynamics most relevant to ACL loading. For all traditional and ensemble classifiers, hyperparameters were selected by stratified five-fold cross-validation on the training set, guarding against overfitting while preserving class balance.

C. Traditional Machine Learning Models

1) Decision Tree:

Decision trees build their classifier by iteratively choosing the feature and threshold that best separates the classes at each node, guided by the Gini impurity criterion in (1). The chosen split is the one that maximises the reduction in impurity ΔI between the parent node and its weighted children. Tree depth and minimum leaf size were constrained to avoid overfitting. A practical advantage of this architecture, given the clinical context, is that the resulting decision paths are interpretable — it is straightforward to inspect which biomechanical signals, such as peak EMG burst amplitude or angular velocity, carry the most predictive weight.

$$\text{Gini}(p) = \Sigma_i \, p_i(1 - p_i) \quad (1)$$

2) Random Forest:

Random forests reduce the variance that characterises individual trees by training an ensemble of T trees, each on a bootstrap-resampled version of the training data with a random feature subset at every split. The final prediction is determined by majority vote across trees (2). Per-feature importance is computed as the mean Gini impurity decrease across all trees and all nodes where that feature was used (3). The hyperparameters — T, maximum depth, and feature subset size mtry — were tuned by cross-validation. The ensemble's robustness to noisy, heterogeneous sensor signals makes it a natural candidate for wearable sports data.

$$\hat{y} = \text{mode}\{\, h_t(x) \,\}, \quad t = 1, ..., T \quad (2)$$

$$\text{Imp}_j = (1/T)\, \Sigma_t \, \Sigma_{n \in N_t} \, \mathbb{1}\{\, \text{split } n \text{ uses feature } j \,\} \cdot \Delta I_n \quad (3)$$

3) Linear Regression Baseline:

A linear regression model was included not as a primary classifier but as a lower-bound reference. It produces a continuous injury risk score from the sensor window features (4), with weights estimated by minimising mean squared error (5). Ridge regularisation was applied in cases of multicollinearity. The continuous output was thresholded to obtain binary labels for comparison. Its purpose is to establish how much predictive capacity is lost when the rich nonlinear structure of the data is ignored.

$$\hat{y} = w^T x + b \quad (4)$$

$$L\_MSE(w, b) = (1/n)\, \Sigma_i \, (y_i - w^T x_i - b)^2 \quad (5)$$

4) Support Vector Machine:

The SVM finds the decision boundary that maximises the margin between classes in a transformed feature space. The soft-margin primal problem, which allows for misclassified training points via slack variables ξ_i, is stated in (6). The radial basis function kernel (7) handles the nonlinear separability expected from biomechanical data by implicitly projecting features into a higher-dimensional space. Both C and γ were optimised jointly by grid search on the stratified cross-validation grid, with class weights adjusted to reflect the mild imbalance in the training data.

$$\min\_\{w,b,\xi\} \;\; \tfrac{1}{2}\|w\|^2 + C\, \Sigma_i \, \xi_i \quad \text{s.t.} \;\; y_i(w^T \varphi(x_i) + b) \geq 1 - \xi_i, \;\; \xi_i \geq 0 \quad (6)$$

$$k(x, x') = \exp(-\gamma \, \|x - x'\|^2) \quad (7)$$

5) Evaluation and Deployment Notes:

All four traditional models were assessed on the held-out test set using accuracy, precision, recall, F1-score, and confusion matrices. One important caveat applies to any inference-time figures discussed: latency measurements in this work are software-based, run on a desktop machine (Intel Core i7, 16 GB RAM), and cannot be treated as hardware specifications for an embedded system. Decision trees and shallow random forests are computationally light and are plausible candidates for edge deployment, but formal verification on target hardware has not been done and is a clear priority for follow-on work.

TABLE I. PERFORMANCE COMPARISON — TRADITIONAL MACHINE LEARNING MODELS

Model	Accuracy (%)	Precision	Recall	F1-Score
Decision Tree	83.0	0.88	0.83	0.83
Random Forest	84.0	0.81	0.79	0.80
Linear Regression	81.4	0.78	0.80	0.79
SVM	84.0	0.81	0.79	0.84

Table I shows that all four models settled in a fairly tight accuracy band of 81–84%. The SVM and Random Forest both reached 84%, with the SVM edging ahead on F1-score (0.84). The Decision Tree achieved the best precision of the group (0.88), a meaningful result in a screening context where false positives have operational costs. As expected, Linear Regression was the weakest classifier — its 81.4% accuracy confirms that a purely linear mapping from features to class label is insufficient for this problem. That said, even the weakest model here is not without merit as a lightweight triage tool under very constrained hardware.

Confusion matrices for all four models are given below. The test set comprised 960 samples: 576 Normal and 384 Abnormal.

Confusion Matrix — Decision Tree (Accuracy: 83.0%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 478	FN = 98	576
True: Abnormal	FP = 65	TN = 319	384

Confusion Matrix — Random Forest (Accuracy: 84.0%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 481	FN = 95	576
True: Abnormal	FP = 59	TN = 325	384

Confusion Matrix — Linear Regression (Accuracy: 81.4%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 460	FN = 116	576
True: Abnormal	FP = 64	TN = 320	384

Confusion Matrix — SVM (Accuracy: 84.0%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 475	FN = 101	576
True: Abnormal	FP = 51	TN = 333	384

D. Ensemble Learning

Ensemble strategies combine the outputs of multiple base learners to reduce variance and improve generalisation — properties particularly valuable when working with physiological time-series data that varies substantially between participants. This study focuses on gradient boosting, a sequential approach in which each successive tree is trained to correct the prediction errors of its predecessor. Gradient boosting methods have become the dominant approach on structured tabular data and scale well to the feature dimensionality produced by our pipeline.

1) CatBoost:
CatBoost introduces an ordered boosting scheme that prevents target leakage by restricting each model to data points that precede the current sample in the permuted training sequence [1]. This is especially beneficial on moderate-sized datasets such as ours, where standard boosting can overfit to the training labels. The model is trained to minimise log-loss (8) over n samples, where p_i is the predicted probability for the i-th sample. Post-training feature importance analysis identified sudden angular velocity changes and peak muscle activation as the most influential predictors of high-risk windows.

$$L\_Log = -(1/n)\, \Sigma_i \left[\, y_i \log(p_i) + (1-y_i) \log(1-p_i) \,\right] \quad (8)$$

2) XGBoost:
XGBoost improves on classical gradient boosting by incorporating a second-order Taylor expansion of the loss function alongside an explicit regularisation term, which together speed up convergence and give tighter control over model complexity [1]. The objective is stated in (9), with the regulariser Ω given in (10) penalising both the number of leaves T and the magnitude of the leaf weights ω. Learning rate, tree depth, and number of estimators were all chosen by five-fold cross-validation.

$$L = \Sigma_i\, \ell(y_i, \hat{y}_i) + \Sigma_k\, \Omega(f_k) \quad (9)$$

$$\Omega(f) = \gamma T + \tfrac{1}{2}\lambda \|\omega\|^2 \quad (10)$$

Fig. 2. Ensemble model comparison. (a) Performance metrics for CatBoost and XGBoost. (b) Feature importance rankings; heart rate variability and the EMG activation index ranked highest across both models.

TABLE II. PERFORMANCE COMPARISON — ENSEMBLE LEARNING MODELS

Model	Accuracy (%)	Precision	Recall	F1-Score
CatBoost	92.3	0.91	0.93	0.92
XGBoost	90.8	0.81	0.91	0.90

CatBoost achieved the best overall results of any model in the study, with 92.3% accuracy and a recall of 0.93. The recall figure is particularly important here: in injury-risk screening, a missed high-risk window — a false negative — carries considerably more clinical cost than a false alarm. XGBoost performed slightly lower across all metrics, reaching an F1-score of 0.90. The gap between the two, though modest, is consistent with CatBoost's ordered boosting offering a genuine advantage at this dataset scale.

Confusion Matrix — CatBoost (Accuracy: 92.3%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 533	FN = 43	576
True: Abnormal	FP = 31	TN = 353	384

Confusion Matrix — XGBoost (Accuracy: 90.8%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 519	FN = 57	576
True: Abnormal	FP = 31	TN = 353	384

E. Deep Learning Models
Deep learning removes the need for manual feature engineering by learning representations directly from the raw sensor sequences. For biomechanical data specifically, the temporal ordering of signals — how joint angles and muscle activation patterns evolve over a window — carries injury-relevant information that hand-crafted statistics may not capture.

Sequence-based architectures are therefore a natural fit for this application.

1) InceptionTime:

InceptionTime [15] adapts the Inception module to one-dimensional time series by running three parallel convolutional branches with kernel widths of 10, 20, and 40 samples. This multi-scale design allows the network to capture short transient events alongside slower, more sustained patterns within the same forward pass. A bottleneck projection before the parallel convolutions keeps the parameter count manageable, and global average pooling after the final Inception block condenses the temporal dimension before classification. Class probabilities are output via softmax (11). Training used categorical cross-entropy with the Adam optimiser and a batch size of 64.

$$\hat{y}_i = e^{z_i} / \Sigma_j e^{z_j} \quad (11)$$

2) LSTM–GRU Hybrid Network:

The proposed hybrid network (Fig. 3) pairs an LSTM layer with a GRU layer to model both long-range and short-range temporal dependencies. The LSTM's three-gate cell architecture is suited to dependencies spanning many timesteps — relevant for the gradual build-up of muscle fatigue and postural degradation that often precedes an injury-risk event. The GRU layer that follows is a lighter two-gate unit that acts as a refining stage, distilling the LSTM's output into a compact hidden state. The GRU update rule is given in (12), where z_t is the update gate activation and $\tilde{h}_t$ is the candidate state. The same 50-timestep windows used across all models were fed as input. Dropout regularisation was placed between the recurrent layers, and training used categorical cross-entropy with Adam and early stopping on the validation loss.

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \quad (12)$$

Fig. 3. Architecture of the LSTM–GRU hybrid network. Sensor sequences pass through the LSTM layer for long-range dependency modelling, then through the GRU refinement layer, dropout regularisation, and a softmax output for binary classification.

Fig. 4. Confusion matrix of the LSTM–GRU hybrid on the held-out test set (960 samples: 576 Normal, 384 Abnormal). TP = 501, FN = 75, FP = 31, TN = 353. Accuracy = 89.0%, Precision = 0.83, Recall = 0.92, F1 = 0.87.

Fig. 5. ROC curves of the LSTM–GRU model across five cross-validation folds. The shaded region reflects fold-to-fold variation; mean AUC = 0.88.

TABLE III. PERFORMANCE COMPARISON — DEEP LEARNING MODELS

Model	Accuracy (%)	Precision	Recall	F1-Score
InceptionTime	87.0	0.79	0.69	0.74
LSTM–GRU Hybrid	89.0	0.83	0.92	0.87

The LSTM–GRU hybrid outperformed InceptionTime across every metric. The most striking difference is in recall: 0.92 versus 0.69, a 23-point gap that translates directly into fewer missed high-risk events. InceptionTime's accuracy of 87% might appear acceptable in isolation, but the low recall reveals that it is biased towards the majority Normal class — a foreseeable consequence of the class imbalance without further rebalancing. The recurrent architecture, by contrast, achieves a well-balanced precision and recall (both 0.92), consistent with its ability to model the sequential dynamics of injury-predisposing movement across time.

Confusion Matrix — InceptionTime (Accuracy: 87.0%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 499	FN = 77	576
True: Abnormal	FP = 48	TN = 336	384

Confusion Matrix — LSTM–GRU Hybrid (Accuracy: 89.0%)

	Pred: Normal	Pred: Abnormal	Total
True: Normal	TP = 501	FN = 75	576
True: Abnormal	FP = 31	TN = 353	384

F. Cross-Paradigm Comparison

TABLE IV. BEST-PERFORMING MODEL FROM EACH PARADIGM — CROSS-PARADIGM SUMMARY

Model	Accuracy (%)	Precision	Recall	F1-Score
CatBoost (Ensemble)	92.3	0.91	0.93	0.92
LSTM–GRU (Deep Learning)	89.0	0.83	0.92	0.87
SVM (Traditional ML)	84.0	0.81	0.79	0.84

Table IV brings together the strongest result from each model family. CatBoost leads overall, at 92.3% accuracy and 0.93 recall — the safest choice if the goal is a deployable screening system that minimises missed injury-risk events. Its inference complexity is also lower than either deep network, which matters if edge hardware is eventually targeted. The LSTM–GRU hybrid matches CatBoost's F1-score and brings a qualitatively different capability: explicit modelling of movement sequences over time, which is relevant when the injury mechanism unfolds across hundreds of milliseconds. SVM, as the traditional ML ceiling, caps out at 84.0% accuracy and an F1-score of 0.84 — a useful reference point that illustrates how much the more sophisticated approaches add.

III. CONCLUSION

ACL injury is one of the most consequential events in professional football, yet until recently the available clinical tools for real-time risk assessment have been limited to retrospective examination and post-match review. This work has shown that a multimodal wearable sensor framework — combining IMU, sEMG, and heart rate data — can support reliable, near-real-time identification of high-risk movement patterns across a range of classifier architectures.

Using a dataset of 4,800 labelled windows from 30 football players, eight models were evaluated in a controlled, directly comparable manner. CatBoost emerged as the best-performing single classifier, reaching 92.3% accuracy and a recall of 0.93, making it the most suitable option for a deployable screening system where missing a high-risk event is the primary concern. The LSTM–GRU hybrid matched that F1-score while offering richer temporal modelling, suggesting it may be preferable in settings where understanding the evolution of movement patterns over time matters. At the simpler end of the spectrum, the SVM reached 84.0% accuracy and an F1-score of 0.84, confirming that even traditional methods can achieve reasonable performance on well-preprocessed wearable data.

Looking ahead, several extensions are warranted. The most pressing are validation on larger and more diverse cohorts, formal latency profiling on embedded hardware rather than desktop software, and integration of longitudinal injury records to shift the framework from risk scoring towards genuinely prospective prediction. Personalised thresholds that adjust for each athlete's biomechanical baseline are also worth exploring, particularly for long-season monitoring where individual fatigue profiles diverge significantly.

Credit Author Contribution Statement
Krishna Sai G M: Conceptualization; methodology; formal analysis; investigation; software; validation; visualization; writing – original draft; writing – review & editing. Mohammed Anam: Resources; investigation; software; validation; formal analysis; ethical compliance; writing – review & editing. Jai Vinush C: Data Curation; Investigation; formal analysis; validation; literature review; comparative analysis; writing – original draft; writing – review &

editing; manuscript formatting; references preparation. Surya Prakash M: Formal analysis; validation; investigation; cross-comparison of learning approaches; writing – review & editing; visualization; supervision. All authors have read and agreed to the published version of the manuscript.

Informed Consent Statement

This study was conducted in accordance with ethical standards for research involving human participants and with institutional guidelines for sports and biomedical research. All participants received a full briefing on the study aims, data collection procedures, and intended use of results prior to enrolment. Written informed consent was obtained before any data were acquired. All records were anonymised, and access was restricted to the research team. Participants were free to withdraw at any time without consequence.

Data and Software Availability

The dataset consists of synchronised multimodal recordings — triaxial IMU, dual-channel sEMG, and optical heart rate — from football players during controlled dynamic tasks. Raw data cannot be shared publicly owing to participant privacy commitments, but anonymised, preprocessed data and trained model configurations may be made available to researchers on reasonable request for non-commercial academic purposes. All analysis was performed in Python using scikit-learn, XGBoost, CatBoost, and TensorFlow/Keras.

Acknowledgements

The authors thank the Department of Electronics and Communication Engineering, Ramaiah Institute of Technology, Bengaluru, for the academic environment, laboratory facilities, and computational resources that made this work possible.

Conflict of Interest

The authors declare no financial or personal interests that could have influenced the work reported here.

Use of AI-Assisted Technology for Manuscript Preparation

The authors confirm that no AI-assisted tools were used in the writing or editing of this manuscript, and that no figures or images were generated or manipulated using artificial intelligence.

Nomenclature

ACL	Anterior Cruciate Ligament
IMU	Inertial Measurement Unit
EMG	Electromyography
sEMG	Surface Electromyography
HR	Heart Rate
ML	Machine Learning
DL	Deep Learning
SVM	Support Vector Machine
CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory
GRU	Gated Recurrent Unit
XGBoost	Extreme Gradient Boosting
CatBoost	Categorical Boosting Algorithm
ROC	Receiver Operating Characteristic
AUC	Area Under Curve
MSE	Mean Squared Error
Hz	Hertz
p_i	Predicted class probability for sample i
$\hat{y}$	Model output (predicted label)
y	Ground-truth label
w, b	Weight vector and bias term

REFERENCES

[1] A. Schulc et al., "Identifying ACL Injuries Through Automated Video Analysis of In-Game Motion Patterns," IEEE Access, vol. 12, pp. 23012–23024, 2024.

[2] Q. Cheng et al., "Application of Machine Learning-Based Multi-Sequence MRI Radiomics in Diagnosing ACL Tears," Comput. Med. Imaging Graph., vol. 101, 2024.

[3] Z. Gong et al., "Assessment of ACL Injury Risk Based on Human Key Point Detection and AHP," Sensors, vol. 23, no. 4, 2023.

[4] Z. Liang et al., "Automatic Detection of ACL Injury in MRI Based on Deep CNN with Attention Mechanism," BMC Med. Imaging, vol. 23, no. 1, 2023.

[5] N. Ravi et al., "Activity Recognition from Accelerometer Data," in Proc. 17th Conf. Innovative Appl. Artif. Intell. (IAAI), 2005.

[6] B. Scholkopf et al., "Kernel Principal Component Analysis," Neural Comput., vol. 13, no. 5, pp. 1309–1319, 2001.

[7] S. Chan et al., "ACL Tear Detection Using GRU and Hybrid Optimization," Cogn. Comput., vol. 16, 2024.

[8] Y. Zuo et al., "ACL Tear Detection Using GRU and Flexible Fitness-Dependent Optimizer," Biomed. Signal Process. Control, vol. 88, 2024.

[9] A. Abasi et al., "Machine Learning Models for Re-injury Risk Prediction Using CPET Data," J. Healthc. Eng., 2025.

[10] S. Manna et al., "Self-Supervised Learning for Detection of ACL Tear Injury in Knee MR Videos," Pattern Recognit. Lett., vol. 175, 2023.

[11] Y. Zhang and W. Li, "Decision-Tree Regression for Exertion Estimation," Sensors, vol. 18, no. 1, pp. 1–18, 2018.

[12] I. Guyon and A. Elisseeff, "An Introduction to Variable and Feature Selection," J. Mach. Learn. Res., vol. 3, pp. 1157–1182, 2003.

[13] N. D. Lane et al., "DeepX: A Software Accelerator for Low-Power Deep Learning Inference on Mobile Devices," in Proc. ACM UbiComp, 2015.

[14] F. Li et al., "Automated Diagnosis of ACL via a Weighted Multi-View Network," J. Biomed. Imaging, 2023.

[15] D. Wang and Y. Yan, "Improving InceptionV4 Model Using Fractional-Order Optimization for ACL Diagnosis," IEEE J. Biomed. Health Inform., vol. 28, no. 1, 2024.

[16] P. Schulc et al., "Video-Based ACL Injury Detection Using Pose Estimation and Recurrent Deep Learning," IEEE Access, vol. 12, pp. 56342–56354, 2024.