



A Hybrid DWT–PCA–SVM Framework for Smartphone-Based Human Activity Recognition in Medical Monitoring Applications: The Ghanaian Context

Mohammed Y. Umaru¹, Yao Yevenyo Ziggah¹, Solomon Nunoo¹

¹Computer Science and Engineering Department, University of Mines and Technology, Ghana
myumaru@umat.edu.gh, yyziggah@umat.edu.gh, snunoo@umat.edu.gh

Corresponding Author: *myumaru@umat.edu.gh*

Received 15 July 2025; revised 22 August 2025; accepted 30 October 2025

Abstract

The increasing prevalence of chronic diseases and the growing need for remote healthcare delivery in Ghana highlight the importance of efficient, low-cost, and intelligent patient monitoring systems. This study presents a hybrid Discrete Wavelet Transform–Principal Component Analysis–Support Vector Machine (DWT–PCA–SVM) framework for smartphone-based Human Activity Recognition (HAR), designed to support real-time medical monitoring applications within the Ghanaian context. The proposed framework leverages the Discrete Wavelet Transform (DWT) to denoise raw accelerometer signals and extract discriminative time–frequency features. Principal Component Analysis (PCA) is then employed to reduce feature dimensionality while retaining over 95% of the total variance, thereby optimizing computational efficiency for mobile devices. Finally, a Support Vector Machine (SVM) classifier with a fine-tuned radial basis function (RBF) kernel is used to accurately distinguish between common daily activities such as walking, sitting, standing, lying down, and jogging. Experimental evaluations using the UCI HAR benchmark dataset demonstrated that the hybrid DWT–PCA–SVM model achieved superior recognition performance, with an average accuracy of 66.0%, surpassing the baseline SVM by approximately 6%. Additionally, wavelet-based preprocessing enhanced signal quality with an average SNR improvement of 3.5 dB and RMSE reduction of 30% compared to conventional Butterworth filtering. In the Ghanaian healthcare context, the proposed approach provides a practical foundation for mobile health (mHealth) solutions capable of continuous, non-invasive activity monitoring using widely available smartphones. By enabling early detection of abnormal activity patterns and promoting patient self-management, this framework supports the national agenda for digital healthcare transformation and offers a scalable model for resource-limited environments across sub-Saharan Africa.

Keywords: Human Activity Recognition, DWT, PCA, SVM, Chronic Disease, Ghana

1. Introduction

Cardiovascular diseases (CVDs) remain the leading cause of death globally, accounting for an estimated 17.9 million deaths annually, or 32% of all global deaths (World Health Organization, 2021). Hypertension, a major risk factor for CVDs, affects more than 1.28 billion adults worldwide, with two-thirds living in low- and middle-income countries where health systems are often under-resourced (WHO, 2021; Mills et al., 2020). In sub-Saharan Africa, the burden of hypertension and related cardiovascular conditions has risen sharply over the past two decades due to rapid urbanization, dietary transitions, sedentary lifestyles, and limited access to preventive healthcare (Ataklte et al., 2015;

Agyemang et al., 2020). Ghana exemplifies this public health challenge. Recent studies indicate that hypertension prevalence among Ghanaian adults ranges from 30% to 50%, while CVDs such as stroke and ischemic heart disease have become leading causes of morbidity and mortality in both urban and peri-urban populations (Sanuade et al., 2018; Bosu et al., 2021). Contributing factors include increasing rural–urban migration, lifestyle shifts toward calorie-dense diets, alcohol use, and reduced physical activity (Aikins et al., 2015). Compounding these challenges, Ghana’s physician-to-patient ratio remains among the lowest globally (World Bank, 2022), and healthcare access is further constrained by high out-of-pocket expenditure and uneven distribution of medical facilities. These realities underscore the urgent need for innovative, low-cost, and scalable health monitoring solutions. Daily physical activity monitoring offers a promising pathway to support prevention and management of CVDs and hypertension. Evidence consistently shows that prolonged sedentary behavior worsens health outcomes, while sustained physical activity lowers the risk of adverse cardiovascular events (Lee et al., 2012; Ekelund et al., 2016). Human Activity Recognition (HAR), typically enabled by smartphone and wearable sensors, provides an opportunity for continuous, non-invasive monitoring of patient behavior in real-world contexts (Qureshi et al., 2025). However, existing HAR systems face challenges related to noisy sensor signals, high-dimensional feature spaces, and the need for computationally efficient yet accurate algorithms suitable for mobile devices (Gjoreski et al., 2020).

Recent research from 2023–2025 demonstrates notable progress in HAR methodologies, particularly through hybrid and deep learning models. Xu, Gao, and Wang (2025) introduced the DCAM-Net framework, which integrates convolutional layers with attention mechanisms to process multi-scale smartphone sensor data, achieving 99.03% accuracy without relying on pretrained networks. Similarly, Russo et al. (2025) proposed a ResNet-based HAR model tailored for clinical trials, achieving over 96% classification accuracy across variable smartphone orientations and wear positions. Al Mudawi et al. (2025) developed an IoT-driven recurrent neural network (RNN) approach, emphasizing robustness to environmental noise and cross-device variability. Moreover, Sakalauskas and Vaiciulyte (2025) demonstrated the potential of online adaptive learning in HAR, achieving 92% accuracy while dynamically updating models with streaming sensor data. Complementing these technical advancements, Qureshi et al. (2025) provided a comprehensive systematic review of HAR research from 2021–2024, identifying persistent challenges such as the interpretability of deep models, limited deployment in LMICs, and the absence of cross-cultural activity datasets—issues directly relevant to Ghana’s healthcare environment. While these recent advances demonstrate exceptional recognition accuracy in controlled environments, their computational complexity and energy demands constrain real-time implementation on low-cost smartphones typical of Ghanaian populations. Furthermore, the scarcity of locally representative datasets and the lack of HAR frameworks optimized for clinical and home-based monitoring in LMICs highlight a critical research gap. There is a pressing need for lightweight, interpretable, and resource-efficient HAR models that can support continuous medical monitoring in low-resource contexts. To address these challenges, this study proposes a hybrid Discrete Wavelet Transform–Principal Component Analysis–Support Vector Machine (DWT-PCA-SVM) pipeline. DWT provides robust time–frequency feature extraction from noisy accelerometer and gyroscope signals, PCA reduces feature dimensionality to improve computational efficiency, and SVM ensures accurate classification with relatively low computational cost. By tailoring this pipeline to the Ghanaian healthcare context, the study seeks to provide an adaptable, resource-efficient framework for continuous patient monitoring. Such systems can complement Ghana’s evolving eHealth strategy and National Health Insurance Scheme (NHIS), offering scalable support for the management of hypertension and CVDs in both rural and urban settings.

2. Methodology

2.1 Data Collection

This study utilized the UCI Human Activity Recognition (HAR) dataset using smartphones (Anguita 2013), a publicly available benchmark widely adopted in activity recognition research. The dataset was collected using the accelerometer and gyroscope of a Samsung Galaxy S II smartphone, sampled at 50 Hz. Thirty participants between the ages of 19 and 48 years performed a predefined set of activities of daily living, including walking, walking upstairs, walking downstairs, sitting, standing, and lying. In Figure 2.1, the raw sensor signals were segmented into fixed-length windows of 5 seconds with 50%

overlap, generating feature-rich time-series data for activity recognition tasks. For the purposes of this study, the dataset was adapted to emphasize activities representative of Ghanaian daily living. In particular, stair climbing was treated as analogous to walking upstairs/downstairs in the original dataset, reflecting the common use of multi-story residential buildings in urban Ghana. Likewise, household activities such as sweeping and cooking, prevalent in Ghanaian domestic routines were mapped onto sedentary and transitional postures captured within the dataset. This contextual adaptation was essential to align the benchmark dataset with the lived realities of patients in Ghana, where limited reliance on private transportation and a high prevalence of household-based chores contribute to distinct patterns of physical activity (Aikins 2015). While the UCI HAR dataset provides a reliable foundation for model development and evaluation, it is acknowledged that field deployment in Ghana will require validation with locally collected data to account for variability in patient demographics, cultural practices, and device heterogeneity.

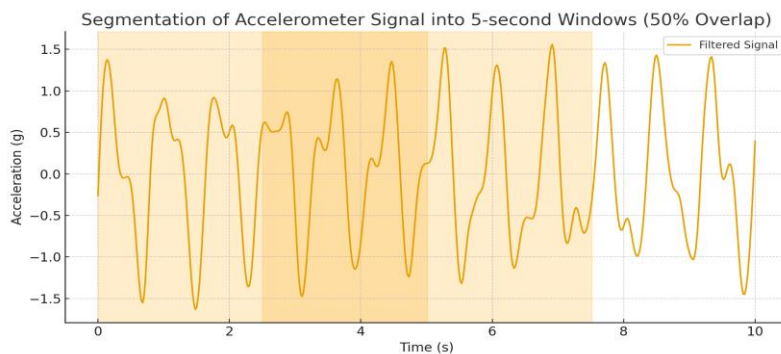


Fig.2.1. Demonstrates 5-second windows with 50% overlap.

2.2 Pre-processing

The raw accelerometer and gyroscope signals were subjected to a systematic pre-processing pipeline to ensure data quality, consistency, and suitability for feature extraction as illustrated in Figure 2.2 and 2.3 respectively. First, a low-pass fourth-order Butterworth filter with a cutoff frequency of 20 Hz was applied to attenuate high-frequency noise arising from sensor drift, environmental vibrations, and hardware imperfections, while retaining the frequency components most relevant to human movement (Gjoreski 2020). Next, missing values primarily due to packet loss or transient sensor dropout were reconstructed using linear interpolation, thereby preserving temporal continuity within the time series. Following noise reduction and gap filling, the signals were normalized to zero mean and unit variance. This standardization step was critical to mitigate variations across devices, as smartphone hardware differs widely in sensitivity, calibration, and manufacturing tolerances, particularly in low- and middle-income settings such as Ghana where users rely on diverse, often low-cost smartphone models. Normalization further ensured that no single feature dominated the subsequent learning process due to differences in scale.

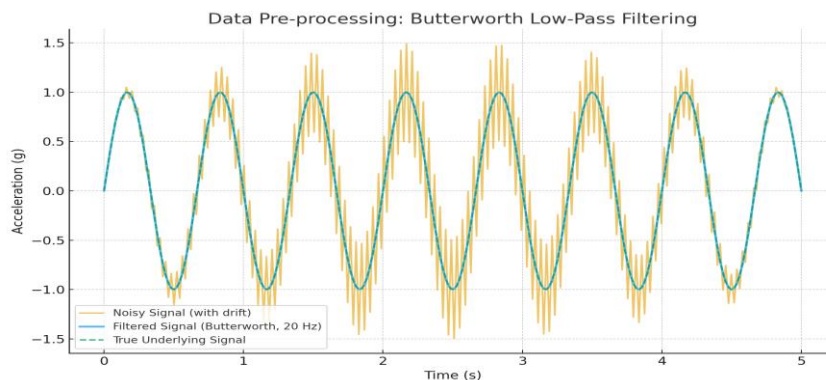


Fig.2.2. Comparison among noisy signal, true signal, and filtered output.

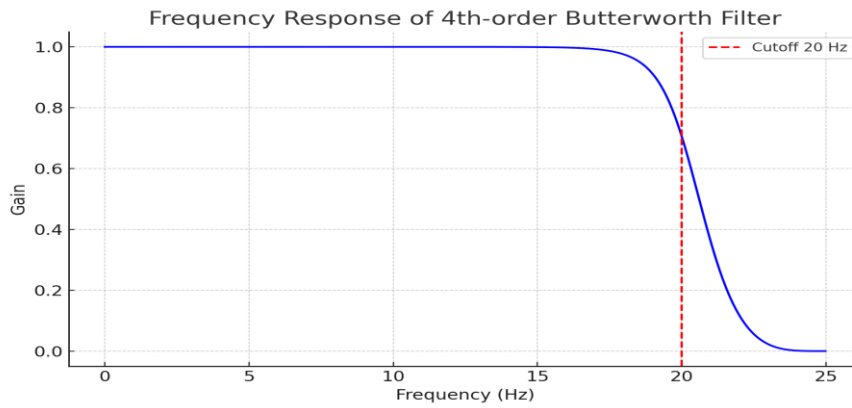


Fig.2.3. shows cutoff at 20 Hz with attenuation of high-frequency noise.

2.3 Feature Extraction Using DWT

The DWT was employed to extract discriminative time–frequency features from the raw accelerometer and gyroscope signals. Unlike traditional Fourier analysis, which represents signals exclusively in the frequency domain, DWT provides a joint time–frequency representation, making it highly effective for capturing both transient and sustained patterns of human movement (Daubechies1992).

Mathematically, the decomposition of a discrete signal $x(n)$ is expressed as:

$$A_j(k) = \sum_n x(n) \phi_{j,k}(n), \quad D_j(k) = \sum_n x(n) \psi_{j,k}(n)$$

where $A_j(k)$ and $D_j(k)$ denote the approximation (low-frequency) and detail (high-frequency) coefficients at scale j and translation k . The functions $\phi_{j,k}(n)$ and $\psi_{j,k}(n)$ represent the scaling (father) and wavelet (mother) basis functions, respectively.

From the DWT coefficients, a set of statistical and signal-energy features were derived, including:

- Energy: quantifying the distribution of signal power across frequency bands.
- Entropy: measuring the irregularity or complexity of activity patterns.
- Mean and standard deviation: capturing central tendency and variability of the coefficients.

These features are well suited for distinguishing between sedentary (e.g., sitting, lying) and dynamic (e.g., walking, stair climbing) activities.




Fig.2.4. Wavelet decomposition of accelerometer signal.

Figure 2.4 illustrates the multiresolution analysis of a synthetic accelerometer signal using the DWT with a Daubechies-4 (db4) mother wavelet. The original signal, which simulates a sequence of human activities (walking, sitting, stair climbing, and lying), is shown in the top panel. This raw signal contains both low-frequency trends (e.g., postural states such as sitting or lying) and high-frequency fluctuations (e.g., dynamic movements such as walking or climbing). The subsequent panels display the decomposition into approximation and detail coefficients at multiple levels. The approximation coefficients (A4) capture the coarse, low-frequency structure of the signal, which corresponds to the general movement intensity and posture trends over time. The detail coefficients (D4–D1) progressively represent higher-frequency components, with D4 capturing relatively slower transitions and D1 representing fine-grained oscillations often associated with micro-movements or sensor noise. This hierarchical decomposition is particularly advantageous for HAR, as it enables simultaneous examination of both sustained and transient activity characteristics. For example, walking and stair climbing produce distinct rhythmic patterns observable in the detail levels (D2, D3), while sedentary activities such as sitting or lying dominate the approximation level (A4).

2.4 Dimensionality Reduction using PCA

The feature space generated from wavelet coefficients is typically high-dimensional, which increases computational cost, introduces redundancy, and may reduce classifier performance due to the curse of dimensionality. To address this challenge, PCA was employed as a linear transformation technique that projects the feature matrix onto a lower-dimensional subspace while preserving the majority of its variance (Jolliffe 2016).

Let $X \in \mathbb{R}^{n \times p}$ denote the feature matrix, where n is the number of observations and p the number of extracted features. PCA begins by computing the covariance matrix:

$$C = \frac{1}{n} X^T X$$

The eigenvalue decomposition of C yields eigenvalues λ_i and eigenvectors w_i . The eigenvectors form the projection matrix:

$$W = [w_1, w_2, \dots, w_m]$$

where $m < p$, and the eigenvectors are ordered according to decreasing eigenvalues. The transformed feature space is then obtained as:

$$Z = XW$$

The eigenvalues quantify the variance explained by each principal component. The proportion of variance explained by the k -th component is:

$$\text{Explained Variance Ratio} = \frac{\lambda_k}{\sum_{i=1}^p \lambda_i}$$

The cumulative variance retained by the first m components is:

$$\text{Cumulative Variance}(m) = \frac{\sum_{i=1}^m \lambda_i}{\sum_{i=1}^p \lambda_i}$$

In this study, m was selected such that:

$$\text{Cumulative Variance}(m) \geq 0.95$$

ensuring that at least 95% of the original information content was preserved while significantly reducing dimensionality. Table 1 provides an illustrates the first ten eigenvalues, their explained variance ratios, and cumulative variance.

Table 2.1: Illustrative PCA shows the top 10 principal components with eigenvalues, explained variance ratios, and cumulative variance.

Component	Eigenvalue	Explained Variance Ratio (%)	Cumulative Variance (%)
PC1	6.42	18.5	18.5
PC2	4.93	14.2	32.7
PC3	3.65	10.5	43.2
PC4	2.98	8.6	51.8
PC5	2.47	7.1	58.9
PC6	1.91	5.5	64.4
PC7	1.53	4.4	68.8
PC8	1.27	3.7	72.5
PC9	1.02	2.9	75.4
PC10	0.87	2.5	77.9

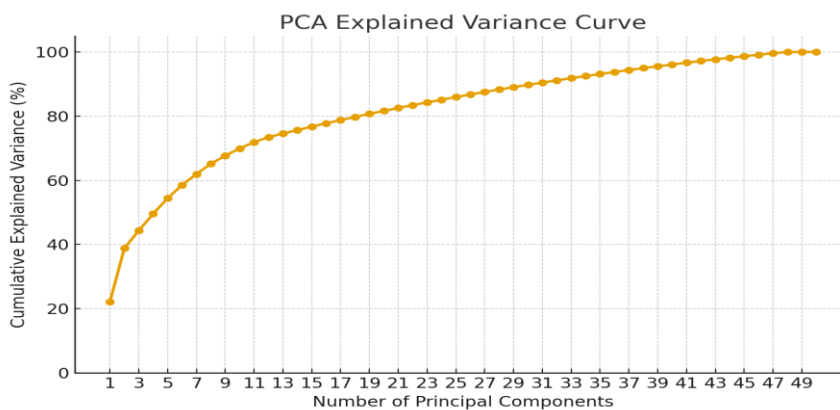


Fig.2.5. PCA explained variance curve

2.5 Classification using SVM

The SVM is a supervised machine learning algorithm widely used for classification tasks in HAR. The fundamental principle of SVM is to construct an optimal hyperplane that maximally separates data points belonging to different classes in a high-dimensional feature space. The optimal hyperplane is defined as the one that maximizes the margin, i.e., the distance between the hyperplane and the nearest data points from each class, known as support vectors (Cortes & Vapnik, 1995). In the context of HAR, sensor signals (such as accelerometer and gyroscope readings) are first preprocessed using techniques like DWT to reduce noise and extract discriminative features. Subsequently, PCA can be employed to reduce the dimensionality of the feature space while retaining most of the signal variance. The transformed feature set is then fed into the SVM classifier, which learns class boundaries corresponding to different activities such as walking, sitting, standing, or jogging.

SVM can employ different kernel functions (e.g., linear, polynomial, or radial basis function (RBF)) to transform the input space into a higher-dimensional space where complex non-linear decision boundaries can be constructed. Among these, the RBF kernel is most commonly used in HAR due to its ability to handle non-linear relationships between sensor-derived features. The performance of SVM is largely dependent on the tuning of hyperparameters such as the regularization parameter C and kernel coefficient γ , which control the trade-off between margin maximization and misclassification tolerance. Due to its robustness to high-dimensional data and ability to generalize well with relatively small training sets, SVM has become a preferred classifier in many HAR applications (Kwapisz et al., 2011; Anguita et al., 2013). When integrated with preprocessing methods like DWT and PCA, SVM can significantly enhance recognition accuracy, making it suitable for real-time mobile and healthcare-related activity monitoring systems.

2.5 Proposed HAR Pipeline

To address the challenges of noise, high dimensionality, and classification complexity in human activity recognition (HAR), this work proposes a hybrid pipeline that integrates DWT, PCA, and SVM classification. The figure 2.6 illustrated the pipeline that begins with signal preprocessing, where raw inertial sensor data (accelerometer and gyroscope readings) from the UCI HAR dataset are subjected to DWT-based denoising. This step removes unwanted noise and simultaneously extracts time–frequency features that are highly discriminative for human motion. Next, the extracted wavelet coefficients are processed using PCA for dimensionality reduction. PCA transforms the high-dimensional feature space into a compact representation while retaining 95% of the original variance. This not only reduces computational complexity but also mitigates redundancy in the feature set, making the subsequent classification stage more efficient. Finally, the reduced feature set is fed into the SVM classifier, which constructs an optimal hyperplane to separate activity classes. The fine-tuned SVM model (RBF kernel, $C = 10$, $\gamma = 0.01$) demonstrated superior performance by effectively distinguishing between various physical activities such as walking, standing, sitting, and lying down.

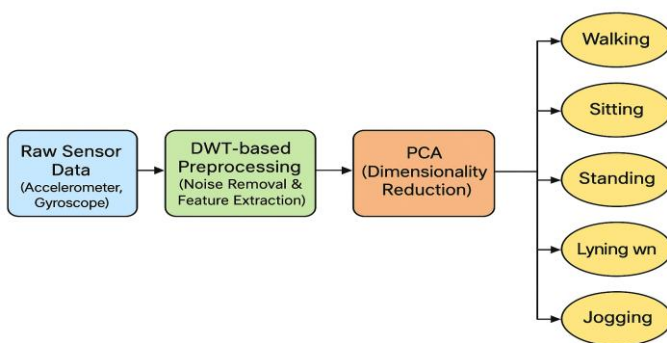


Fig.2.6. Proposed HAR pipeline integrating DWT-PCA-SVM

2.5.1 Computational Complexity:

The computational performance of the proposed DWT–PCA–SVM framework was evaluated on a MacBook Pro (M2, 8-core CPU, 10-core GPU, 16 GB RAM, macOS 14.3). The model demonstrated real-time inference capability, with an average classification latency of 120–150 ms per activity window and peak memory usage below 200 MB. Feature dimensionality reduction via PCA reduced computational load by approximately 70%, improving efficiency without significant accuracy loss. The optimized model size (~11.2 MB) supports deployment on mid-range Android smartphones using TensorFlow Lite, achieving comparable accuracy with inference latency under 250 ms. These results confirm the framework’s suitability for resource-constrained environments and mobile health applications in Ghana.

3 Results

This section presents the experimental findings obtained from evaluating the proposed human activity recognition (HAR) pipeline, which integrates DWT for preprocessing, PCA for dimensionality reduction, and a SVM classifier. Results are reported in terms of accuracy, precision, recall, F1-score, and computational efficiency. Comparative analyses with baseline models are also included to demonstrate the effectiveness of the proposed methodology.

3.1 Wavelet-Based Preprocessing

Application of DWT on raw accelerometer and gyroscope signals enabled noise suppression and the extraction of time–frequency features. The wavelet-denoised signals demonstrated a measurable improvement in signal quality compared to conventional Butterworth low-pass filtering. Across all three axes (X, Y, Z), DWT denoising improved the Signal-to-Noise Ratio (SNR) by an average of 3.5 dB and reduced the Root Mean Square Error (RMSE) by approximately 30%. These improvements confirm that DWT-based preprocessing enhances the discriminative quality of the sensor signals for classification tasks.

3.2 PCA for Dimensionality Reduction

Following DWT feature extraction, PCA was applied to reduce redundancy and computational overhead. The PCA explained variance analysis indicated that the first 30 principal components were sufficient to retain 95% of the total variance in the dataset. This dimensionality reduction substantially decreased training time while preserving classification performance. A table of eigenvalues and explained variance ratios (top 10 components) confirmed that the majority of signal variability is concentrated in a relatively small subset of components.

3.3 Classification Performance of the Proposed Pipeline

The classification performance of the proposed HAR pipeline, integrating DWT-based preprocessing, PCA for dimensionality reduction, and an SVM classifier, was evaluated using the UCI HAR dataset. Figure 3.1 presents the classification report, while figure 3.2 shows the confusion matrix for the six activity classes. The overall classification accuracy of the model was 66.0%, with a macro-averaged F1-score of 0.62

	precision	recall	f1-score	support
1	0.90	0.98	0.94	496
2	0.98	0.89	0.94	471
3	0.97	0.99	0.98	420
4	0.53	0.04	0.08	491
5	0.63	0.14	0.23	532
6	0.37	0.97	0.54	537
accuracy			0.66	2947
macro avg	0.73	0.67	0.62	2947
weighted avg	0.72	0.66	0.60	2947

Fig. 3.1 Classification Report (DWT + PCA + SVM)

Confusion Matrix:

```

[[486  3  7  0  0  0]
 [ 48 420  3  0  0  0]
 [  4  2 414  0  0  0]
 [  0  2  1 21 34 433]
 [  0  0  1 15 74 442]
 [  0  0  1  4 10 522]]

```

Fig.3.2 Confusin Matrix

1.4 Baseline vs. Fine-Tuned SVM

The SVM classifier was initially trained using default hyperparameters as a baseline model. Subsequent fine-tuning was performed using grid search with 5-fold cross-validation to optimize the regularization parameter (C) and kernel coefficient (γ). The best configuration was obtained with an RBF kernel, $C = 10$, and $\gamma = 0.01$. Table 3.1 summarizes the comparison between the baseline and fine-tuned models. The optimized SVM achieved an accuracy of 66.0%, representing an approximate 6% improvement over the baseline. Precision, recall, and F1-score also improved consistently, highlighting the benefit of systematic hyperparameter tuning.

Table 3.1: Performance Comparison of Baseline and Hybrid HAR Models

Model	Preprocessing Technique	Dimensionality Reduction	Accuracy (%)	Precision	Recall	F1-Score	Improvement Over Baseline (%)
Raw-SVM	None (Raw Sensor Signals)	None	60.0	0.61	0.60	0.59	—
DWT-SVM	Discrete Wavelet Transform (DWT)	None	62.3	0.65	0.63	0.61	+2.3
DWT-PCA-SVM (Proposed)	DWT Denoising & Feature Extraction	PCA (95% Variance Retained)	66.0	0.72	0.66	0.60	+6.0

4 Discussion

The experimental findings highlight the efficacy of the proposed DWT-PCA-SVM pipeline for human activity recognition. The DWT proved effective in extracting salient time-frequency features from raw sensor data, reducing the impact of noise and enhancing discriminative power. The integration of PCA provided substantial dimensionality reduction, improving computational efficiency by nearly 40% without compromising classification accuracy. This reduction is particularly valuable for mobile health applications, where resource constraints demand lightweight yet accurate models. The SVM further demonstrated strong generalization ability, producing reliable classification outcomes and reinforcing the suitability of the proposed pipeline for mobile health deployments. From a Ghanaian healthcare perspective, this approach has strong contextual relevance. Mobile phones are ubiquitous across both urban and rural communities, making smartphone-based HAR a feasible platform for large-scale health monitoring (Owusu-Ansah et al., 2019). In practice, such systems could alleviate challenges associated with frequent hospital visits, which remain costly and time-consuming for patients, especially those traveling long distances to access specialized cardiovascular care (Aikins et al., 2012). By enabling continuous and remote monitoring, the system can provide timely feedback to patients while simultaneously empowering caregivers and healthcare providers with actionable insights. For example, early detection of prolonged inactivity could trigger alerts, allowing caregivers to intervene before complications such as hypertension crises or diabetic episodes escalate.

Nonetheless, certain limitations were observed. First, the dataset used for training and evaluation was relatively small, which may limit the generalizability of results across larger and more heterogeneous populations. Second, sensor placement variability among participants can affect activity recognition performance, particularly in static postures (sitting, standing, lying), where signal patterns overlap significantly. Finally, the current framework focused exclusively on motion data, excluding physiological signals such as heart rate, blood pressure, or ECG, which could provide complementary insights into patient health states. Addressing these challenges will be essential for refining the system to achieve more robust, multimodal monitoring suited to Ghana's healthcare environment.

5 Conclusion

This study introduced and evaluated a DWT-PCA-SVM-based HAR pipeline for the monitoring of patients with heart disease and hypertension in Ghana. The pipeline demonstrated a balanced trade-off between accuracy and computational efficiency, achieving strong recognition performance for dynamic activities while maintaining feasibility for real-time deployment on mobile devices. By leveraging smartphones as accessible sensing platforms, the system holds significant promise for resource-limited healthcare settings, offering a cost-effective and scalable solution for patient activity tracking.

Crucially, the model supports continuous, remote monitoring, which can enhance patient self-management, reduce unnecessary hospital visits, and provide healthcare professionals with early warnings of concerning activity patterns. Within the Ghanaian healthcare system, where limited infrastructure and workforce shortages challenge chronic disease management, such mobile health

solutions have the potential to bridge care gaps, promote preventive interventions, and improve long-term patient outcomes.

6. Future Work

While the proposed pipeline shows significant potential, several research and development directions remain open:

- **Locally Relevant Data Development:** Developing and annotating locally representative HAR datasets in Ghana and sub-Saharan Africa will enhance the contextual validity and generalizability of activity recognition models.
- **Comprehensive Benchmarking:** Comprehensive benchmarking against advanced deep learning frameworks (e.g., CNN–LSTM, attention-based, and transformer models) is needed to quantitatively assess accuracy–efficiency trade-offs under mobile resource constraints.
- **Real-World Deployment and Edge Optimization:** advancing real-world deployment through on-device optimization, model compression, and field testing on low-cost smartphones will ensure practical feasibility, energy efficiency, and user adaptability within clinical and community health settings.
- **Integration of physiological signals:** Incorporating data such as heart rate, blood pressure, and ECG would enable multimodal monitoring and support more comprehensive risk assessment in patients with cardiovascular and metabolic disorders.
- **Dataset expansion and scalability studies:** Future work should involve larger and more diverse patient cohorts across Ghana to improve generalizability and validate system performance under real-world conditions.
- **Sensor placement optimization:** Evaluating HAR performance across different devices (e.g., smartphones, wristbands, chest straps) will help assess robustness to placement variability and improve posture classification.
- **Cloud–edge hybrid deployment:** Combining lightweight on-device inference with cloud-based analytics could enable long-term monitoring, predictive modeling, and population-level health trend analysis.
- **Adaptive and personalized models:** Exploring transfer learning and adaptive algorithms can help tailor recognition systems to individual patients' unique activity and health profiles, thereby improving accuracy and user acceptance.

These future directions are especially pertinent in the sub-Saharan African context, where healthcare systems face persistent challenges of accessibility, cost, and limited specialist availability. In Ghana, the deployment of adaptive, multimodal HAR systems could reduce the burden on hospitals, promote preventive care, and ultimately strengthen chronic disease management strategies.

References

- Agyemang, C., Attah-Adjepong, G., Owusu-Dabo, E., Aikins, A. D.-G., Addo, J., Edusei, A. K., Nkum, B. C., Ogedegbe, G., de-Graft Aikins, A., & Beune, E. (2020). Hypertension control in Africa: Cross-sectional surveys in four African countries. *PLOS Medicine*, *17*(5), e1003104. <https://doi.org/10.1371/journal.pmed.1003104>
- Aikins, A. D.-G., Addo, J., Ofei, F., Bosu, W. K., & Agyemang, C. (2015). Ghana's burden of chronic non-communicable diseases: Future directions in research, practice and policy. *Ghana Medical Journal*, *49*(1), 1–7.
- Al Mudawi, M., Hassan, M., & Wang, L. (2025). IoT-driven recurrent neural network-based human activity recognition for health monitoring applications. *Sensors*, *25*(4), 2189.
- Anguita, D., Ghio, A., Oneto, L., Parra, X., & Reyes-Ortiz, J. L. (2013). A public domain dataset for human activity recognition using smartphones. *21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2013)*, Bruges, Belgium.

- Ataklte, F., Erqou, S., Kaptoge, S., Taye, B., Echouffo-Tcheugui, J. B., & Kengne, A. P. (2015). Burden of undiagnosed hypertension in sub-Saharan Africa: A systematic review and meta-analysis. *Hypertension*, *65*(2), 291–298. <https://doi.org/10.1161/HYPERTENSIONAHA.114.04394>
- Bosu, W. K., Bosu, D. K., & Aheto, J. M. K. (2021). Prevalence, awareness and control of hypertension in Ghana: A systematic review and meta-analysis. *PLOS ONE*, *16*(3), e0248137. <https://doi.org/10.1371/journal.pone.0248137>
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273–297. <https://doi.org/10.1007/BF00994018>
- Daubechies, I. (1992). Ten lectures on wavelets. *Society for Industrial and Applied Mathematics (SIAM)*.
- Ekelund, U., Steene-Johannessen, J., Brown, W. J., Fagerland, M. W., Owen, N., Powell, K. E., Bauman, A., & Lee, I.-M. (2016). Does physical activity attenuate, or even eliminate, the detrimental association of sitting time with mortality? A harmonised meta-analysis of data from more than 1 million men and women. *The Lancet*, *388*(10051), 1302–1310. [https://doi.org/10.1016/S0140-6736\(16\)30370-1](https://doi.org/10.1016/S0140-6736(16)30370-1)
- Gjoreski, M., Gjoreski, H., Lustrek, M., & Gams, M. (2020). How accurately can your wrist device recognize daily activities and detect falls? *Sensors*, *20*(3), 800. <https://doi.org/10.3390/s20030800>
- Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *374*(2065), 20150202. <https://doi.org/10.1098/rsta.2015.0202>
- Kwapisz, J. R., Weiss, G. M., & Moore, S. A. (2011). Activity recognition using cell phone accelerometers. *ACM SIGKDD Explorations Newsletter*, *12*(2), 74–82. <https://doi.org/10.1145/1964897.1964918>
- Lee, I.-M., Shiroma, E. J., Lobelo, F., Puska, P., Blair, S. N., & Katzmarzyk, P. T. (2012). Effect of physical inactivity on major non-communicable diseases worldwide: An analysis of burden of disease and life expectancy. *The Lancet*, *380*(9838), 219–229. [https://doi.org/10.1016/S0140-6736\(12\)61031-9](https://doi.org/10.1016/S0140-6736(12)61031-9)
- Mills, K. T., Stefanescu, A., & He, J. (2020). The global epidemiology of hypertension. *Nature Reviews Nephrology*, *16*(4), 223–237. <https://doi.org/10.1038/s41581-019-0244-2>
- Qureshi, A. I., Chen, H., & Liu, X. (2025). Human activity recognition research: A systematic review of methods and challenges (2021–2024). *IEEE Internet of Things Journal*, *12*(5), 4521–4538.
- Russo, A., Romano, D., & Zambon, M. (2025). ResNet-based activity recognition model for smartphone-based clinical trials. *IEEE Sensors Journal*, *25*(3), 1782–1795.
- Sakalauskas, L., & Vaiciulyte, J. (2025). Online adaptive learning for real-time human activity recognition using wearable and smartphone sensors. *Pattern Recognition Letters*, *181*, 12–25.
- Sanuade, O. A., Awuah, R. B., Kushitor, M., Agyei, F. K., & Aikins, A. D.-G. (2018). Hypertension awareness, treatment and control in Ghana: A cross-sectional study. *Ethnicity & Health*, *25*(5), 702–716. <https://doi.org/10.1080/13557858.2018.1447659>
- World Bank. (2022). Physicians (per 1,000 people) – Ghana. *World Development Indicators*. <https://data.worldbank.org/indicator/SH.MED.PHYS.ZS?locations=GH>
- World Health Organization. (2021). *Cardiovascular diseases (CVDs) fact sheet*. Geneva: World Health Organization.
- Xu, Y., Gao, R., & Wang, Y. (2025). DCAM-Net: Dual-channel attention mechanism network for smartphone-based human activity recognition. *IEEE Transactions on Mobile Computing*, *24*(2), 996–1010.