

Lightweight Machine Learning System Embedded in Smart Wearable Devices for Personalized Sport Guidance

Ming Mo, Wenfeng Shi, Jun Wang, Ye Yang, Xuyin Xu

How to cite: Mo M, Shi W, Wang J, Yang Y, Xu X. Lightweight Machine Learning System Embedded in Smart Wearable Devices for Personalized Sport Guidance. Textile & Leather Review. 2025; 8:1210-1230. <https://doi.org/10.31881/TLR.2025.1210>

How to link: <https://doi.org/10.31881/TLR.2025.1210>

Published: 31 December 2025



Lightweight Machine Learning System Embedded in Smart Wearable Devices for Personalized Sport Guidance

Ming MO¹, Wenfeng SHI^{2*}, Jun WANG¹, Ye YANG¹, Xuyin XU¹

¹Directional Training College, Changsha Aeronautical Vocational and Technical College, Changsha 410124, Hunan, China

²Department of Physical Education, Hunan Railway Professional Technology College, Zhuzhou 412001, Hunan, China

*15073378855@163.com

Article

<https://doi.org/10.31881/TLR.2025.1210>

Received 15 August 2025; Accepted 24 September 2025; Published 31 December 2025

ABSTRACT

The convergence of the textile industry and electronics presents new frontiers in functional apparel. This research details a system for personalized athletic guidance, rooted in advanced textile processing and materials science. The foundation is a smart garment fabricated from a blend of synthetic fibers using industrial knitting techniques. These methods are also applicable to natural fibers such as cotton or wool for enhanced comfort. The core innovation is a novel piezoresistive yarn, developed by dip-coating a base polyester yarn with poly (3,4-ethylenedioxythiophene):polystyrene sulfonate (PEDOT:PSS), a conductive polymer. This process transforms conventional fiber into a motion-sensing element. A lightweight machine learning model is embedded within the garment's microcontroller, enabling on-device analysis of data from the yarn sensors. This approach avoids energy-intensive cloud computing and represents a step toward sustainable development in electronic textiles by minimizing the system's overall energy footprint. The system accurately interprets complex movements, providing real-time feedback. This work demonstrates the integration of computational intelligence directly into fiber products, showcasing a scalable manufacturing pathway for a new generation of interactive textiles and moving beyond traditional applications of materials such as leather for wearable enclosures.

KEYWORDS

smart textiles, piezoresistive yarn, textile processing, sustainable development, fiber products

INTRODUCTION

The proliferation of wearable technology has fundamentally altered the landscape of personal health and fitness monitoring [1]. Devices ranging from wristbands to smartwatches have empowered individuals to track a myriad of physiological and biomechanical parameters, fostering greater awareness of their physical well-being [2,3]. Concurrently, the field of smart textiles has emerged as a transformative force, enabling the

seamless integration of electronic sensing capabilities directly into the fabric of everyday apparel [4,5]. This convergence of electronics and textiles has given rise to a new class of wearable systems that offer unobtrusive, continuous, and high-fidelity monitoring of human movement [6]. Unlike rigid, single-point sensors, smart textile-based wearables can conform to the body, capturing distributed pressure, strain, and motion data across multiple muscle groups and joints, thereby providing a more holistic view of a user's kinematics [7,8].

Despite these advancements in hardware, a critical gap remains in translating raw sensor data into actionable, personalized guidance [9]. Most current commercial systems rely on offloading data to a smartphone or the cloud for processing [10]. This cloud-centric approach introduces significant latency, making real-time feedback for correcting athletic form impractical [11,12]. A basketball player completing a jump shot or a runner adjusting their stride needs immediate, not delayed, corrective cues [13]. Furthermore, the continuous transmission of sensitive health and movement data raises substantial privacy concerns [14-16]. The dependency on an external processing unit also increases the overall power consumption of the system and renders the wearable non-functional in environments without reliable connectivity [17,18]. Therefore, the next frontier in intelligent sportswear lies in the development of on-device intelligence, where data processing and decision-making occur directly within the wearable itself [19,20].

This research addresses the critical need for embedded intelligence by developing a lightweight machine learning system specifically designed for resource-constrained wearable devices integrated into smart textiles. The primary challenge lies in the inherent conflict between the computational intensity of sophisticated machine learning models, such as deep neural networks—which excel at pattern recognition—and the limited processing power, memory, and energy budget of microcontrollers used in wearables. To overcome this hurdle, this paper proposes a holistic system that encompasses both hardware and software innovations. We introduce a smart garment with integrated textile sensors for comprehensive motion capture and, at its core, a lightweight convolutional neural network (Lw-CNN) architecture optimized for on-device inference. This Lw-CNN is engineered to provide not only accurate classification of different athletic movements but also to analyze the quality of those movements, enabling the system to deliver personalized, corrective feedback in real time. By embedding the entire intelligence pipeline—from data acquisition to actionable guidance—within the garment, our system enhances user experience, ensures data privacy, and provides the low-latency feedback essential for effective motor learning and injury prevention. This work contributes to the advancement of smart textiles and applied machine learning, demonstrating a viable pathway toward

creating truly autonomous and intelligent wearable coaching systems.

SYSTEM ARCHITECTURE AND METHODS

This section details the comprehensive architecture of the proposed system, from the design of the smart textile garment and its sensing mechanism to the development and implementation of the lightweight machine learning model. The methodology for data collection and the experimental protocol used to validate the system's performance are also described.

Smart Textile Garment Design and Data Acquisition

The foundation of the system is a form-fitting, long-sleeved compression shirt fabricated from a blend of nylon (85%) and spandex (15%). This material choice ensures high elasticity and a snug fit, which is critical for maintaining consistent contact between the sensors and the skin, and for accurately transducing body movements into electrical signals.

The sensing modality is based on a piezoresistive yarn developed in-house. The fabrication process involved a dip-coating technique. A multifilament polyester core yarn was continuously passed through a 1.5 wt% aqueous dispersion of poly (3,4-ethylenedioxythiophene) polystyrene sulfonate (PEDOT:PSS). After immersion, the yarn was drawn through a set of precision rollers to ensure a uniform coating thickness and then cured in a tube furnace at 120°C to evaporate the solvent. This process yielded a flexible, conductive yarn suitable for integration into textiles. The resulting yarn exhibits a change in electrical resistance in response to mechanical deformation such as stretching or pressure. These functional yarns were strategically integrated into the garment during the knitting process using an industrial flat-bed knitting machine (Stoll). The sensor patches were placed in locations chosen to capture the kinematics of complex compound movements based on established biomechanical principles. The goal was to monitor not only the primary mover muscles but also key stabilizers that are critical for maintaining correct form. The specific rationale is as follows:

Pectoralis Major and Anterior Deltoids: These are primary movers in pressing movements such as the push-up. Sensors here directly measure muscular effort and contraction patterns throughout the exercise.

Elbow Joint: A sensor across this joint provides a direct measurement of flexion and extension, which is crucial for assessing the range of motion in push-ups.

Latissimus Dorsi: This large back muscle is a key torso stabilizer. In push-ups, its activity indicates core

engagement, helping to detect errors such as hip sag. In squats and lunges, it is essential for maintaining an upright posture and preventing the common error of the torso leaning too far forward.

This distributed configuration was deemed sufficient for this study's scope, as it provides a holistic view of upper body action and postural stability, which are critical for correctly performing all three chosen exercises. Following fabrication, the piezoresistive yarns underwent rigorous electromechanical characterization to verify their sensitivity and stability prior to integration into the garment, as detailed in the subsequent section. The data acquisition hardware is a custom-designed, miniature electronic module (3 cm × 2.5 cm) securely housed in a small pocket on the shirt's hem. The module is built around a low-power microcontroller (ARM Cortex-M4), chosen for its balance of processing capability and energy efficiency. The microcontroller interfaces with a 16-channel analog-to-digital converter (ADC) to sample the resistance values from the textile sensor arrays. The sampling rate was set to 50 Hz, which is sufficient to capture the dynamics of typical athletic movements without generating excessive data. The module also incorporates a small vibration motor for providing haptic feedback and a Bluetooth Low Energy (BLE) module for initial setup and potential data logging, though all primary processing occurs on-device. The entire system is powered by a small, rechargeable 150 mAh Li-Po battery, providing over eight hours of continuous operation.

Piezoresistive Yarn Characterization

To validate the sensor's suitability for human motion capture, the electromechanical properties of the fabricated PEDOT:PSS yarn were systematically characterized using a universal testing machine (Instron 5943) coupled with a Keithley 2450 SourceMeter.

Sensitivity (Gauge Factor): The sensor's sensitivity was evaluated under quasi-static stretching conditions (0% to 50% strain). The relative resistance change ($\Delta R/R_0$) exhibited high linearity ($R^2 > 0.98$) within the working range of typical skin deformation (0–30% strain). The Gauge Factor (GF), calculated as $GF = (\Delta R/R_0) / \epsilon$, was measured to be 18.5. This high sensitivity confirms the yarn's capability to detect subtle muscle movements, far exceeding conventional metal foil strain gauges.

Hysteresis and Response Time: Low hysteresis is critical for differentiating between flexion and extension phases. During loading-unloading cycles at 30% strain, the yarn demonstrated a low hysteresis error of <6.5%. The sensor response time was measured at approximately 15 ms, which is well below the sampling interval of the system (20 ms at 50 Hz), ensuring no signal lag during rapid athletic movements.

Cyclic Durability and Washability: Long-term stability was tested over 1,000 stretching cycles at 30% strain.

The baseline resistance drift was $<4.2\%$, indicating excellent mechanical fatigue resistance. Furthermore, to address the requirement for sustainable, everyday use, the yarn underwent standardized washing tests (ISO 6330). After ten wash cycles, although the baseline resistance increased by 15% due to partial coating degradation, the sensitivity (GF) remained stable (variation $<8\%$). This stability validates the effectiveness of the user calibration routine described in Section Data Preprocessing and Feature Engineering, which compensates for baseline shifts while relying on the stable GF for motion detection.

Data Preprocessing and Feature Engineering

Raw data from piezoresistive sensors are often noisy due to motion artifacts, electromagnetic interference, and baseline drift. Therefore, a multi-stage preprocessing pipeline was implemented directly on the microcontroller. First, a digital second-order Butterworth low-pass filter with a cutoff frequency of 10 Hz was applied to each sensor channel to remove high-frequency noise. Subsequently, the signal from each sensor was normalized to a range of $[0, 1]$ based on the minimum and maximum resistance values observed during a brief calibration routine performed by the user at the start of each session. This normalization step compensates for inter-user variability and slight differences in sensor placement. Although the sensor exhibits excellent mechanical stability (Section Piezoresistive Yarn Characterization), minor baseline shifts due to sweat or thermal expansion are inevitable in wearable environments. The start-of-session calibration routine (detailed in Section Pilot Study Protocol and Personalized Guidance Validation) effectively neutralizes these shifts by establishing a fresh $[0, 1]$ normalization range for each user session. This ensures that the Lw-CNN receives consistent input distributions regardless of absolute resistance drift.

The continuous stream of normalized data was then segmented into fixed-size windows. A window size of 128 samples (corresponding to 2.56 seconds) with a 50% overlap (64 samples) was determined to be optimal for capturing a complete repetition of the target exercises without introducing excessive latency. As a preliminary step to establish a baseline and justify the use of a deep learning model, a set of time-domain features was extracted from each data window for use with traditional machine learning classifiers (e.g., Support Vector Machines). These extracted features included mean, variance, standard deviation, root mean square (RMS), and the 25th and 75th percentiles. While this feature engineering approach provides a compact data representation, our final Lw-CNN model, as detailed below, was designed to operate directly on the raw windowed sensor data to leverage end-to-end feature learning.

Lw-CNN Architecture

The core of the embedded intelligence is the Lw-CNN, a model specifically designed to balance high accuracy with the stringent computational constraints of the wearable's microcontroller. The architecture is inspired by modern efficient networks such as MobileNet but is further scaled down.

The Lw-CNN was designed to learn features directly from the sensor signals. As such, the input to the network is the raw, normalized, and windowed time-series data, formatted as a 2D matrix of size $N \times W$, where N is the number of sensor *channels* and W is the window size (128). The network structure is as follows:

Input Layer: 16 (channels) x 128 (samples)

Convolutional Block 1: A 1D convolutional layer with 8 filters of kernel size 5, followed by a ReLU activation function and batch normalization. This is followed by a max-pooling layer of size 2. This block is designed to learn low-level temporal features from the sensor signals.

Convolutional Block 2: A second 1D convolutional layer with 16 filters of kernel size 3, again followed by ReLU and batch normalization, and a max-pooling layer of size 2. This block captures more complex feature combinations.

Flatten Layer: The output from the convolutional blocks is flattened into a 1D vector.

Dense Layer: A fully connected layer with 32 neurons and a ReLU activation function.

Output Layer: A final dense layer with C neurons, where C is the number of classes ($C = 9$ in this implementation). A softmax activation function is used to produce the probability distribution over the classes.

To make this model lightweight, two primary optimization techniques were employed post-training:

Knowledge Distillation: To maximize the accuracy of our small model, we used knowledge distillation. First, a larger, more complex “teacher” model was trained on an offline dataset. The Lw-CNN (the “student” model) was then trained to mimic the teacher's outputs. This process transfers what is often called “dark knowledge.” An analogy is a master craftsman (teacher) training an apprentice (student). Instead of just showing the final, correct product (the hard label, e.g., “Correct Squat”), the master explains their nuanced thought process. The teacher model does something similar: its output might indicate that a specific movement is classified as a Correct Squat with 90% probability, but also 7% similar to a “Squat with Insufficient Depth” and 3% similar to a “Correct Lunge.” This information about class similarities—the dark knowledge—is incredibly valuable. It teaches the lightweight student model a more nuanced understanding of the data, allowing it to achieve higher accuracy than if it were trained only on the simple right or wrong labels. For this study, the teacher

model was a larger CNN architecture comprising four convolutional blocks, each with double the number of filters of the student model, and two dense layers with 128 and 64 neurons, respectively. Both the teacher and student models were trained using the Adam optimizer with an initial learning rate of 0.001 and a batch size of 64. The knowledge distillation process used a temperature parameter (τ) of 2 and a loss function that was a weighted average of the standard cross-entropy loss (weight $\alpha = 0.3$) and the distillation loss (weight $1 - \alpha = 0.7$). Training was conducted for 100 epochs with an early stopping protocol that monitored the validation loss with a patience of 10 epochs to prevent overfitting.

Post-Training Quantization: After training, the model's weights and activations, typically stored as 32-bit floating-point numbers, were quantized to 8-bit integers (INT8). This significantly reduces the model's memory footprint by a factor of four and allows for faster computation on microcontrollers that have specialized instructions for integer arithmetic, leading to lower latency and reduced power consumption. The final quantized model had a memory footprint of less than 100 KB, making it suitable for on-chip deployment.

Pilot Study Protocol and Personalized Guidance Validation

To evaluate the feasibility and system-level performance of the proposed Lw-CNN architecture, a pilot study was conducted with 10 healthy participants (6 male, 4 female, age 25 ± 4 years) with intermediate fitness experience. This specific demographic was selected to establish a baseline for sensor performance and algorithmic efficacy in a controlled environment before expanding to broader populations. Each participant was fitted with the smart shirt and instructed to perform three common exercises: squats, push-ups, and lunges. The protocol for each participant was as follows:

Phase 1: Sensor Normalization (30 s): To account for electrode shift and skin impedance changes, the user performs a brief 30-second routine involving prescribed range-of-motion gestures (e.g., T-pose, torso twists). This establishes the dynamic minimum and maximum resistance values ($R_{\min}$, $R_{\max}$) for the signal normalization step described in Section Data Preprocessing and Feature Engineering.

Correct Form Data Collection: Under the guidance of a certified personal trainer, each participant performed 20 repetitions of each exercise with correct form. To ensure consistency, a metronome was used to standardize the pace (e.g., a 4-second count for each repetition). All sessions were video-recorded. The trainer used a timestamped input to label the start and end of each "correct" repetition in real time. This ground truth was later verified by a second certified trainer using the video footage to ensure inter-rater reliability.

Incorrect Form Data Collection: To generate error data, a standardized elicited protocol was employed. The

trainer instructed participants to perform 10 repetitions of specific, archetypal errors (e.g., “valgus collapse”). This approach was strictly chosen to prioritize high-fidelity ground truth labeling. Unlike natural fatigue-induced errors, which are often subtle and progressive (leading to ambiguous labeling boundaries), these elicited archetypes provide distinct biomechanical signatures. This ensures that the initial training of the Lw-CNN is based on clear, uncontaminated kinematic targets. For squats, this included insufficient depth and excessive forward knee movement. For push-ups, it included hip sag and incomplete range of motion. For lunges, it involved the front knee extending past the toes. This data was crucial for training the model to recognize specific error states.

Model Training and Personalization: The system employs a transfer learning approach to ensure both generalizability and personalization.

(1) **Base Model Pre-training:** First, a general base Lw-CNN model was trained offline using the aggregated data from all participants in the study. This base model learns the general features of correct and incorrect forms across a diverse set of movements.

(2) **On-Device Fine-Tuning (Model Personalization):** Following sensor normalization (Phase 1), the user completes a guided Enrollment Set (approx. 2 min). This pre-trained base model is loaded onto the device, and the user performs five repetitions of each target exercise (squat, push-up, lunge) with correct form, standardized by a metronome (4 s/rep). Crucially, this phase utilizes only correct examples to ensure user safety. These specific data points are used to fine-tune the dense layers of the Lw-CNN (while keeping feature extractors frozen), anchoring the model to the user’s unique biomechanics without the need to train a full model from scratch.

(3) **Real-Time Testing:** Following the brief fine-tuning, the now-personalized model is used for real-time classification. To ensure a fair and unbiased evaluation of the personalized model’s real-time performance, the data collected for the subsequent testing phase was strictly separated from the initial data used for on-device fine-tuning. Participants took a brief rest period between the fine-tuning data collection and the final testing sequence to ensure temporal isolation, thereby preventing any overlap or data leakage between the training and testing sets for the personalization assessment. When an incorrect form was detected for more than two consecutive repetitions, the vibration motor provided a specific haptic pattern corresponding to the detected error.

Following the session, each participant completed a short questionnaire to provide qualitative feedback on their experience with the system, focusing on the intuitiveness and effectiveness of the haptic feedback. For

a comprehensive evaluation, we also benchmarked the Lw-CNN against two alternative models: a traditional Support Vector Machine (SVM) trained on the extracted time-domain features, and a scaled-down version of the well-established MobileNetV2 architecture.

Haptic Feedback Design

The effectiveness of the system hinges on the user's ability to intuitively understand the haptic feedback. We designed a set of distinct vibration patterns, each corresponding to a specific, common error. The patterns were designed to be simple and easily distinguishable during a workout:

Insufficient Squat Depth: A series of three short, quick pulses (buzz-buzz-buzz), designed to feel like a "tapping" cue to encourage the user to go lower.

Excessive Forward Knee Movement (Squat/Lunge): A single, long, continuous vibration (buzzzzzz), intended as a clear "warning" or "hold" signal to prompt the user to correct their knee alignment.

Hip Sag (Push-up): A distinct double pulse (buzz-buzz), acting as an alert to "tense up" or "re-engage" the core muscles.

Incomplete Range of Motion (Push-up): A slow, pulsing pattern (buzz...buzz...), indicating that the movement was not fully completed.

During the study, after the data collection phase, participants were familiarized with these patterns. While a formal usability study was outside the scope of this initial research, informal feedback indicated that users were able to correctly distinguish between the patterns and associate them with the intended correction after a brief training period.

RESULTS

The performance of the embedded lightweight machine learning system was evaluated based on its classification accuracy, computational efficiency, and the effectiveness of its real-time feedback mechanism. The results demonstrate the system's viability for on-device personalized sport guidance.

Classification Performance

The Lw-CNN model, after being trained and deployed on the microcontroller, was tasked with classifying the user's movements into nine distinct classes to enable granular feedback. These included: Squat (Correct), Squat (Depth Error), Squat (Knee Error), Push-up (Correct), Push-up (Hip Sag), Push-up (ROM Error), Lunge

(Correct), Lunge (Knee Error), and a “Null” class representing transitions or no activity. To ensure a robust and unbiased evaluation of the model’s generalizability and to prevent data leakage between training and validation sets, a Leave-One-Subject-Out (LOSO) cross-validation strategy was employed. For each fold of the validation, data from nine participants were used for training the model, while the data from the remaining one participant were used for testing. This process was repeated ten times, with each participant serving as the test subject once, and the final performance metrics were averaged across all folds. The performance of the model, averaged across all ten participants, is presented in the confusion matrix in Table 1.

Table 1. Detailed Confusion Matrix of the Lw-CNN Model’s Performance

Predicted True	Squat (Correct)	Squat (Depth Error)	Squat (Knee Error)	Push-up (C)	Push-up (Sag Error)	Push-up (ROM Error)	Lunge (C)	Lunge (Knee Error)	Null
Squat (Correct)	97.5%	1.0%	1.0%	0.0%	0.0%	0.0%	0.5%	0.0%	0.0%
Squat (Depth Error)	2.5%	95.0%	1.5%	0.0%	0.0%	0.0%	0.0%	1.0%	0.0%
Squat (Knee Error)	2.0%	1.8%	94.2%	0.0%	0.0%	0.0%	1.0%	1.0%	0.0%
Push-up (Correct)	0.0%	0.0%	0.0%	98.0%	0.5%	1.0%	0.0%	0.0%	0.5%
Push-up (Sag Error)	0.0%	0.0%	0.0%	1.5%	96.0%	1.0%	0.0%	0.0%	1.5%
Push-up (ROM Error)	0.0%	0.0%	0.0%	2.0%	1.5%	95.0%	0.0%	0.0%	1.5%
Lunge (Correct)	1.0%	0.0%	0.5%	0.0%	0.0%	0.0%	96.5%	2.0%	0.0%
Lunge (Knee Error)	0.0%	0.5%	1.0%	0.0%	0.0%	0.0%	3.0%	94.0%	1.5%
Null	0.5%	0.0%	0.0%	0.5%	0.0%	0.0%	0.0%	0.0%	99.0%

The diagonal values represent the percentage of correct classifications for each class. To strictly evaluate the model’s robustness across different users, we analyzed the inter-subject variability from the LOSO cross-validation. The overall accuracy of the system was calculated to be 96.2% ± 1.4% (mean ± SD), with a 95% confidence interval (CI) of [95.2%, 97.2%]. This low variance confirms that the model generalizes effectively across individuals, with no single subject exhibiting an accuracy below 93.5%. To address the potential for metric inflation due to the high-performing Null class (99.0%), we further calculated the Active Class Accuracy (excluding the Null state). The mean accuracy across the eight active exercise classes remains robust at 95.7%. This minimal discrepancy (<0.6%) confirms that the model’s high performance is driven by genuine discrimination of complex biomechanical patterns, rather than being skewed by the easier detection of inactivity. Consequently, the model demonstrated excellent performance in distinguishing among the three main types of exercises. The most common confusions occurred between the correct and incorrect forms of the same exercise, which is expected. For example, 2.3% of incorrect squats were misclassified as correct squats. This is a reasonable trade-off for a lightweight model, but it highlights an area for future improvement. The Null class was identified with 99.0% accuracy, indicating the system’s robustness in ignoring transitional movements and periods of rest.

The key performance metrics, including precision, recall, and F1-score, are summarized in Table 2. The high F1-scores across all classes (ranging from 0.941 to 0.975) indicate a strong balance between precision and recall, confirming that the system is not only accurate but also reliable in its predictions.

Table 2. Performance Metrics of the Lw-CNN Model (Averaged)

Class	Precision	Recall	F1-Score
Squat (Correct)	0.956	0.975	0.965
Squat (Depth Error)	0.955	0.950	0.952
Squat (Knee Error)	0.941	0.942	0.941
Push-up (Correct)	0.970	0.980	0.975
Push-up (Sag Error)	0.965	0.960	0.962
Push-up (ROM Error)	0.950	0.950	0.950

Lunge (Correct)	0.962	0.965	0.963
Lunge (Knee Error)	0.949	0.940	0.944
Overall	0.956	0.958	0.957

Impact of On-Device Personalization (Ablation Study)

To quantify the efficacy of the transfer learning framework, we compared the Base Model (trained only on offline data) against the Personalized Model (after on-device fine-tuning). As shown in Table 3, the Base Model achieved an average accuracy of 88.5% ± 3.1%. Following the rapid calibration phase, performance significantly improved to 96.2% ± 1.4%. A paired sample t-test confirmed this improvement is statistically significant (t(9) = 5.42, p < 0.001). Furthermore, the effect size (Cohen’s d = 1.71) indicates a substantial practical benefit, confirming that personalization is critical for adapting to individual biomechanics.

Table 3. Ablation Study: Impact of Personalization

Model Stage	Accuracy (Mean ± SD)	p-value	Effect Size (Cohen’s d)
Base Model	88.5% ± 3.1%	<0.001	1.71
Personalized (Lw-CNN)	96.2% ± 1.4%		

Computational Performance and Power Consumption

A primary goal of this research was to ensure the system was lightweight enough for on-device execution. The computational performance of the final quantized Lw-CNN model was benchmarked on the ARM Cortex-M4 microcontroller.

Memory Footprint: The final model size after quantization was 94.7 KB, which is well within the typical on-chip flash memory capacity of modern microcontrollers (often 256 KB or more).

Inference Time: The average time taken to perform a forward pass (i.e., classify one 2.56-second window of data) was 38 milliseconds. This low latency is critical for real-time applications, as it allows the system to

process incoming data much faster than it is generated (50 Hz sampling).

Power Consumption: The system exhibits a dynamic power profile dependent on user performance.

Monitoring State (Baseline): During standard operation (data acquisition, processing, and BLE maintenance), the average current consumption is 16.5 mA.

Feedback State (Active): When the haptic motor is activated, current consumption spikes to approximately 85 mA.

Consequently, battery life is user-dependent. For a Proficient User (monitoring mode with <1% feedback activation), the system achieves the reported ~9 hours of operation. However, under a Heavy Correction scenario (e.g., a novice user receiving feedback 10% of the time), the weighted average current rises to ~23.5 mA, reducing the operational life to approximately 6.4 hours. Even in this worst-case scenario, the 150 mAh battery remains sufficient for extended workout sessions.

For comparison, a standard, non-quantized CNN model with a similar architecture but using 32-bit floats had a memory footprint of over 380 KB and an inference time of 152 ms, making it unsuitable for real-time deployment on the target hardware. This comparison highlights the critical importance of the model optimization techniques employed.

Real-Time Feedback Latency

The effectiveness of a coaching system depends on the timeliness of its feedback. For this study, we measured the processing latency, defined as the time from the moment a complete 2.56-second data window is available to the activation of the haptic feedback motor.

This was empirically measured to be 215 ms on average. A breakdown of this latency reveals the distinct contributions of computational and physical factors:

Model Inference (38 ms): The core classification task on the Cortex-M4.

Data Preprocessing & Signal Acquisition (~25 ms): Includes sequential ADC readout of 16 channels, filtering, and normalization.

BLE Stack Overhead (~50 ms): Due to the high-priority interrupt handling of the Bluetooth Low Energy stack for maintaining connection intervals, process preemption introduces intermittent delays.

Haptic Actuator Rise Time (~100 ms): Critically, the measurement accounts for the mechanical inertia of the vibration motor. Unlike a digital LED which switches instantly, the eccentric rotating mass (ERM) motor requires approximately 100 ms to ramp up from a standstill to the perceptible vibration intensity required for

feedback.

Therefore, while the *computational* latency is minimal (~63 ms), the *system-level* latency is dominated by the physical constraints of the haptic hardware, which represents a hardware-level optimization target for future iterations (e.g., switching to Linear Resonant Actuators).

It is important to note that the total system latency comprises the inherent data acquisition time for the window overlap (1.28 s) plus the processing time (0.215 s), resulting in a total feedback loop of approximately 1.5 s. While this exceeds the strict sub-second threshold for intra-movement correction (modifying a motion while it is occurring), it is highly effective for “repetition-level feedback.” Given that a typical repetition in our protocol lasts approximately four seconds (Section Pilot Study Protocol and Personalized Guidance Validation), a 1.5-second delay ensures the user receives haptic cues immediately upon completing the eccentric phase or during the transition, acting as an effective active spotter for the subsequent repetition.

User Experience Feedback

In addition to the system’s quantitative performance, user experience feedback was collected through post-session questionnaires, which included both quantitative ratings and qualitative comments. The results indicated a highly positive reception of the system. On a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree), participants agreed that the haptic feedback was helpful in correcting their form (average score: 4.6) and that the guidance felt intuitive after a brief familiarization period (average score: 4.4).

Open-ended comments frequently highlighted the immediacy of the feedback. One participant noted, “The long buzz when my knee went too far forward was an instant reminder to pull back, much faster than a coach yelling.” Another commented, “I felt more confident that I was doing the squats correctly, especially with the go deeper tapping signal.” This qualitative feedback supports the conclusion that the system not only functions accurately but also provides a positive and effective user experience.

Benchmark Comparison with Alternative Models

To validate the effectiveness and efficiency of our proposed Lw-CNN, we compared its performance against both a traditional machine learning model (SVM) and another lightweight deep learning architecture (a scaled-down MobileNetV2). The results, summarized in Table 4, highlight the trade-offs between accuracy, model size, and computational cost.

Table 4. Performance Benchmark of Lw-CNN Against Alternative Models

Model	Overall Accuracy (%)	Model Size (KB)	Inference Time (ms)
SVM (with feature engineering)	91.5	N/A	8
Scaled-down MobileNetV2	97.1	312	115
Proposed Lw-CNN (Quantized)	96.2	94.7	38

The SVM, while extremely fast, achieved lower accuracy, struggling to differentiate the subtle variations between correct and incorrect forms. While the scaled-down MobileNetV2 achieved a marginally higher average accuracy (97.1%) compared to our Lw-CNN (96.2%), a Wilcoxon signed-rank test revealed no statistically significant difference between the two models ($p > 0.05$). This lack of statistical significance is a crucial finding: it validates that the proposed Lw-CNN achieves predictive performance statistically equivalent to deeper architectures, yet it does so with a 3.3× reduction in model size and a 3× faster inference speed. This confirms Lw-CNN as the statistically optimal choice for resource-constrained embedded constraints. These results validate that our custom Lw-CNN architecture provides the optimal balance of high accuracy and computational efficiency for this embedded application.

DISCUSSION

The results presented in this study demonstrate the successful development and validation of a lightweight machine learning system embedded in a smart textile garment for personalized sport guidance. Our findings have several important implications for the future of wearable technology, smart textiles, and personal fitness. The high classification accuracy of 96.2% achieved by the on-device Lw-CNN model is a significant outcome. It indicates that highly optimized, lightweight deep learning models can indeed rival the performance of more complex, cloud-based counterparts for specific tasks such as activity recognition. This challenges the prevailing paradigm that high-performance AI requires powerful, energy-intensive hardware, and it opens up new possibilities for embedding sophisticated intelligence at the very edge of the network—directly into our clothing. The ability to accurately distinguish not only among different exercises but also between correct and incorrect execution of the same exercise is a critical step toward creating truly smart coaching systems. The comparison between the quantized Lw-CNN and a standard, non-quantized version—as well as the

benchmark against the SVM and MobileNetV2—underscores the necessity of a custom, optimized architecture. While traditional models like SVM are fast but less accurate, and established deep models like MobileNetV2 are powerful but too resource-intensive, our Lw-CNN occupies a crucial middle ground. This aspect of our work provides a practical blueprint for other researchers and engineers looking to deploy machine learning models on resource-constrained embedded systems. It highlights that the design of the AI software is as crucial as the design of the wearable hardware. The low power consumption and resulting multi-hour battery life further confirm the feasibility of this approach for consumer-grade products that require both intelligence and practicality.

A core contribution of this research is the system's capacity for real-time, personalized feedback. The measured latency of 215 ms is a crucial performance metric. In motor learning, timeliness is critical. While some literature suggests sub-second feedback is ideal for continuous tracking, our system provides near real-time guidance. By delivering haptic feedback within 1.5 seconds—well within the duration of a typical exercise repetition—the system functions as an active spotter, notifying the user of errors just as they complete a movement segment. This allows for inter-repetition correction, enabling the user to adjust their form before the next repetition, creating a tight and effective feedback loop. This immediacy allows the user to mentally connect the haptic sensation with their proprioceptive feeling of the incorrect movement, facilitating faster and more effective correction. The personalization aspect, achieved through rapid on-device fine-tuning, is another vital element. By adapting a robust, pre-trained base model to an individual's specific data, the system accounts for unique biomechanical differences. This transfer learning approach provides tailored guidance that is more accurate than a generic model, while remaining practical and scalable for a consumer product, avoiding the need for lengthy, data-intensive training sessions for each new user. This accounts for individual biomechanical differences, ensuring that the guidance provided is tailored to the user's unique body and movement patterns, rather than a generic, one-size-fits-all template.

However, we must also acknowledge the significant limitations of this study, which frame the directions for future research. While the sensor characterization (Section Piezoresistive Yarn Characterization) confirmed the yarn's robustness and high sensitivity ($GF \approx 18.5$) for detecting athletic movements, the current study implies practical challenges for mass deployment.

Furthermore, it is acknowledged that the current validation relied on a homogeneous cohort of young, healthy individuals with intermediate fitness levels. This represents a best-case biomechanical scenario. Additionally, the use of elicited, archetypal error patterns limits the direct generalizability to the high-entropy

movement patterns typical of natural, fatigue-induced form degradation.

However, these constraints were accepted to prioritize internal validity and safety. Validating the system on archetypes establishes a necessary baseline without the injury risks associated with inducing failure-level fatigue. Moreover, the inherent limitation of small sample sizes and specific error types highlights the necessity of the proposed transfer learning approach. Instead of relying on a massive, universal dataset to cover every possible human variation (which is computationally expensive and impractical for embedded systems), our architecture relies on the on-device fine-tuning mechanism. This design allows the system to adapt the general features learned from the pilot group to the specific biomechanics of a new user—effectively mitigating the issue of demographic overfitting. This suggests that the system is architecturally ready to scale to diverse populations, provided the individual calibration step is performed. Furthermore, the personalization protocol was intentionally designed to train exclusively on correct form data. Requiring novice users to perform specific incorrect movements for calibration would introduce significant safety risks and confusion. By freezing the feature extraction layers that encode the archetypal error patterns (learned from the elicited dataset), the system effectively detects deviations from the user-calibrated correct baseline. This ensures reliable error detection without exposing the end-user to the risk of performing faulty repetitions during setup. Finally, while the system's *processing* latency is low, the window-based approach introduces an inherent data acquisition delay that must be addressed, and the system has not been benchmarked against commercial wearables such as smartwatches. Future work must therefore focus on robust sensor characterization, developing intuitive self-calibration for novice users, exploring more efficient “few-shot” personalization methods, and conducting larger-scale validation studies with a direct comparison to existing market solutions.

In conclusion, this research serves as a proof of concept for a new generation of intelligent wearables. By synergistically combining smart textile sensors for rich data capture with highly optimized, embedded machine learning for on-device intelligence, we have created a system that is more than just a passive tracker—it is an active, personalized coach that provides real-time guidance. This shift from data logging to actionable intelligence represents a significant leap forward, paving the way for wearable systems that not only inform users about their past activity but actively shape and improve their future performance in a safe and effective manner.

CONCLUSION

This paper presented a complete embedded system for personalized sport guidance, integrating a smart textile garment with a lightweight machine learning model. We have successfully demonstrated that by leveraging advanced sensor-integrated fabrics and applying rigorous model optimization techniques such as knowledge distillation and quantization, it is feasible to deploy an intelligent system on a low-power microcontroller without sacrificing performance. The developed system achieves high accuracy in classifying complex athletic movements and, more importantly, in identifying common form errors. The on-device Lw-CNN model provides the computational efficiency required for a minimal memory footprint and low-latency inference, enabling a long battery life suitable for practical use.

The key contribution of this work is the realization of a closed-loop feedback system entirely within a wearable garment. By eliminating the reliance on external processing units such as smartphones or the cloud, our system ensures user privacy, minimizes latency, and offers a seamless user experience. The ability to provide immediate haptic feedback transforms the smart garment from a passive data logger into an active coaching tool that can guide users toward better form, potentially enhancing workout efficacy and reducing the risk of injury. The findings of this research provide a robust framework and a significant step toward the commercialization of truly intelligent apparel, highlighting a promising future where our clothing actively collaborates with us to achieve our health and fitness goals. The methodologies and principles established herein offer valuable insights for the continued development of the next generation of smart textiles and embedded AI systems.

Availability of Data and Materials

The datasets used and/or analysed during the current study were available from the corresponding author on reasonable request.

Author Contributions

Ming Mo and Ye Yang designed the study; all authors conducted the study; Wenfeng Shi and Xuyin Xu collected and analyzed the data. Jun Wang and Wenfeng Shi participated in drafting the manuscript, and all authors contributed to critical revision of the manuscript for important intellectual content. All authors gave final approval of the version to be published. All authors participated fully in the work, took public responsibility for appropriate portions of the content, and agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or completeness of any part of the work were appropriately investigated and resolved.

Ethics Approval and Consent to Participate

This survey was conducted in compliance with Ethics Committee of Hunan Railway Professional Technology College. Participants were informed of the study's purpose and data usage prior to participation, and responses were collected anonymously. No personally identifiable information was stored.

Acknowledgments

Not applicable.

Funding

Supported by Hunan Provincial Natural Science Foundation of China. (Projectnumber: 2024JJ8032)

Project of Hunan social science achievement evaluation committee. (Projectnumber: XSP25YBC513)

Conflict of Interest

The authors declare no conflict of interest.

REFERENCES

- [1] Del-Valle-Soto C, Briseño RA, Valdivia LJ, Nolasco-Flores JA. Unveiling wearables: Exploring the global landscape of biometric applications and vital signs and behavioral impact. *BioData Mining*. 2024; 17(1):15. doi: 10.1186/s13040-024-00368-y

- [2] Del-Valle-Soto C, López-Pimentel JC, Vázquez-Castillo J, Nolasco-Flores JA, Velázquez R, Varela-Aldás J, et al. A comprehensive review of behavior change techniques in wearables and IoT: Implications for health and well-being. *Sensors*. 2024; 24(8):2429. doi: 10.3390/s24082429
- [3] Asgari Mehrabadi M. Holistic Health Monitoring and Personalized Intervention for Well-being Promotion: UC Irvine; 2022. ProQuest ID: AsgariMehrabadi_uci_0030D_17885. Merritt ID: ark:/13030/m5kb0j7v. Available from: <https://escholarship.org/uc/item/2dv4h97k>
- [4] Shi J, Liu S, Zhang L, Yang B, Shu L, Yang Y, et al. Smart textile-integrated microelectronic systems for wearable applications. *Advanced Materials*. 2020; 32(5):1901958. doi: 10.1002/adma.201901958
- [5] Abtew MA, Atalie D, Dejene BK, McBee-Black K. Intelligent and electronic textile materials for adaptive apparel: Innovations, functional design, and future directions. *Journal of Industrial Textiles*. 2025; 55:15280837251346789. doi: 10.1177/15280837251346789
- [6] Yin L, Sun X. Textile-based sensors for human motion sensing: Recent developments and future perspectives. *Nanocomposites*. 2025; 11(1):79-98. doi: 10.1080/20550324.2025.2477392
- [7] Golgouneh A, Dunne LE. A review in on-body compression using soft actuators and sensors: Applications, mechanisms, and challenges. *IEEE Reviews in Biomedical Engineering*. 2022; 17:166-179. doi: 10.1109/RBME.2022.3220505
- [8] Mondal R, Hasan MAM, Zhang R, Olin H, Yang Y. Nanogenerators-based self-powered sensors. *Advanced Materials Technologies*. 2022; 7(12):2200282. doi: 10.1002/admt.202200282
- [9] Diraco G, Rescio G, Caroppo A, Manni A, Leone A. Human action recognition in smart living services and applications: Context awareness, data availability, personalization, and privacy. *Sensors*. 2023; 23(13):6040. doi: 10.3390/s23136040
- [10] Akherfi K, Gerndt M, Harroud H. Mobile cloud computing for computation offloading: Issues and challenges. *Applied Computing and Informatics*. 2018; 14(1):1-16. doi: 10.1016/j.aci.2016.11.002
- [11] D'Agati L. Compute Continuum in Bioengineering: Enhanced Motion Capture and Real-Time Data Processing on Cloud-Edge Infrastructures. 2024. Available from: <https://hdl.handle.net/20.500.14242/188170>
- [12] Trigka M, Dritsas E. Edge and cloud computing in smart cities. *Future Internet*. 2025; 17(3):118. doi: 10.3390/fi17030118
- [13] Krause JV, Nelson C. Basketball Skills & Drills. Champaign, IL, USA: Human Kinetics; 2019.

- [14] Li M, Lou W, Ren K. Data security and privacy in wireless body area networks. *IEEE Wireless Communications*. 2010; 17(1):51-58. doi: 10.1109/MWC.2010.5416350
- [15] Lane J, Schur C. Balancing access to health data and privacy: A review of the issues and approaches for the future. *Health Services Research*. 2010; 45(5p2):1456-1467. doi: 10.1111/j.1475-6773.2010.01141.x
- [16] Mone V, Shakhlo F. Health data on the go: Navigating privacy concerns with wearable technologies. *Legal Information Management*. 2023; 23(3):179-188. doi: 10.1017/S1472669623000427
- [17] Gatouillat A. Towards Smart Services with Reusable and Adaptable Connected Objects: An Application to Wearable Non-Invasive Biomedical Sensors. Lyon, France: Université de Lyon; 2018. Available from: <https://theses.hal.science/tel-02090738/file/these.pdf>
- [18] Ben Ammar M, Ben Dhaou I, El Houssaini D, Sahnoun S, Fakhfakh A, Kanoun O. Requirements for energy-harvesting-driven edge devices using task-offloading approaches. *Electronics*. 2022; 11(3):383. doi: 10.3390/electronics11030383
- [19] Baburaj A, Jayadevan S, Aliyana AK, Sk NK, Stylios GK. AI-Driven TENGs for Self-Powered Smart Sensors and Intelligent Devices. *Advanced Science*. 2025; 12(20):2417414. doi: 10.1002/advs.202417414
- [20] Fu W, Yu B, Ji D, Zhou Z, Li X, Wang R, et al. Intelligent fibers and textiles for wearable biosensors. *Responsive Materials*. 2024; 2(4):e20240018. doi: 10.1002/rpm.20240018