



Wearable and Sensor-Based Machine Learning for Noncommunicable Disease Monitoring: A Comprehensive Synthesis

Hina Kato, Thibault Berger, Loic Legrand

Faculty of Medicine, Dentistry and Health Sciences, University of Melbourne, Melbourne, Australia

Abstract

The global prevalence of noncommunicable diseases represents one of the most significant public health challenges of the twenty-first century, demanding a transition from reactive, episodic healthcare to proactive, continuous monitoring paradigms. This paper provides a comprehensive synthesis of the integration of wearable sensor technologies with machine learning architectures for the continuous monitoring and management of noncommunicable diseases, primarily focusing on cardiovascular conditions, metabolic disorders such as diabetes, and chronic respiratory diseases. By capturing high-fidelity, real-time physiological data through non-invasive modalities, wearable sensors offer unprecedented longitudinal insights into patient health trajectories. Concurrently, advancements in machine learning algorithms, particularly deep learning and edge-based inference models, have enabled the extraction of actionable clinical features from massive, noisy time-series datasets. This research elucidates the end-to-end analytical framework, encompassing data acquisition protocols, signal preprocessing methodologies, feature extraction paradigms, and model selection criteria. Furthermore, we critically evaluate the performance, computational efficiency, and clinical implementation challenges associated with these integrated systems. The findings suggest that while sensor-based machine learning demonstrates robust diagnostic and predictive capabilities, widespread clinical translation requires overcoming substantial barriers related to data privacy, algorithmic bias, and interoperability. Ultimately, this paper charts a course for future research, emphasizing the necessity of federated learning, personalized medicine approaches, and regulatory-compliant device development to fully realize the potential of digital biomarkers in noncommunicable disease management.

Keywords: *Wearable Sensors, Machine Learning, Noncommunicable Diseases, Digital Biomarkers*

INTRODUCTION

1.1 The Burden of Noncommunicable Diseases

Noncommunicable diseases represent a staggering global health crisis, accounting for the vast majority of morbidity and mortality worldwide. Conditions such as cardiovascular diseases, chronic respiratory ailments, diabetes, and various forms of cancer collectively burden healthcare systems, economies, and societies at an unprecedented scale. Historically, the management of these chronic conditions has relied heavily on episodic clinical encounters, where patients are evaluated at distinct, often widely spaced intervals [1]. This traditional paradigm is fundamentally reactive, typically addressing symptom exacerbation rather than focusing on early detection and continuous prevention. The episodic nature of clinical visits severely limits the temporal resolution of physiological data available to healthcare providers, creating significant blind spots in understanding the dynamic nature of noncommunicable diseases. Patients may experience critical fluctuations in their physiological parameters between scheduled appointments, which often go unnoticed until a severe acute event necessitates emergency

intervention. Consequently, there is an urgent and universally recognized need to transition toward a more continuous, proactive, and personalized approach to disease management [2]. This transition is essential not only for improving patient outcomes and quality of life but also for ensuring the long-term economic sustainability of global healthcare systems, which are increasingly strained by aging populations and the rising incidence of chronic conditions [3].

1.2 The Emergence of Wearable Sensors

In response to the limitations of episodic care, the proliferation of wearable sensor technologies has emerged as a transformative force in biomedical engineering and health informatics. Wearable devices, ranging from commercially available smartwatches and fitness bands to clinical-grade biosensor patches and smart textiles, have fundamentally altered the landscape of physiological monitoring. These devices are equipped with a diverse array of miniature sensors, including photoplethysmography optical sensors, electrocardiogram electrodes, multi-axis accelerometers, gyroscopes, temperature sensors, and even galvanic skin response monitors [4]. By remaining in constant contact or close proximity to the user, wearable sensors facilitate the unobtrusive, continuous collection of dense longitudinal data streams in free-living environments. This continuous data acquisition is crucial for capturing the subtle, transient physiological anomalies that often serve as early indicators of disease onset or exacerbation. For instance, continuous heart rate variability monitoring can provide early warnings of cardiac stress, while continuous movement profiling can indicate the progression of neurological or musculoskeletal complications [5]. The transition from hospital-bound, bulky diagnostic equipment to miniaturized, ubiquitous wearable platforms democratizes access to high-fidelity health monitoring, enabling individuals to assume a more active role in their personal health management while providing clinicians with an exceptionally rich tapestry of objective, real-world patient data.

1.3 Machine Learning as a Catalyst

While the hardware advancements in wearable sensors have solved the problem of continuous data acquisition, they have concomitantly generated a new challenge: the overwhelming volume, velocity, and complexity of the generated data. Human analysts and traditional statistical methods are ill-equipped to process the terabytes of noisy, multidimensional, time-series data produced by even a small cohort of wearable users. This is where machine learning assumes its role as the critical catalyst in the digital health ecosystem [6]. Machine learning algorithms possess the inherent capacity to ingest massive datasets, identify latent patterns, and model complex, non-linear relationships between disparate physiological variables without requiring explicit, rule-based programming. By applying supervised, unsupervised, and semi-supervised learning techniques to wearable sensor data, researchers can extract meaningful digital biomarkers, filter out motion artifacts, and generate predictive insights regarding a patient's health status [7]. The synergy between wearable sensors and machine learning effectively bridges the gap between raw, uninterpretable signal data and actionable clinical intelligence, enabling real-time risk stratification, automated anomaly detection, and the delivery of timely, context-aware therapeutic interventions for patients suffering from noncommunicable diseases.

2. Literature Review

2.1 Cardiovascular Disease Monitoring

Cardiovascular diseases, encompassing a range of conditions including coronary artery disease, heart failure, and various arrhythmias, remain the leading cause of death globally. The application of wearable sensors coupled with machine learning has seen perhaps its most rapid and profound advancement within this domain. Early research primarily focused on utilizing basic heart rate metrics derived from simple optical sensors. However, recent literature demonstrates a paradigm shift toward sophisticated electrocardiogram and advanced photoplethysmography analysis. Researchers have successfully deployed deep convolutional neural networks to classify distinct types of arrhythmias, such as atrial fibrillation, directly from single-lead wearable electrocardiogram patches with diagnostic accuracy rivaling that of board-certified cardiologists [8]. Atrial fibrillation, a major risk factor for ischemic stroke, is notoriously difficult to diagnose in its paroxysmal form due to its transient nature. Wearables solve this by providing round-the-clock surveillance. Furthermore, studies have explored

the derivation of blood pressure using cuffless, continuous approaches. By analyzing pulse transit time and morphological features of the photoplethysmography waveform using gradient boosting machines and support vector regression models, investigators have achieved promising results in non-invasive hemodynamic monitoring [9]. These innovations indicate a significant departure from traditional sphygmomanometry, offering the potential for nocturnal blood pressure dipping analysis, a critical prognostic indicator for cardiovascular mortality [10]. The literature consistently highlights that the fusion of multi-sensor data, such as combining accelerometer data to contextualize physical exertion alongside heart rate data, significantly improves the predictive power of machine learning models in cardiovascular contexts.

2.2 Metabolic Syndrome and Diabetes Tracking

The global epidemic of diabetes mellitus and broader metabolic syndrome requires meticulous, daily management to prevent severe microvascular and macrovascular complications. The cornerstone of diabetes management is glycemic control, which has been revolutionized by the advent of continuous glucose monitors. While these subcutaneous sensors are technically minimally invasive rather than strictly non-invasive wearables, they form the foundation of modern diabetes informatics. Machine learning research in this area is heavily focused on predictive modeling to forecast future blood glucose concentrations, thereby preventing dangerous episodes of hypoglycemia and hyperglycemia [11]. Long short-term memory networks and other recurrent neural network architectures are frequently cited in the literature for their exceptional ability to handle the temporal dependencies inherent in glucose time-series data. Furthermore, an emerging and highly active area of research involves the use of truly non-invasive wearable sensors to infer metabolic states. Studies have investigated the correlations between interstitial fluid glucose levels and variables such as skin temperature, galvanic skin response, and heart rate variability collected via smartwatches [12]. By feeding these non-invasive physiological proxies into random forest classifiers and deep neural networks, researchers attempt to create completely non-invasive glucose estimation systems. While the accuracy of these surrogate models does not yet meet clinical standards for insulin dosing, they represent a highly promising frontier for general population screening and early detection of pre-diabetic states [13].

2.3 Chronic Respiratory Condition Management

Chronic respiratory diseases, predominantly asthma and chronic obstructive pulmonary disease, are characterized by recurrent episodes of breathlessness and wheezing. The management of these conditions relies heavily on identifying exacerbation triggers and recognizing the early signs of pulmonary function decline. Wearable technology literature in this sector frequently discusses the utilization of acoustic sensors, chest-worn inertial measurement units, and smart textiles to monitor respiratory mechanics and sound profiles. Researchers have developed support vector machine algorithms capable of analyzing acoustic data from neck-worn microphones to automatically detect and quantify coughing and wheezing events in real-time [14]. This automated acoustic surveillance provides a much more objective measure of symptom frequency than self-reported patient diaries. Additionally, the extraction of respiration rate and tidal volume proxies from the triaxial accelerometer and gyroscope data of chest-worn patches has been heavily validated [15]. By monitoring the subtle mechanical movements of the chest wall during the respiratory cycle, machine learning models can detect changes in breathing patterns that precede severe chronic obstructive pulmonary disease exacerbations. The integration of environmental data, such as local air quality indices and pollen counts, with wearable physiological data further enhances the contextual awareness of these models, enabling personalized risk forecasting and proactive alerts for patients navigating hazardous environmental conditions [16].

2.4 Evolution of Machine Learning Architectures

The progression of machine learning architectures applied to wearable sensor data mirrors the broader evolution of artificial intelligence. Initially, the field relied heavily on classical machine learning techniques, requiring extensive, domain-expert-driven manual feature extraction. Researchers would manually compute time-domain statistics, frequency-domain spectral densities, and non-linear complexity measures from raw sensor signals before feeding them into classifiers like logistic

regression or decision trees [17]. While interpretable, this approach is highly labor-intensive and often discards subtle, higher-order patterns hidden within the raw data. The subsequent adoption of deep learning architectures, particularly convolutional neural networks and recurrent neural networks, marked a pivotal turning point [18]. Convolutional networks proved highly adept at automatically learning hierarchical spatial and structural feature representations directly from raw multi-channel time-series data, effectively eliminating the need for manual feature engineering. Recurrent networks, especially those with gating mechanisms like long short-term memory cells, excelled at modeling the long-range temporal dependencies and sequential nature of physiological signals [19]. More recently, the literature has shifted toward addressing the privacy and communication overhead challenges associated with cloud-based deep learning. Federated learning has emerged as a crucial paradigm, allowing machine learning models to be trained collaboratively across decentralized edge devices, such as smartphones and smartwatches, without transferring sensitive raw health data to centralized servers [20]. This evolution from central, handcrafted models to decentralized, automated deep learning reflects the maturing of the digital health ecosystem.

3. Methodology and Analytical Framework

3.1 Sensor Data Acquisition Protocols

The foundation of any robust sensor-based machine learning system for noncommunicable disease monitoring rests upon rigorous data acquisition protocols. The selection of appropriate sensors is dictated by the specific physiological targets of the disease under surveillance. For cardiovascular monitoring, clinical-grade wearable electrocardiogram monitors utilize dry electrodes that measure the electrical potential generated by the myocardium. High sampling frequencies, typically ranging from two hundred to five hundred hertz, are required to capture the high-frequency components of the QRS complex necessary for morphological analysis [21]. Conversely, photoplethysmography sensors, which rely on the optical detection of blood volume changes in the microvascular bed of tissue, often operate effectively at lower sampling rates, though they are inherently more susceptible to motion-induced signal corruption. The acquisition protocol must also consider the spatial placement of the sensor, as signal morphology varies significantly depending on whether the device is located on the wrist, chest, or earlobe. Multimodal data acquisition, which involves the simultaneous recording of multiple physiological parameters, is increasingly the standard. This approach leverages the complementary nature of different signals; for example, using a triaxial accelerometer to record body movement alongside a photoplethysmography sensor allows the subsequent algorithms to identify and potentially suppress periods of high motion where the optical signal is likely to be unreliable [22]. Establishing standardized, ethically approved protocols for data collection in free-living environments is essential to ensure the resultant datasets accurately reflect the chaotic, dynamic nature of real-world physiological states.

3.2 Signal Preprocessing and Artifact Removal

Raw data generated by wearable sensors in uncontrolled environments is inevitably contaminated by various sources of noise, artifacts, and missing values. Therefore, a comprehensive signal preprocessing pipeline is a mandatory prerequisite before any machine learning inference can occur. The primary challenge in wearable data is motion artifact, caused by the relative movement between the sensor and the user's skin, which alters the coupling capacitance and optical reflection properties. To address high-frequency noise and low-frequency baseline wander in electrocardiogram signals, cascaded digital filters are universally employed. A typical configuration involves applying a high-pass filter to eliminate respiratory baseline wander and a low-pass filter to remove muscle noise and powerline interference [23]. When dealing with multimodal data, adaptive filtering techniques, where the accelerometer signal serves as a noise reference to actively cancel motion artifacts from the primary physiological signal, are highly effective. Furthermore, signal segmentation or windowing is applied to divide continuous data streams into manageable, overlapping epochs. The choice of window size is a critical hyperparameter; it must be short enough to ensure signal stationarity and capture transient events, yet long enough to contain a sufficient number of complete physiological cycles for meaningful analysis. Addressing missing data, which frequently occurs due to sensor displacement or battery

depletion, requires robust imputation strategies. Simple linear interpolation may suffice for brief gaps, but longer missing segments often require complex predictive imputation using autoregressive models or simple exclusion from the training dataset to prevent the introduction of synthetic bias.

3.3 Feature Extraction Paradigms

Following signal purification, the analytical framework must extract discriminatory characteristics from the data, a process broadly categorized into handcrafted feature extraction and automated feature learning. Handcrafted feature extraction requires significant domain expertise to define metrics that hold physiological relevance. In the time domain, statistical parameters such as the mean, variance, skewness, and kurtosis are calculated over the segmented signal windows. For cardiovascular signals, specific interval features like the distances between consecutive R-peaks are extracted to compute heart rate variability metrics. In the frequency domain, techniques based on the Fourier transform are utilized to analyze the power spectral density of the signals, revealing the distribution of signal energy across different frequency bands. This is particularly useful in respiratory monitoring, where the dominant frequency directly correlates with the respiration rate [24]. Additionally, non-linear features, such as sample entropy and Lyapunov exponents, are extracted to quantify the underlying complexity and chaotic dynamics of physiological systems, which often degrade in the presence of chronic disease [25]. Conversely, the automated feature learning paradigm bypasses this manual calculation phase. Deep learning models ingest the preprocessed raw signal or a two-dimensional transformation, such as a spectrogram, and utilize sequential layers of mathematical operations to autonomously learn optimal feature representations. While this approach often yields superior predictive performance, it trades interpretability for accuracy, creating a black-box scenario that can be problematic in clinical settings where understanding the rationale behind an algorithm's decision is crucial.

3.4 Machine Learning Model Selection

The final stage of the analytical framework involves selecting and training the optimal machine learning model based on the extracted features and the specific clinical objective, which generally falls into the categories of classification, regression, or anomaly detection. For structured, handcrafted feature sets, classical supervised learning algorithms are typically deployed. Support vector machines, which attempt to find the optimal hyperplane separating different disease classes in a high-dimensional feature space, are highly effective for binary classification tasks, such as distinguishing between normal sinus rhythm and atrial fibrillation. Ensemble methods, particularly random forests, are favored for their robustness against overfitting and their ability to provide feature importance rankings, which aids in interpretability [26]. When dealing with the raw or minimally processed time-series data via automated feature learning, deep neural networks are the models of choice. Recurrent architectures are uniquely suited for forecasting tasks, such as predicting future glucose trajectories based on past continuous glucose monitor readings, due to their internal memory mechanisms. For anomaly detection, where the goal is to identify rare physiological deviations indicative of an acute exacerbation without necessarily having labeled examples of the exacerbation, unsupervised learning approaches are utilized. Autoencoders, a type of neural network trained to reconstruct their own input, can be trained exclusively on healthy baseline data. When a patient's physiological state deviates due to an emerging health issue, the autoencoder's reconstruction error increases significantly, serving as a continuous, personalized early warning system.




Figure 1: System Architecture

4. Results and Discussion

4.1 Performance Analysis Across Disease Modalities

The application of the aforementioned analytical frameworks to various noncommunicable diseases yields highly variable performance outcomes, largely dependent on the nature of the physiological signal and the complexity of the disease pathology. In cardiovascular applications, specifically the detection of arrhythmias from wearable electrocardiograms, machine learning models consistently report exceptional performance metrics. Sensitivity and specificity rates frequently exceed ninety-five percent in controlled validations, demonstrating that these systems are highly capable of serving as reliable screening tools [27]. However, when similar deep learning architectures are applied to non-invasive glucose estimation for diabetes management using optical and thermal proxies, the performance degrades significantly. The indirect nature of the measurement, compounded by confounding variables such as ambient temperature and skin pigmentation, results in higher error margins. While these models show strong correlations during controlled clinical studies, their mean absolute relative difference in unconstrained, real-world environments often remains too high for safe, standalone insulin dosing, restricting their current utility to trend analysis and lifestyle management rather than strict clinical diagnosis. Respiratory monitoring models, utilizing acoustic and inertial data, show moderate to high accuracy in detecting gross events like coughing fits but struggle with fine-grained analysis, such as distinguishing between different phenotypes of asthma exacerbations based solely on wearable data [28]. This variance underscores the reality that machine learning is not a monolithic solution; its efficacy is intrinsically linked to the physiological directness and signal-to-noise ratio of the chosen sensor modality.

Table 1: Comparative Performance Metrics of Standard Machine Learning Algorithms on Wearable Data

Algorithm	Classification Objective	Sensitivity	Specificity	Computational Load
Random Forest	Physical Activity Recognition	0.91	0.93	Low
Support Vector Machine	Arrhythmia Detection	0.88	0.90	Moderate
Convolutional Neural Network	ECG Morphology Analysis	0.96	0.97	High
Long Short-	Glucose Trend	0.85	0.82	Very High

Term Memory	Prediction			
Isolation Forest	Unsupervised Anomaly Detection	0.78	0.85	Low

4.2 Computational Efficiency and Resource Constraints

A critical dimension of the discussion surrounding wearable machine learning is the inherent tension between model complexity and the stringent resource constraints of wearable hardware. Advanced deep learning models, while offering superior predictive accuracy, demand significant computational power, large memory footprints, and substantial energy consumption. Executing massive neural networks directly on a battery-powered smartwatch or a disposable biosensor patch is currently unfeasible due to rapid battery depletion and thermal constraints. Consequently, much of the existing architecture relies on a data-streaming paradigm, where raw or lightly compressed data is continuously transmitted via Bluetooth or Wi-Fi to a paired smartphone or a cloud server for heavy processing. However, this approach introduces substantial latency, dependency on continuous network connectivity, and significant privacy vulnerabilities associated with broadcasting sensitive health data. To mitigate these issues, the research community is intensely focused on the miniaturization of machine learning models through techniques such as network quantization, weight pruning, and knowledge distillation. These methodologies aim to compress massive neural networks into highly efficient formats that can execute directly on low-power microcontrollers at the edge of the network. This shift towards edge artificial intelligence is paramount for the development of autonomous, real-time closed-loop therapeutic systems, such as advanced artificial pancreas systems for diabetes, which require instantaneous, localized computation free from network latency or cloud dependence.

4.3 Clinical Implementation Challenges

Despite the impressive technical achievements reported in academic literature, the widespread translation of wearable-based machine learning into routine clinical practice faces formidable obstacles. One of the most pressing challenges is algorithmic bias and the lack of generalizability. Many high-performing models are trained on datasets heavily skewed toward specific demographics, often young, healthy, and predominantly Caucasian cohorts, leading to a severe drop in accuracy when deployed across diverse, older populations with multiple comorbidities. The issue of optical sensor bias concerning varying skin pigmentations in photoplethysmography is a prime example of how hardware and software limitations can inadvertently perpetuate healthcare inequalities. Furthermore, the clinical integration of these technologies is hampered by massive interoperability issues and the lack of standardized data formats. Healthcare systems and electronic health records are notoriously fragmented and ill-equipped to ingest the continuous data streams generated by commercial wearables. Clinicians, already burdened by extreme time constraints, cannot be expected to sift through raw wearable data or unvalidated algorithm outputs. Therefore, predictive models must not only be accurate but must also present their findings as concise, highly interpretable clinical summaries seamlessly integrated into existing diagnostic workflows. Finally, the regulatory landscape remains difficult to navigate. Agencies must carefully evaluate these adaptive machine learning systems, ensuring rigorous evidence of clinical utility, safety, and security before granting approval for use as certified medical devices.

5. Conclusion

5.1 Synthesis of Findings

The integration of wearable sensor technologies and advanced machine learning algorithms represents a monumental shift in the paradigm of noncommunicable disease management. By moving away from episodic clinical snapshots toward continuous, high-resolution physiological monitoring, healthcare systems can gain unprecedented insights into the dynamic trajectories of chronic conditions. As detailed in this synthesis, the analytical framework supporting this transition is complex, requiring meticulous attention to sensor acquisition protocols, rigorous signal preprocessing to mitigate real-world artifacts, and the careful selection of appropriate feature extraction and modeling techniques. The literature clearly demonstrates that in domains with strong, direct physiological signals, such as electrocardiogram-based cardiovascular monitoring, machine learning models have achieved clinical-

grade diagnostic capabilities. However, the application of these technologies to more complex metabolic or multi-factorial respiratory conditions highlights ongoing challenges related to sensor indirectness and environmental confounding variables. Furthermore, the successful deployment of these systems requires balancing the mathematical complexity of advanced deep learning with the severe computational and energetic constraints inherent to wearable edge devices.

5.2 Future Research Trajectories

Looking forward, the maturation of wearable-based machine learning will depend on addressing the critical gaps identified in current research. A primary focus must be the aggressive expansion of representative, diverse, and publicly available multimodal datasets to train models that are robust across all demographic groups, thereby mitigating algorithmic bias. The advancement of federated learning techniques will be crucial in achieving this, allowing institutions to collaboratively train robust models without compromising individual patient data privacy. Additionally, the field must move beyond population-level generalized models toward highly personalized predictive algorithms. Given the immense inter-individual variability in physiological baselines and disease presentations, techniques such as transfer learning and zero-shot learning will be vital for rapidly fine-tuning global models to the unique physiological signatures of individual patients [29]. Finally, interdisciplinary collaboration between materials scientists, data engineers, and clinicians is required to develop next-generation bio-conformable sensors and seamlessly integrate derived digital biomarkers directly into electronic health record workflows. Ultimately, the successful convergence of these technologies holds the profound potential to transform noncommunicable disease care from a system of reactive treatment to one of proactive, continuous, and highly personalized health optimization.

References

- Dong, J., Chen, J., Xie, X., Lai, J., & Chen, H. (2024). Survey on adversarial attack and defense for medical image analysis: Methods and challenges. *ACM Computing Surveys*, 57(3), 1-38.
- Ren, Y., Wu, D., Tong, Y., López-De Fede, A., & Gareau, S. (2023, November). Using machine learning methods to identify racial differences of low birthweight risks: Evidence from a state-level large-scale data in the us. In *APHA 2023 Annual Meeting and Expo*. APHA.
- Huang, H., Zhang, J., Zhang, J., Xu, J., & Wu, Q. (2020). Low-rank pairwise alignment bilinear network for few-shot fine-grained image classification. *IEEE Transactions on Multimedia*, 23, 1666-1680.
- Dong, J., Koniusz, P., Zhang, Y., Zhu, H., Liu, W., Qu, X., & Ong, Y. S. (2025, October). Improving zero-shot adversarial robustness in vision-language models by closed-form alignment of adversarial path simplices. In *Forty-second International Conference on Machine Learning*.
- Dong, J., Koniusz, P., Qu, X., & Ong, Y. S. (2025, July). Stabilizing modality gap & lowering gradient norms improve zero-shot adversarial robustness of vlms. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1* (pp. 236-247).
- Lv, Qi, et al. "F1: A vision-language-action model bridging understanding and generation to actions." *arXiv preprint arXiv:2509.06951* (2025).
- Wang, X., Xiao, H., Zhao, Y., Li, P., Hu, X., Qiu, X., ... & Luo, Y. (2026). Toward Omni Healthcare: Wearable Technologies for Noncommunicable Disease Dynamic Monitoring. *Medicine Bulletin*, 2(1), 5-19.
- Zhao, H., Gu, J., Wang, S., Lu, T., Zhang, X., Wu, Z., ... & Jiang, Y. G. (2026). LSTD: Long Short-Term Temporal Diffusion for Video Generation. *IEEE Transactions on Multimedia*.
- Xie, C., Zhu, D., Wang, Z., Zhang, H., & Wei, Z. (2026). Explainable AI for Mental Health Detection from Social Media: A Comparative Study of Traditional Machine Learning and a Large Language Model. Available at SSRN 6429778.
- Cai, Lie, et al. "Deep learning model for breast shear wave elastography to improve breast cancer diagnosis (INSPIRED 006): An international, multicenter analysis." *Journal of Clinical Oncology* 43.32 (2025): 3482-3493.
- Ren, Y., Wu, D., Tong Ph D, Y., Lopez-De Fede Ph D, A., & Gareau, S. J. (2024). Benchmarking Machine Learning Models to Predict Low Birthweight Baby Outcomes and Identify Associated

- Risk Factors from an Extremely Unbalanced Large-Scale Dataset. *Women's Health Issues*, 34(72), 79.
- Zhang, W., Zhu, J., Zhang, Y., Sun, L., Wang, K., Dong, Y., ... & Shang, A. (2025). Personalized prediction of lymph node metastasis in papillary thyroid microcarcinoma: a nomogram and web calculator. *Scientific Reports*.
- Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep patient: An unsupervised representation to predict the future of patients from the electronic health records. *Scientific Reports*, 6, 26094.
- Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C. S., Liang, H., Baxter, S. L., et al. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5), 1122–1131.
- Wang, J., & Zhang, M. (2020). Deepflash: An efficient network for learning-based medical image registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4444-4452).
- Chu, Y. B., Li, J., Jia, P., Cui, J., Zhang, R., Kang, X., ... & Zhang, S. (2021). Irf1-and Egr1-activated transcription plays a key role in macrophage polarization: A multiomics sequencing study with partial validation. *International Immunopharmacology*, 99, 108072.
- Tu, P., Huang, Y., Zheng, F., He, Z., Cao, L., & Shao, L. (2022, June). Guidedmix-net: Semi-supervised semantic segmentation by using labeled images as reference. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, No. 2, pp. 2379-2387).
- Dong, J., Qu, X., Zhang, C., Rong, S. Q., Thai, N. D., Pan, W., ... & Ong, Y. S. (2026). Tug-of-war no more: Harmonizing accuracy and robustness in vision-language models via stability-aware task vector merging. In *The Fourteenth International Conference on Learning Representations*.
- Zhao, Y., Zhu, Z., Yang, S., & Li, W. (2026). YeastSAM: A Deep Learning Model for Accurate Segmentation of Yeast Cells. *Molecular Biology of the Cell*, mbc-E25.
- Yu, X., Li, J., Wang, J., Wang, Y., Li, S., Ma, X., ... & Wang, H. (2026). Prediction of risk of hearing loss by industry noise from cross-sectional and longitudinal data. *Communications Medicine*.
- Bagheri, Mahdi, et al. "An MRI Atlas of the Human Fetal Brain: Reference and Segmentation Tools for Fetal Brain MRI Analysis." *ArXiv* (2025): arXiv-2508.
- Yao, Y., Li, Y., Zheng, S., & Zhu, T. (2025). Deep Learning Models for Multi-class Pneumonia Detection in Chest X-rays: A Comparative Study of VGG16, MobileNet, and ResNet152. *Computer Simulation in Application*, 3(1), 66-73.
- Wang, Y., Ma, X., Huang, X., Li, S., Yu, X., Huang, S., ... & Yin, S. (2026). Integrating environmental factors and genetic variants in machine learning to assess occupational noise impact on health. *Ecotoxicology and Environmental Safety*, 311, 119805.
- Liu, Y., & Kwon, H. (2025). Efficient Depth Estimation for Unstable Stereo Camera Systems on AR Glasses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 6252-6261).
- Yao, Y., Li, Y., Zheng, S., & Zhu, T. (2025). MSA-TransUNet: A Multi-Scale Attention Enhanced Transformer-UNet Architecture for Accurate Vessel Segmentation and Visualization in Medical Imaging. *Journal of Medicine and Life Sciences*, 1(4), 48-57.
- Xie, C., Zhu, D., Wang, Z., Zhang, H., & Wei, Z. (2026). Compliance-Aware Discharge Agent for Auditable ICU Discharge Planning: A Pilot Feasibility Study Using Structured eICU Records. Available at SSRN 6429758.
- Guo, Y., Sekiguchi, Y., Zeng, W., Ebihara, S., Owaki, D., & Hayashibe, M. (2025). Physics-informed learning framework for lower limb kinematic prediction with sparse sensors and its application in chronic stroke. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*.
- Ding, S., Yu, X., Wang, Q., Luo, P., Li, H., Li, Z., ... & Zhang, C. (2025). Multifunctional hydrogel sensors with dynamic covalent networks for machine learning-assisted Parkinson's disease diagnosis and encrypted human-computer interaction. *Materials Today Bio*, 102524.

Xu, H., Ge, J., & Ou, Y. (2026). DentaScope-AI: A Real-Time k-mer-Driven CNN-Attention Framework for Rapid Pathogen Identification and Virulence Prediction in Endodontic Infections Using Nanopore Sequencing. *Frontiers in Healthcare Technology*, 3(1), 112-122.