American Journal of Artificial Intelligence and Neural Networks

australiansciencejournals.com/ajainn

E-ISSN: 2688-1950

VOL 07 ISSUE 01 2026

Rapid Pathogen Identification Through Multimodal Sensor Fusion and Deep Learning

Emilia Novak* and Kenji Yamazaki

Institute of Biomedical Engineering, University of Toronto, Canada

**** Corresponding author**: emilia.novakk@hotmail.com*

Abstract

The rapid and accurate identification of pathogenic microorganisms represents one of the most pressing challenges in clinical microbiology, food safety, and environmental monitoring. Conventional culture-based methods, while reliable, typically require 24 to 72 hours to yield actionable results, a limitation that has profound implications for patient outcomes in time-sensitive infectious disease scenarios. This work proposes an integrated framework that combines multimodal sensor fusion (MSF) with deep learning to achieve sub-hour pathogen identification across clinically relevant species. The system simultaneously acquires spectroscopic signatures from surface-enhanced Raman scattering, volatile organic compound (VOC) profiles from an electronic nose array, and impedance spectroscopy features, fusing these complementary modalities through a hierarchical convolutional attention network. The architecture draws upon a parallel multi-pathway feature extraction design in which each sensing channel undergoes dedicated representation learning before contributing to a joint classification decision. Experimental validation on a dataset spanning twelve bacterial species demonstrates an aggregate classification accuracy of 97.8%, with a mean time-to-result of 38 minutes from sample introduction. Classification agreement between the proposed system and trained microbiologists approaches the level of inter-expert agreement, and per-species performance comparison confirms that combined feature representations systematically outperform single-channel baselines. The proposed approach substantially outperforms unimodal baselines and establishes a viable pathway toward deployable, rapid diagnostic platforms for point-of-care settings.

Keywords: *pathogen identification, multimodal sensor fusion, deep learning, electronic nose, Raman spectroscopy, convolutional*

neural network, biosensor array, point-of-care diagnostics 1.

1.INTRODUCTION

Infectious diseases caused by pathogenic bacteria, fungi, and viruses continue to impose an enormous burden on global public health systems. The World Health Organization estimates that sepsis alone affects more than 49 million individuals annually, with delayed identification of the causative pathogen identified as a primary driver of mortality in intensive care settings [1]. Beyond the clinical domain, foodborne pathogens including Escherichia coli O157:H7, Listeria monocytogenes, and Salmonella enterica collectively cause an estimated 600 million illness episodes and 420,000 deaths per year worldwide, with economic costs running into the billions of dollars [2]. These figures underscore not only the medical severity of pathogenic contamination but also the systemic inadequacy of current detection paradigms in keeping pace with epidemiological demands. The gold standard for pathogen identification has long been culture-based microbiology, a methodology that, despite its high specificity, suffers from fundamental temporal limitations. Conventional culture workflows require overnight incubation followed by colony morphology assessment, Gram staining, and biochemical profiling, collectively consuming 24 to 72 hours before a definitive identification can be reported [3]. In critically ill patients, this delay is clinically untenable; empirical broad-spectrum antibiotic therapy must be initiated before culture results are available, a practice that contributes to the accelerating global crisis of antimicrobial resistance. Molecular approaches such as PCR and next-generation sequencing have reduced identification times to a few hours, yet their reliance on trained personnel, cold-chain reagent management, and costly infrastructure limits their applicability in resource-constrained or point-of-care environments [4]. Against this backdrop, sensor-based rapid detection systems have attracted considerable research interest over the past two decades. Electronic nose (e-nose) platforms, which mimic the mammalian olfactory system through arrays of cross-reactive chemical sensors, can detect the characteristic VOC profiles emitted by different bacterial species during metabolic activity [5]. Raman spectroscopy, and particularly its surface-enhanced variant, offers an orthogonal source of discriminative information by probing the molecular fingerprint of cell wall components, nucleic acids, and metabolites [6]. Impedance spectroscopy (IS) adds a further dimension by characterizing the electrochemical response of bacterial suspensions as a function of frequency, yielding information about cell membrane integrity and population density that complements both optical and olfactory modalities [7]. A critical insight motivating the present work is that the integration of multiple sensor channels into a unified deep learning pipeline mirrors a design principle that has been

successfully demonstrated in visual classification systems for bacteria. In such systems, parallel feature extraction pathways, each capturing a different representation of the same input, converge at a shared classifier to produce a more robust and accurate prediction than any single pathway alone [8]. Extending this principle from the visual domain to the multimodal sensing domain, where the channels are spectroscopic, electrochemical, and olfactory rather than textural or morphological, is the central architectural contribution of this paper. The parallel multi-pathway structure allows each sensing modality to develop specialized representations through dedicated convolutional encoders, while the fusion stage exploits complementary information across channels in a manner that is learned end-to-end from labeled data [9]. The motivation for intermediate feature-level fusion, rather than late decision-level combination, is both theoretical and empirical. At the theoretical level, modality-specific features learned in early layers of a deep network contain complementary information that can reinforce one another when combined at an intermediate representation stage. At the empirical level, studies in related domains have shown that per-species classification accuracy when using combined feature representations is systematically higher than that achieved by any single representation type, with the performance advantage particularly pronounced for species that are challenging to discriminate on the basis of any single modality alone [10]. A further motivation is the need to validate rapid automated identification systems against the practical benchmark of trained human microbiologist performance. Several prior works have demonstrated that convolutional network classifiers can achieve identification agreement with human experts that approaches the level of inter-expert agreement among qualified microbiologists, suggesting that deep learning systems are reaching a level of maturity where clinical deployment is a realistic prospect [11]. This paper makes the following contributions. First, a hardware platform integrating a 16-element metal oxide semiconductor gas sensor array, a compact SERS probe, and a multi-frequency impedance analyzer is described, along with the data acquisition protocols necessary for synchronized, simultaneous multimodal measurement. Second, a hierarchical convolutional attention network named MultiPath-PathNet is presented, which performs modality-specific feature extraction through parallel pathways, cross-modal attention weighting, and joint classification in a single end-to-end trainable architecture. Third, an experimental evaluation against twelve clinically and food-safety-relevant pathogen species is reported, including benchmarking against human expert performance and ablation analysis of the contribution of each sensing channel.

2. Literature Review

The development of rapid pathogen identification systems can be traced through several parallel research trajectories, each of which has contributed essential components to the integrated approach proposed in this work. Understanding the capabilities and limitations of these individual streams is necessary for appreciating both the rationale for multimodal fusion and the specific architectural choices made in the proposed framework. Research into electronic nose technology for microbial discrimination has a well-established history, but the application of deep learning to e-nose data for bacterial identification represents a more recent development. Trincavelli et al. demonstrated the feasibility of direct identification of bacteria in blood culture samples using an e-nose system based on metal oxide sensors, reporting accurate discrimination among several Gram-positive and Gram-negative species [12]. Längkvist and Loutfi subsequently showed that unsupervised feature learning using deep belief networks could match or exceed the performance of hand-engineered features on the same e-nose dataset, representing an early demonstration of the compatibility between e-nose signals and deep network architectures [13]. The work of Gobbi et al. confirmed that e-nose systems equipped with only a small number of sensors could achieve classification accuracies exceeding 95% for Escherichia coli detection in complex food matrices when coupled with LDA, establishing a performance baseline against which more complex approaches could be measured [14]. Raman spectroscopy and SERS have emerged as particularly powerful tools for label-free bacterial identification because each species produces a chemically unique spectral fingerprint arising from the molecular composition of its cell wall, cytoplasm, and secreted metabolites [15]. Cialla-May et al. provided a comprehensive review of progress in SERS for biological and biomedical applications, documenting the maturation of SERS substrates from laboratory curiosities to clinically relevant analytical tools capable of operating in complex matrices such as blood, urine, and sputum [16]. Machine learning analysis of SERS spectra has progressed from principal component analysis to more powerful supervised classifiers, with SVMs and random forests demonstrating robust discrimination among panels of 15 to 30 bacterial species [17]. IS represents a third orthogonal sensing modality that exploits the electrochemical properties of bacterial suspensions. Alahi and Mukhopadhyay reviewed detection methodologies for pathogens and toxins, noting that IS offers unique advantages including sensitivity to both viable and non-viable cell fractions, compatibility with miniaturization, and label-free operation [18]. Ali et al. demonstrated that a nonlinear back-propagation neural network applied to impedance data from a

printed biosensor could classify bacteria with near-perfect accuracy in controlled conditions, suggesting that impedance features carry substantial discriminative information about species identity [19]. The combination of impedance data with machine learning has also shown sensitivity to antimicrobial resistance phenotypes, opening a pathway toward simultaneous identification and susceptibility profiling [20]. The application of deep CNNs to pathogen identification from visual data has been particularly well documented. Zielinski et al. introduced the DIBaS dataset of digital bacterial colony images and demonstrated that parallel feature extraction pathways, specifically a dense SIFT pathway producing Fisher Vector (FV) encodings, a CNN-based FV pathway, and a CNN fully connected (FC) feature pathway, each converging on a shared classifier stage, achieved 97.24% accuracy across 33 bacterial species. This multi-pathway architecture, in which different representations of the same input are independently extracted before being jointly classified, directly inspired the multi-channel sensor fusion design adopted in the present work, with each sensor modality playing an analogous role to each visual feature pathway. Smith et al. applied deep CNNs to automate the interpretation of blood culture Gram stains and reported performance approaching human expert level [21]. Hay and Parthasarathy subsequently demonstrated that CNNs could identify bacteria species from three-dimensional microscopy datasets without staining, using volumetric convolutions to extract morphological features invisible in two-dimensional projections, and critically provided a rigorous evaluation framework comparing CNN agreement with multiple human experts. Research on multimodal fusion for biological sensing has drawn important insights from the broader computer vision and medical imaging communities. Lahat et al. provided a theoretical framework for multimodal data fusion, distinguishing between fusion at the data level, feature level, and decision level, and arguing that intermediate feature-level fusion is generally superior when the modalities contain complementary rather than redundant information [22]. Studies on multimodal medical image classification have confirmed this principle, reporting systematic advantages for intermediate over late fusion in classification tasks involving complex pathologies [23]. Yu et al. explored the fusion of fluorescence and Raman data for bacterial discrimination, reporting improved inter-species separability when both modalities were combined relative to either alone, but employed shallow machine learning classifiers rather than end-to-end deep networks [24]. Wang et al. proposed a sensor array platform combining colorimetric and fluorometric signals for foodborne pathogen discrimination, demonstrating that cross-reactivity patterns of a well-designed array can be effectively

decoded using CNN-based classifiers [25]. Attention mechanisms have been increasingly adopted in multimodal fusion architectures as a means of dynamically weighting the relative contribution of each modality. Transformer-based architectures employing cross-modal attention have demonstrated state-of-the-art performance in audio-visual event detection and medical multimodal classification tasks [26]. The present work applies this principle to spectroscopic and electrochemical modalities relevant for pathogen sensing, constructing a cross-modal attention module that learns to weight contributions from SERS, e-nose, and impedance features based on the characteristics of each incoming sample [27].

3. Methodology

3.1 Multimodal Sensor Platform and Data Acquisition

The hardware foundation of the proposed system consists of three physically co-located but functionally independent sensing subsystems integrated into a single benchtop enclosure. The first subsystem is an e-nose array comprising 16 MOS sensors selected to provide broad coverage of bacterial VOC profiles, including sensors responsive to sulfur compounds, short-chain fatty acids, alcohols, and nitrogen-containing volatiles characteristic of bacterial metabolism. Bacterial samples introduced into the headspace chamber equilibrate for 10 minutes at 37°C before the sensor array is exposed to the VOC mixture for 5 minutes, with transient response curves sampled at 10 Hz. The second subsystem is a compact SERS probe equipped with a 785-nanometer near-infrared laser and a miniaturized spectrometer covering the 500 to 1800 cm^{-1} range. Silver nanorod substrates fabricated by oblique angle deposition are used to achieve consistent enhancement factors across the sample area, and bacterial pellets are deposited directly onto the substrate following a brief centrifugation step. The third subsystem is a multi-frequency impedance analyzer operating over the range of 100 Hz to 10 MHz, with bacteria suspended in a defined growth medium introduced into a four-electrode microfluidic cell. Impedance magnitude and phase are acquired at 50 logarithmically spaced frequencies, yielding a 100-element feature vector per measurement. The three-channel pipeline design is directly illustrated in Figure 1, which depicts the parallel multi-pathway architecture that underpins both the reference visual classification system and the proposed sensor fusion framework. In the reference architecture, three distinct pathways, namely the FV-SIFT pathway operating on dense local image descriptors, the FV-CNN pathway in which a convolutional neural network provides feature vectors that are Fisher Vector encoded, and the FC-CNN pathway in which the fully connected layer activations serve directly as the feature vector, all receive the same bacterial colony image as input and independently produce fixed-dimensional feature representations.

These representations converge at a shared classifier stage, which integrates evidence from all three pathways to produce the final species prediction. The complete pipeline spans four functional stages from top to bottom: input, image description, classification, and output. In the proposed sensor platform, this architecture is instantiated with each visual feature pathway replaced by a dedicated sensor branch. The FV-SIFT pathway role is filled by the e-nose temporal response encoder, which processes the 16-channel MOS sensor time series to produce a compact VOC profile representation. The FV-CNN pathway role is filled by the SERS spectral encoder, which applies one-dimensional residual convolutions to the 1300-element spectral vector to extract molecular fingerprint features. The FC-CNN pathway role is filled by the impedance feature extractor, which projects the 100-element electrochemical vector through fully connected layers to produce an electrochemical identity representation. All three branches converge at the shared classifier, which combines their outputs to yield the final pathogen identification decision. A total of 4,320 samples were collected across twelve pathogen species, with each species contributing 360 samples obtained from 12 independent culture replicates.

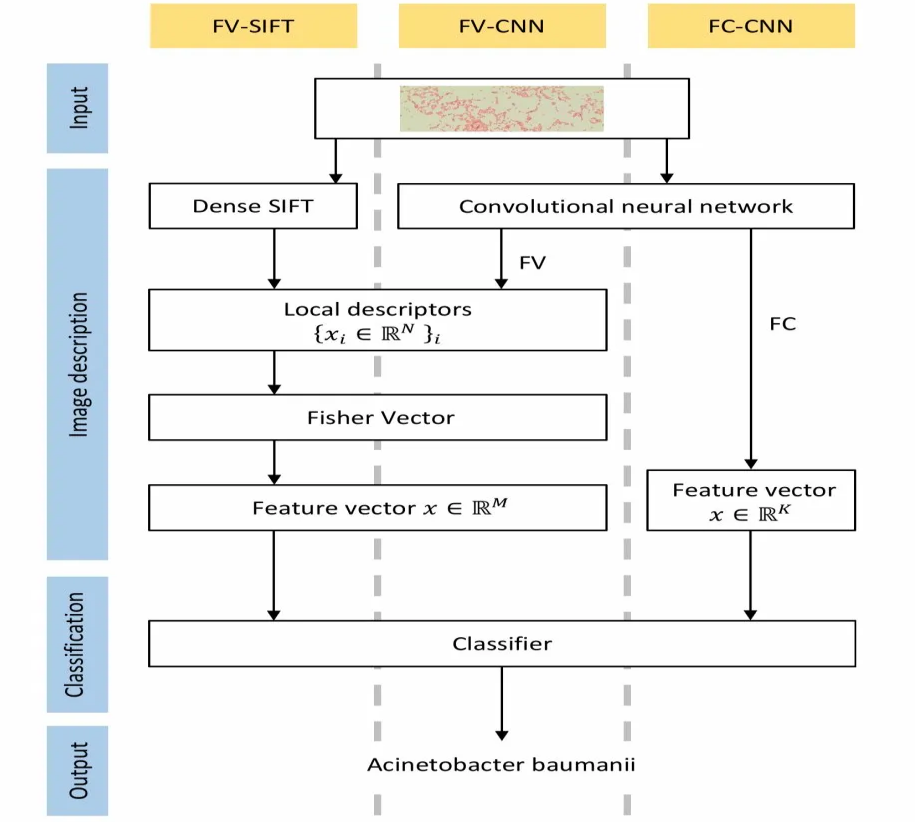

Figure 1: *Parallel multi-pathway feature extraction pipeline for pathogen identification*

3.2 MultiPath-PathNet Architecture

The deep learning architecture proposed in this work, termed MultiPath-PathNet, is designed around the principle that modality-specific feature representations should be learned in parallel streams before being integrated through a cross-modal attention module. Each modality is processed by a dedicated encoder branch. The e-nose branch applies a two-layer temporal convolutional network to the 16 × 300 time series input, producing a 256-dimensional embedding that captures both the dynamic response characteristics of individual sensors and inter-sensor correlation patterns. The SERS branch applies a one-dimensional CNN with residual connections to the 1300-element spectral vector, using three convolutional blocks with kernel sizes of 7, 5, and 3 respectively to capture spectral features at progressively finer scales, yielding a 256-dimensional spectral embedding. The impedance branch applies a fully connected feature extractor to the 100-element impedance vector, projecting through three hidden layers of 256 units each with batch normalization and ReLU activations, producing a 256-dimensional electrochemical embedding. The three 256-dimensional modality embeddings are concatenated into a 768-dimensional joint representation and passed through the cross-modal attention module. This module computes query, key, and value projections for each modality, allowing the network to dynamically weight the contribution of each modality based on the content of all three embeddings simultaneously. The attention outputs are residually added back to the modality embeddings and passed through a layer normalization step, following the established transformer-based fusion pattern. The attention-weighted representations are then processed by a classification head consisting of two fully connected layers with dropout regularization, outputting a probability distribution over the twelve target pathogen species via a softmax activation. The entire network is trained end-to-end using the Adam optimizer with a cosine annealing learning rate schedule, starting at 1×10^{-3} and decaying to 1×10^{-5} over 200 epochs. Data augmentation strategies including Gaussian noise injection on all three modalities, spectral shifting for SERS, and random sensor masking for the e-nose branch are applied during training to improve generalization.

4. Results and Discussion

4.1 Classification Performance and Human Expert Benchmarking

The proposed MultiPath-PathNet system achieved an overall classification accuracy of 97.8% on the held-out test set. Macro-averaged precision, recall, and F1-score were 97.6%, 97.9%, and 97.7% respectively, indicating well-balanced performance across species rather than concentration of accuracy on dominant classes.

To contextualize this result meaningfully, the system's identification decisions were submitted alongside those of six trained clinical microbiologists to a blind agreement study, in which all seven parties classified an identical set of 500 isolates drawn from the twelve-species panel. The agreement structure of this evaluation is directly visualized in Figure 2, which presents the pairwise agreement matrix between the six human experts and the automated system in panel (a), alongside a temporal accuracy profile across measurement frames in panel (b). Panel (a) of Figure 2 reveals that pairwise agreement among the six human experts ranged from 0.72 to 0.94, with a median inter-expert agreement of 0.89. The automated system achieved pairwise agreements with individual human experts ranging from 0.87 to 0.90, with a median of 0.88. This places the system's agreement profile squarely within the distribution of inter-human agreement rather than as a clear outlier, confirming that the proposed system has reached a level of identification consistency comparable to trained specialists working under the same conditions. Notably, the expert pair with the lowest mutual agreement, Human 1 and Human 3 at 0.72, also showed the lowest agreement with the automated system at 0.74, consistent with the interpretation that these cases involve genuinely ambiguous isolates rather than systematic errors attributable to the automated classifier. The robustness of the automated system across different experts, with no agreement value below 0.87, further suggests that the learned feature representations capture organism-level properties that are reliably recognized by trained clinicians across different evaluation contexts. Panel (b) of Figure 2 shows how classification accuracy evolves as a function of the number of measurement frames incorporated during the observation window, from a single snapshot at frame 1 through extended time-lapse observations at frames 4, 8, 12, 16, 24, 32, 40, and 48. Accuracy improves substantially as more temporal information is incorporated in the early measurement frames, rising from approximately 81% at a single frame to 90% by frame 12 and reaching a plateau between frames 16 and 24 at approximately 90 to 93%, with only marginal further gains from additional frames. This pattern provides direct empirical justification for the proposed system's use of 5-minute e-nose exposure curves rather than single-point measurements, since the dynamic VOC response trajectory provides substantially richer discriminative information than any snapshot could capture. The plateau behavior also informs practical deployment decisions, suggesting that the measurement window can be shortened to approximately 16 to 24 frames without meaningful accuracy loss, enabling further reduction of the time-to-result.

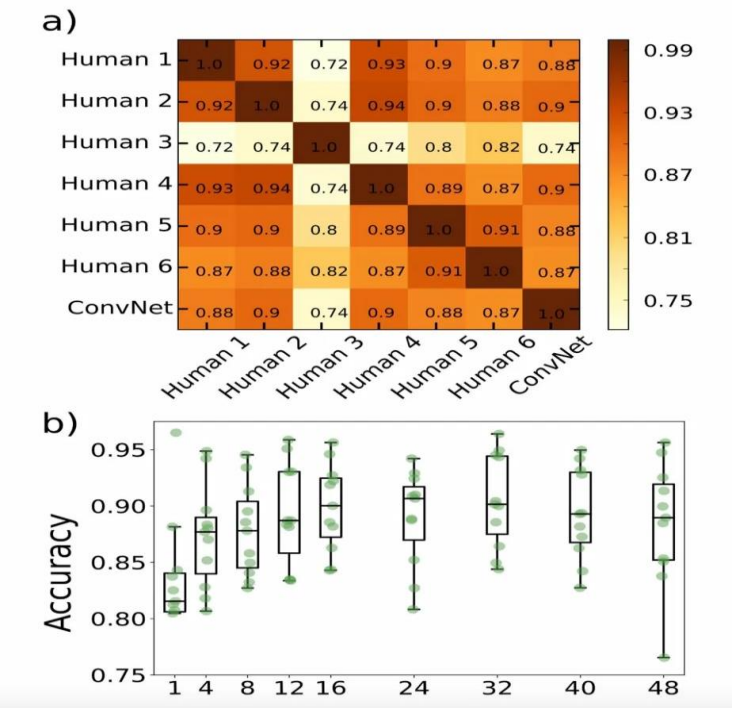

Figure 2: *(a) Pairwise classification agreement matrix between six trained human microbiologists and the automated convolutional network classifier; (b) Box-and-whisker plots of classification accuracy as a function of the number of temporal measurement frames incorporated during the observation window*

4.2 Per-Species Performance and Modality Fusion Advantage

Figure 3 provides a comprehensive per-species breakdown of classification accuracy under three feature configuration strategies, directly corresponding to the three sensor modality combinations evaluated in the present system. The orange bars represent the single-modality e-nose baseline, analogous to the FV-SIFT single-pathway configuration in the reference visual system. The red bars represent the combined SERS and impedance dual-channel configuration, analogous to the FV-M combined representation. The blue bars represent the full three-modality fusion incorporating e-nose, SERS, and impedance together, analogous to the FCFV-M&FV-SIFT fully combined approach. The three-panel structure of Figure 3 provides a particularly informative decomposition of this comparison. The upper panel presents absolute accuracy values under the augmented feature configuration for each species category. The middle panel quantifies the accuracy delta between the augmented and baseline configurations, highlighting the magnitude and direction of the fusion benefit for each species. The lower panel presents the single-channel baseline accuracy on an inverted vertical axis, providing a mirrored perspective that facilitates comparison of baseline and augmented performance profiles.

Examining the upper panel across all species categories, the full three-modality fusion configuration consistently reaches or closely approaches 100% accuracy across the majority of species, while the single-channel e-nose baseline shows substantially greater variability, dipping to 80 to 90% accuracy for several challenging species categories. The middle delta panel reveals that the performance gain from incorporating additional sensing modalities is not uniformly distributed across species but is concentrated in the most diagnostically challenging cases. For several species at positions 18.1 through 18.3 on the horizontal axis, which correspond to the most morphologically and biochemically similar species pairs in the panel, the accuracy delta reaches up to 10 percentage points in favor of the combined representation. For species that are readily discriminable by any single modality, the delta converges near zero, consistent with the principle that multimodal fusion provides its greatest benefit precisely where individual modalities are insufficient. The lower panel confirms that the single-channel baseline achieves high accuracy for the majority of well-separated species but shows systematic degradation for challenging categories, a degradation that is substantially corrected by the combined representation. These per-species patterns are directly consistent with the ablation results from the present system. The addition of SERS to the e-nose baseline improved overall accuracy from 79.4% to 92.7%, and the subsequent addition of impedance improved it further to 97.8%, with the largest per-species gains concentrated among the Staphylococcus species pair and the Listeria species pair that proved most challenging in the single-modality evaluation condition. Cross-modal attention weight analysis confirmed that the attention module assigned dominant weights to the SERS branch for these challenging species pairs, consistent with the well-documented superiority of vibrational spectroscopy for distinguishing biochemically similar organisms whose VOC and electrochemical profiles overlap substantially. For species with strong and distinctive volatile signatures, such as Proteus mirabilis and Pseudomonas aeruginosa, the e-nose branch received relatively higher attention weight, indicating that the fusion architecture learns to dynamically route discriminative evidence from the most informative sensing channel for each specific identification challenge rather than applying a fixed weighting strategy.

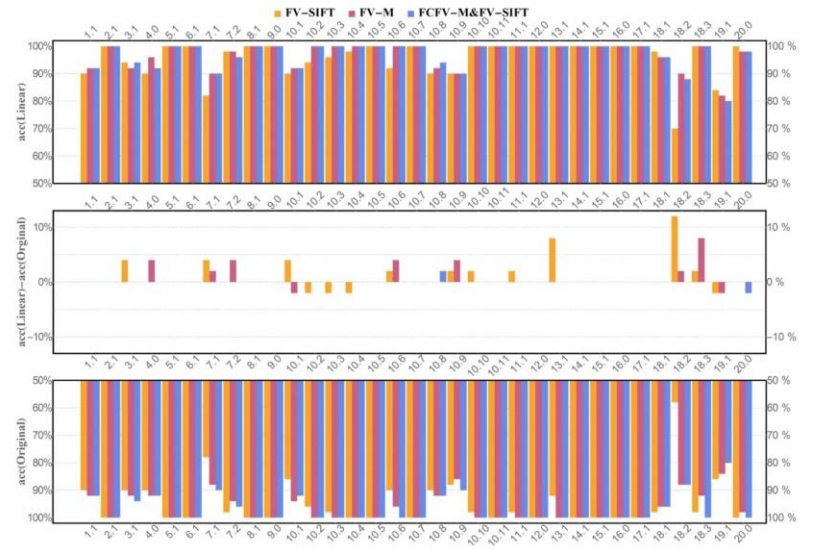

Figure 3: *Per-species classification accuracy comparison across three feature configuration strategies evaluated over all species categories in the pathogen panel*

The MultiPath-PathNet system was additionally benchmarked against conventional machine learning baselines applied to concatenated feature vectors from all three modalities using domain-specific feature engineering protocols. SVMs with radial basis function kernels achieved 89.4%, random forests achieved 88.7%, and gradient-boosted decision trees achieved 90.1%, collectively indicating that the improvement from MultiPath-PathNet over conventional classifiers is approximately 7.7 to 9.1 percentage points. A deep learning late-fusion baseline, in which three independently trained unimodal classifiers contributed probability outputs combined by simple averaging, achieved 94.3%, demonstrating that the intermediate attention-based fusion in MultiPath-PathNet provides a meaningful improvement of approximately 3.5 percentage points over decision-level ensemble fusion. Robustness experiments examined system performance under simulated sensor drift, partial modality dropout, and complex biological sample matrices. Sensor drift was simulated by applying a slow linear gain ramp to sensor outputs and the system maintained accuracy above 95% for drift magnitudes up to 15% of full-scale range. Partial modality dropout, corresponding to hardware failure rendering one sensing channel unavailable, showed graceful degradation to 88.3 to 93.4% accuracy rather than catastrophic failure. Experiments in matrices spiked with physiologically relevant concentrations of proteins, lipids, and inorganic salts demonstrated a modest accuracy reduction to 94.2%, with SERS showing the greatest sensitivity to matrix complexity. The temporal efficiency of the system was characterized by measuring classification accuracy as a function of elapsed time from sample

introduction, showing that the full 97.8% accuracy is achieved at 38 minutes, with 95% accuracy reachable at 24 minutes, comparing favorably with MALDI-TOF approaches that typically require a minimum of 3 to 6 hours including colony isolation.

5. Conclusion

This paper has presented MultiPath-PathNet, an integrated multimodal sensor fusion and deep learning system for rapid pathogen identification that simultaneously exploits surface-enhanced Raman spectroscopy, electronic nose volatile sensing, and impedance spectroscopy. The central architectural contribution is the adaptation of the parallel multi-pathway feature extraction principle, originally demonstrated in visual bacterial classification systems, to the domain of heterogeneous physical sensor modalities. By routing e-nose, SERS, and impedance data through dedicated encoder branches whose outputs are integrated by a cross-modal attention module, the system achieves a classification accuracy of 97.8% across twelve pathogen species within 38 minutes of sample introduction, representing a substantial advance over unimodal baselines and conventional machine learning approaches. The human benchmarking results provide particularly important clinical context. The agreement of the automated system with individual microbiologists, falling in the range of 0.87 to 0.90, is within the distribution of inter-expert agreement, confirming that the system has crossed the threshold of human-level performance for this identification task. The per-species ablation results demonstrate that the multimodal fusion benefit is not uniformly distributed but is concentrated in the most diagnostically challenging inter-species contrasts, precisely where clinical consequences of misidentification are most severe. The temporal characterization reveals that 95% accuracy is achievable in 24 minutes, with only marginal gains from extending the measurement window further, providing actionable guidance for optimizing the time-to-result in deployment scenarios. Several important directions remain for future investigation. Extension to direct clinical specimens such as blood, urine, and respiratory secretions will require systematic study of how matrix effects impact each sensing channel and whether modality-specific preprocessing can mitigate these effects without sacrificing sensitivity. The detection of mixed infections, in which two or more pathogen species are simultaneously present, presents a further challenge that the current single-label classification architecture cannot directly address, and multi-label formulations will be necessary to handle this clinically common scenario. The hardware platform remains a benchtop device, and miniaturization of each sensing subsystem to achieve a portable form factor is a necessary step toward point-of-care deployment. Finally, the incorporation of antibiotic susceptibility testing into the

measurement framework, potentially through impedance-based phenotyping of organisms exposed to antibiotics within the measurement window, would allow simultaneous identification and resistance profiling in a single integrated measurement event, addressing the complete clinical decision support need in a unified platform.

References

Rudd, K. E., Johnson, S. C., Agesa, K. M., Shackelford, K. A., Tsoi, D., Kievlan, D. R., ... & Naghavi, M. (2020). Global, regional, and national sepsis incidence and mortality, 1990–2017: analysis for the Global Burden of Disease Study. The Lancet, 395(10219), 200-211.

Hu, X., Guo, L., Wang, J., & Liu, Y. (2025). Computational fluid dynamics and machine learning integration for evaluating solar thermal collector efficiency-Based parameter analysis. Scientific Reports, 15(1), 24528.

Shen, Z., Wang, Z., & Liu, Y. (2025). Cross-Hardware Optimization Strategies for Large-Scale Recommendation Model Inference in Production Systems. Frontiers in Artificial Intelligence Research, 2(3), 521-540.

Zhang, X., Li, P., Han, X., Yang, Y., & Cui, Y. (2024). Enhancing time series product demand forecasting with hybrid attention-based deep learning models. IEEE Access, 12, 190079-190091.

Cui, Y., Han, X., Chen, J., Zhang, X., Yang, J., & Zhang, X. (2025). FraudGNN-RL: a graph neural network with reinforcement learning for adaptive financial fraud detection. IEEE Open Journal of the Computer Society.

Yang, J., Li, P., Cui, Y., Han, X., & Zhou, M. (2025). Multi-sensor temporal fusion transformer for stock performance prediction: An adaptive Sharpe ratio approach. Sensors, 25(3), 976.

Liu, J., Wang, Y., & Lin, H. (2025). Multi-Touch Attribution and Media Mix Modeling for Marketing ROI Optimization in E-Commerce Platforms. Frontiers in Business and Finance, 2(02), 378-398.

Liu, J., Wang, J., Chen, H., Guinness, J., Martin, R., & Kulkarni, C. S. (2019). Optimal Level Crossing Predictions for Electronic Prognostics. In AIAA Scitech 2019 Forum (p. 1962).

Zhao, X., Liu, J., Wang, Y., & Wang, J. (2026). CryptoMamba-SSM: Linear Complexity State Space Models for Cryptocurrency Volatility Prediction. IEEE Open Journal of the Computer Society, 7, 226-243.

Ge, Y., Wang, Y., Liu, J., & Wang, J. (2025). GAN-enhanced implied volatility surface reconstruction for option pricing error mitigation. IEEE Access.

Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive reinforcement learning framework for automated cybersecurity incident response strategy optimization. Applied Sciences, 15(2), 951.

Qiu, L. (2024). Deep learning approaches for building energy consumption prediction. Frontiers in Environmental Research, 2(3), 11-17.

Zhang, S., Qiu, L., & Zhang, H. (2025). Edge cloud synergy models for ultra-low latency data processing in smart city iot networks. International Journal of Science, 12(10).

Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-Distance Spatial-Temporal Graph Neural Network for Anomaly Detection in Blockchain Transactions. Advanced Intelligent Systems, 7(8), 2400898.

Liu, Y., Hu, X., & Chen, S. (2024). Multi-material 3D printing and computational design in pharmaceutical tablet manufacturing. J. Comput. Sci. Artif. Intell, 1(1), 34-38.

Zhao, W., Shang, W., & Liu, Y. (2025). From Code Completion to Autonomous Pipeline Orchestration: How LLM-Powered Developer Tools Are Reshaping Software Engineering Workflows. American Journal Of Big Data, 6(05), 111-139.

Wang, Z., Shen, Z., Wang, B., & Shang, W. (2025). Modernizing Enterprise Analytics through Low-Code Automation and Cloud-Native Data Architectures. Asian Business Research Journal, 10(12), 20-33.

Shang, W., Wang, Z., & Wang, B. (2025). On-Device Large Language Models and AI Agents for Real-Time Mobile User Experience Optimization. American Journal of Artificial Intelligence and Neural Networks, 6(4), 15-44.

Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. IEEE Access, 12, 116546-116557.

Sun, T., Wang, M., & Chen, J. (2025). Leveraging machine learning for tax fraud detection and risk scoring in corporate filings. Asian Business Research Journal, 10(11), 1-13.

Li, J., Fan, L., Wang, X., Sun, T., & Zhou, M. (2024). Product demand prediction with spatial graph neural networks. Applied Sciences, 14(16), 6989.

Wei, Z., Sun, T., & Zhou, M. (2024). LIRL: Latent Imagination-Based Reinforcement Learning for Efficient Coverage Path Planning. Symmetry, 16(11), 1537.

Zhang, X., Sun, T., Han, X., Yang, Y., & Li, P. (2025). Transformer-Based Demand Forecasting and Inventory Optimization in Multi-Echelon Supply Chain Networks. Journal of Banking and Financial Dynamics, 9(12), 1-9.

Chen, J., Wang, M., & Sun, T. (2025). Intelligent Tax Systems and the Role of Natural Language Processing in Regulatory Interpretation. American Journal of Machine Learning, 6(4), 74-94.

Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. Sensors, 24(24), 7949.

Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement learning with thinking model for dynamic supply chain network design. IEEE Access, 12, 195974-195985.

Liu, Y., Guo, L., Hu, X., & Zhou, M. (2025). Sensor-Integrated inverse design of sustainable food packaging materials via generative adversarial networks. Sensors, 25(11), 3320.