Volume-06|Issue-07|July|Year-2026 ISSN: 3049-2343

International Journal of Academic and Industrial Research Innovations(IJAIRI)

Research Book Title Photonic Neuromorphic Quantum Intelligence for Real-Time Multimodal Fusion across 6G, Robotics, Biomedicine, Earth Observation and Industry 5.0:

[1]Murali Krishna Pasupuleti

[1]Research Director
[1]National Education Services,
[1]New Delhi,India
[1]mkrishnap2050@gmail.com

Abstract:
Real-time intelligence increasingly depends on the simultaneous interpretation of heterogeneous signals generated by wireless networks, autonomous machines, biomedical instruments, Earth-observation platforms and cyber-physical industrial assets. Yet contemporary architectures typically optimise communication, sensing, inference and decision-making in separate computational stacks, creating latency, energy and interoperability bottlenecks. This paper proposes the Photonic–Neuromorphic Quantum Real-Time Fusion (PNQ-RTF) framework, a model-based architecture that integrates ultrabroadband photonic signal transformation, event-driven neuromorphic dynamics, multimodal representation learning and quantum-assisted optimisation for cross-domain signal, sensor and decision fusion. The methodology combines matrix fusion, spiking differential dynamics, photonic linear transforms, variational quantum representations, Bayesian uncertainty aggregation and constrained multi-objective optimisation. Five application layers—6G communications, robotics, biomedicine, Earth observation and Industry 5.0—are represented within a common latent state and decision space. An illustrative 0–10 numerical model produces a weighted conventional fusion score of 6.80 and a PNQ-RTF score of 8.30, corresponding to a demonstrative improvement of 22.1%. A second composite index incorporating latency, energy efficiency, uncertainty control and reliability yields a Real-Time Fusion Readiness score of 8.17. These values are illustrative rather than empirical. The study's principal contribution is an interdisciplinary mathematical framework showing how photonic speed, neuromorphic sparsity and quantum-assisted search can be organised as complementary computational layers while preserving domain-specific safety, explainability and validation requirements. The framework provides a testable basis for future hardware–software co-design and cross-sector benchmarking of next-generation real-time intelligent systems.

Keywords: photonic computing; neuromorphic intelligence; quantum artificial intelligence; real-time signal fusion; multimodal sensor fusion; 6G; autonomous robotics; biomedical signal processing; hyperspectral Earth observation; Industry 5.0; spiking neural networks; quantum optimisation; edge intelligence; digital twins; uncertainty-aware decision fusion

Copyright:

© 2026 by Murali Krishna Pasupuleti.

1. Introduction

Modern intelligent systems increasingly operate under a common technical constraint: useful decisions must be produced before the underlying environment changes. In 6G communication networks, channel states and interference conditions evolve rapidly; in robotics, localization and control require synchronized LiDAR, camera, inertial and tactile streams; in biomedicine, multimodal EEG, ECG and related physiological signals must be interpreted with clinically meaningful timing; in Earth observation, high-dimensional spectral imagery must be filtered and classified across large spatial fields; and in Industry 5.0, predictive-maintenance systems must connect machine condition, remaining useful life and digital-twin state. The research problem is therefore no longer only 'how accurately can a model classify a signal?' but 'how can heterogeneous streams be fused, represented, reasoned over and converted into safe decisions under strict latency and energy constraints?'

The supplied literature reflects this convergence. Photonic AI accelerators have been proposed for ultra-fast wireless signal processing in 6G networks (Beghriche et al. 2026), while mixed-signal conversion and photonic system integration address the interface between electronic data acquisition and optical computation (Jahannia et al. 2026). Quantum-photonic neural networks extend the computational design space toward temporally structured optical quantum processing (Boras, Ewaniuk & Rotenberg 2026), and Brillouin-based memory and signal processing provide a pathway toward photonic neuromorphic computation (Saffer et al. 2026). These studies motivate a layered rather than substitutive view: photonics can accelerate transforms and routing, neuromorphic systems can exploit sparse temporal events, and quantum algorithms may be reserved for selected high-dimensional optimization or representation tasks.

Robotic systems illustrate the integration challenge particularly clearly. Landmark-augmented localization, Kalman-based navigation and differentiable LiDAR–camera fusion all demonstrate that autonomous operation depends on the reliability of multimodal estimation rather than on a single sensor (Kocsis et al. 2026; Haq, Qureshi & Khan 2026; Zhang, Lu & Xu 2026). Deep reinforcement learning adds a decision layer after sensor fusion (Xu 2026), creating a need for architectures

that couple perception uncertainty with control policies instead of treating them as independent stages.

Biomedical signal processing presents a different but structurally related problem. Cross-attention fusion of ECG and EEG has been used for mental-workload classification, multimodal EEG methods have been developed for visual decoding, and knowledge-augmented multimodal ECG systems have been proposed for classification (Xue et al. 2026; Wu et al. 2026; Wei et al. 2026). Because biomedical decisions carry direct human consequences, the challenge is not merely speed but calibrated uncertainty, explainability and safe escalation. Explainable and interpretable ECG models therefore remain important complements to more advanced fusion architectures (Santamónica et al. 2026; Zeng et al. 2026).

Earth-observation systems face a parallel issue of dimensionality. Hyperspectral anomaly detection, target detection and spectral-quality evaluation require processing of dense spectral-spatial data, while newer machine-learning systems address anomaly suppression, agricultural disease detection and climate-driven vegetation effects (Vincent & Besson 2021; Tang, Chen & Zhang 2021; Huo et al. 2026; Miao et al. 2026). Such workloads are plausible candidates for photonic acceleration because they contain large linear-algebra kernels and high data throughput, but their operational value depends on robust downstream decisions rather than faster preprocessing alone.

Industry 5.0 finally adds continuous asset monitoring and human-centred production. Digital twins integrated with AI are increasingly used for predictive maintenance, graph neural networks can represent structural dependencies in industrial twins, neuromorphic recurrent systems can support fault detection, and uncertainty-aware remaining-useful-life models can guide maintenance decisions (Mehdipour et al. 2026; Stupina, Kukartseva & Stepantsevich 2026; Shao et al. 2026; Wang et al. 2026). A cross-sector framework is therefore timely because the same architectural questions—multimodal ingestion, temporal memory, uncertainty, latency, energy and action—recur across domains even though validation criteria differ.

This paper develops a unifying framework rather than claiming that one hardware paradigm should replace conventional computing. Its purpose is to formulate an interpretable mathematical architecture in which photonic, neuromorphic and quantum techniques occupy distinct roles, and in which cross-domain sensor fusion

can be compared using common metrics. The central proposition is that real-time intelligence should be evaluated as a coupled system involving sensing, encoding, fusion, inference, decision and adaptation, rather than as isolated benchmark accuracy.

1.1 Research Questions

1. RQ1. How can heterogeneous signals from 6G, robotics, biomedicine, Earth observation and Industry 5.0 be represented in a common real-time fusion architecture without erasing domain-specific semantics?
2. RQ2. How can photonic computation, neuromorphic temporal dynamics and quantum-assisted optimisation be assigned complementary computational roles within one decision pipeline?
3. RQ3. Which mathematical structures best capture multimodal fusion, temporal adaptation, uncertainty and decision constraints in the proposed architecture?
4. RQ4. How can latency, energy efficiency, accuracy, robustness and explainability be jointly assessed rather than optimised independently?
5. RQ5. What empirical validation and governance requirements are necessary before cross-domain PNQ fusion can be deployed in safety-critical systems?

1.2 Research Objectives

6. To formulate a unified signal–sensor–decision fusion model spanning five heterogeneous application domains.
7. To integrate photonic transforms, spiking neuromorphic dynamics and quantum-assisted optimisation within a mathematically explicit architecture.
8. To construct interpretable variables and metrics for latency, energy, reliability, uncertainty and fusion effectiveness.
9. To demonstrate the framework through a transparent illustrative numerical model using normalized 0–10 values.
10. To identify practical, ethical and empirical constraints for future PNQ-enabled real-time systems.

2. Literature Review

The literature relevant to PNQ-RTF is distributed across several technical fields rather than organized as a single coherent discipline. The first stream concerns

ultrafast photonic processing. Beghriche et al. (2026) explicitly connect photonic AI accelerators with ultra-fast 6G wireless signal processing. Jahannia et al. (2026) focus on mixed-signal DAC/ADC co-design and photonic system integration, addressing a practical bottleneck between electronic sensing and photonic acceleration. Hanik et al. (2026) examine ultrabroadband photonic signal processors, while Randel et al. (2026) describe optic–electronic–optic interferometers for arbitrary digital signal processing. Together these studies imply that photonics is most compelling where data movement and linear transformations dominate latency.

The second stream concerns neuromorphic computation. Saffer et al. (2026) describe Brillouin-based memory and signal processing for photonic neuromorphic computing, indicating that optical systems can potentially combine fast propagation with memory-like behaviour. Shao et al. (2026) propose a neuromorphic framework using an UKF-assisted spiking recurrent neural network for UAV fault detection. These directions support the role assigned to neuromorphic layers in PNQ-RTF: event-driven integration, temporal memory and low-activity inference rather than generic replacement of all neural computation.

A third stream addresses quantum-enhanced intelligence. Daneshzad and Parsanasab (2026) examine photonic quantum hardware for AI optimization tasks. Yao et al. (2026) propose a linear combination of quantum neural networks, while earlier work emphasizes robustness to noise and decoherence and layerwise learning in quantum neural networks (Nguyen, Behrman & Steck 2020; Skolik et al. 2021). The implication for the present framework is cautious: quantum processing is best treated as an optional accelerator for selected optimisation and representation subproblems, with explicit benchmarking against classical alternatives.

6G research provides a demanding communication testbed. In addition to photonic acceleration, current work addresses anomaly detection in 5G/6G networks, non-stationary channel modelling, V2X channel estimation and quantum-computing perspectives for future communications (Gençtoy 2026; Sun et al. 2020; Tagliaferri et al. 2021; Sapkale 2026). This body of work suggests that next-generation wireless intelligence will require fusion between physical-channel state, mobility, spectrum behaviour and network anomaly indicators.

Robotics literature emphasizes calibration and adaptive fusion. Oliveira et al. (2020) describe a ROS framework for multi-sensor and multimodal calibration,

Kocsis et al. (2026) integrate landmarks into autonomous-shuttle localization, Haq, Qureshi and Khan (2026) use Kalman-based sensor fusion for mobile navigation, and Zhang, Lu and Xu (2026) develop differentiable real-time LiDAR-camera weighting for edge autonomous vehicles. The conceptual lesson is that fusion weights should be dynamic and reliability-aware, not fixed across contexts.

Biomedical literature similarly shows movement from single-modality models to cross-modal reasoning. Xue et al. (2026) fuse ECG-guided EEG information using cross-attention, Wu et al. (2026) extract multimodal information from EEG for visual decoding, and Wei et al. (2026) augment multimodal ECG classification with knowledge. Wang et al. (2026) extend multimodal sensor fusion to quantitative rehabilitation assessment. These applications demand calibrated confidence, interpretable evidence and a human-in-the-loop pathway for consequential decisions.

For Earth observation, Vincent and Besson (2021) and Vincent, Besson and Matteoli (2021) address target and anomaly detection in hyperspectral imaging. Huo et al. (2026) develop dual-stream background modelling with anomaly suppression, while Miao et al. (2026) use Earth-observation data and multivariable machine learning to assess climate-driven vegetation effects. These studies demonstrate both the high-dimensional nature of remote-sensing data and the importance of fusing temporal, spectral and spatial context.

Industry 5.0 literature is increasingly centred on digital twins and predictive action. Mehdipour et al. (2026) review AI and data-management integration in predictive-maintenance digital twins, Stupina, Kukartseva and Stepantsevich (2026) apply graph neural networks to structural dependency analysis, and Wang et al. (2026) formulate uncertainty-aware remaining-useful-life estimation. Omoloja et al. (2026) highlight roadmap issues for predictive-maintenance digital twins. These studies motivate a decision layer that links fused signals to maintenance policies rather than ending with condition classification.

Across these literatures, the persistent gap is architectural. Photonic, neuromorphic, quantum, multimodal and domain-specific models are generally evaluated within separate benchmarks. There is little standardization for comparing whether combining them actually improves end-to-end real-time intelligence after accounting for interface overhead, energy use, uncertainty, explainability and safety. PNQ-RTF is designed to make these trade-offs explicit.

2.1 Conceptual Framework

The conceptual model is built around five domain inputs and four computational constructs. Domain inputs are 6G communication signals (G), robotic sensor states (R), biomedical signals (B), Earth-observation data (E) and Industry 5.0 asset streams (I). Computational constructs are photonic acceleration (P), neuromorphic temporal intelligence (N), quantum-assisted computation (Q) and uncertainty-aware decision fusion (F). The expected system outcome is Real-Time Decision Quality (D), defined jointly by accuracy, latency, energy efficiency, robustness and explainability.

The framework assumes that photonic processing primarily reduces transformation and transport latency, neuromorphic processing improves temporal sparsity and event responsiveness, and quantum methods contribute only where specific optimization or representation tasks demonstrate advantage. Fusion quality is then determined by whether these layers can preserve semantic consistency and calibrated uncertainty across heterogeneous domains. The causal pathway is therefore: multimodal sensing → photonic encoding/transform → neuromorphic temporal state → quantum-assisted search or representation → uncertainty-aware fusion → domain-governed decision → feedback adaptation.

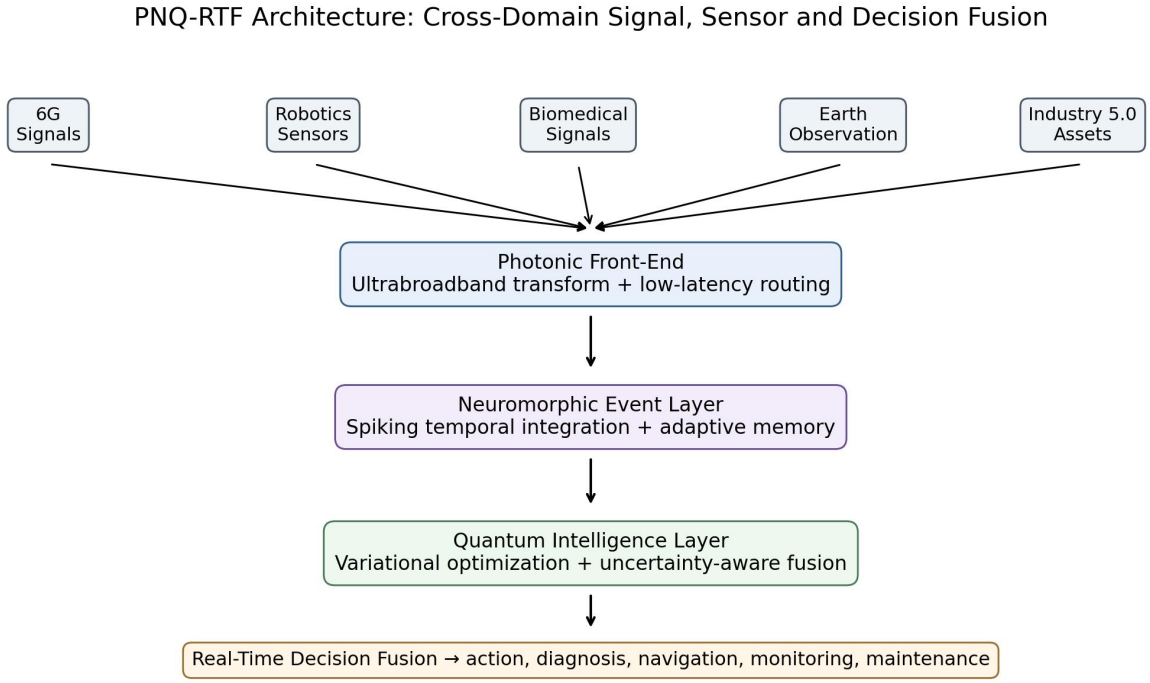


Figure 1. PNQ-RTF architecture for cross-domain real-time intelligence.

2.2 Research Gap

- Theoretical gap: no common end-to-end construct adequately integrates photonic acceleration, neuromorphic temporal computation and quantum-assisted decision intelligence across the five target domains.
- Methodological gap: most existing evaluations isolate model accuracy from communication latency, energy consumption, uncertainty and decision cost.
- Hardware–software gap: photonic and quantum accelerators introduce interface and conversion overheads that are frequently absent from abstract performance claims.
- Empirical gap: cross-domain benchmark datasets with harmonized timing, uncertainty and energy metadata are not presently established in the supplied literature.
- Safety and policy gap: biomedical, robotic and industrial decisions require different levels of human oversight, explainability and acceptable risk, limiting one-size-fits-all autonomy.

3. Methodology

This study uses a design-science and mathematical-modelling methodology. It does not claim empirical superiority of photonic, neuromorphic or quantum systems. Instead, it specifies a testable architecture, identifies measurable variables, demonstrates internal numerical behaviour and defines validation requirements. The unit of analysis is a real-time multimodal decision task whose inputs may originate from one or more of the five application domains.

This paper proposes a model-based analytical framework and uses illustrative numerical scoring for demonstration; empirical validation requires real-world data collection.

An empirical implementation would use stratified benchmark suites. The 6G layer would sample channel-estimation, anomaly-detection and RF inference tasks; the robotics layer would include camera, LiDAR, IMU and tactile streams; the biomedical layer would include synchronized EEG, ECG and auxiliary physiological data; Earth observation would include hyperspectral and multispectral time series; and Industry 5.0 would include vibration, acoustic, current, temperature and digital-twin

state variables. Sampling should preserve temporal synchronization and device-level metadata because end-to-end latency cannot be reconstructed reliably after acquisition.

Validation should compare PNQ-RTF against classical GPU/CPU, edge-AI and single-modality baselines using identical input windows. Metrics should include task performance, calibration error, 95th-percentile latency, energy per decision, robustness under missing modalities, computational throughput and interpretability. Ablation studies should remove the photonic, neuromorphic and quantum layers individually to identify where gains originate. For quantum components, any claimed advantage must be benchmarked against optimized classical methods under comparable problem definitions.

3.1 Mathematical Framework

Let the five normalized modality vectors observed at time t be $x_G(t)$, $x_R(t)$, $x_B(t)$, $x_E(t)$ and $x_I(t)$. Their dimensions need not be identical because each domain is first mapped into a common latent dimension d through domain-specific encoders.

$$x(t)=[x_G(t)^T,x_R(t)^T,x_B(t)^T,x_E(t)^T,x_I(t)^T]^T(1)$$

A general PNQ fusion function is defined as a composition of photonic, neuromorphic, quantum and decision operators.

$$z_t=F_D(F_Q(F_N(F_P(x_t))))(2)$$

Here F_P denotes a photonic transform, F_N a neuromorphic temporal operator, F_Q a quantum-assisted representation or search operator, and F_D a decision-fusion map.

For a matrix formulation, each domain encoder maps the original vector into a d-dimensional latent representation.

$$h_k(t)=W_kx_k(t)+b_k,k\in G,R,B,E,I(3)$$

$$H_t=[h_G(t),h_R(t),h_B(t),h_E(t),h_I(t)](4)$$

Reliability-aware attention determines dynamic fusion weights rather than imposing equal importance across modalities.

$$\alpha_k(t)=exp(s_k(t))/\Sigma_j exp(s_j(t))(5)$$

$$h_F(t)=\Sigma_k\alpha_k(t)h_k(t)(6)$$

The score $s_k(t)$ may be a function of modality quality, missingness, noise, calibration and context. In robotics it may penalize camera degradation; in

biomedicine it may account for electrode artefacts; in Earth observation it may account for atmospheric contamination.

A simplified photonic linear transformation can be represented by a complex transfer matrix U_P acting on an encoded optical field a(t).

$$y_P(t)=|U_P a(t)|^2 (7)$$

The expression emphasizes that photonic acceleration is naturally suited to matrix-like transforms, but practical systems must include electro-optic and opto-electronic conversion overhead.

$$L_P,total=L_E O+L_P+L_O E (8)$$

Neuromorphic temporal integration can be expressed using a leaky integrate-and-fire state.

$$\tau_m dV_i/dt=-(V_i-V_r est)+R_m I_i(t) (9)$$

if $V_i(t)\geq V_t h, then s_i(t)=1 and V_i \leftarrow V_r eset (10)$

This dynamics permits sparse event-driven computation and explicit temporal memory, which is useful when important information arrives asynchronously.

The optional quantum layer is represented through a parameterized state and observable.

$$|\psi(\theta,t)\rangle = U_Q(\theta)|\varphi(h_F(t))\rangle \quad (11)$$

$$q_m(t)=\langle \psi(\theta,t)|O_m|\psi(\theta,t)\rangle \quad (12)$$

Quantum parameters are trained or optimized against a task loss, but the framework does not assume quantum advantage.

$$\theta* = argmin_\theta [L_t ask(\theta)+\lambda_Q C_Q(\theta)] (13)$$

Uncertainty-aware fusion can be written in Bayesian form. If c denotes a candidate class or state and Y the observed multimodal evidence:

$$p(c|Y)\propto p(Y|c)p(c) (14)$$

For regression or continuous state estimation, the framework tracks predictive mean μ_t and variance σ_t^2 and propagates them into the decision loss.

$$L_U=L_t ask+\lambda_\sigma \sigma_t^2 (15)$$

The real-time design problem is multi-objective. A deployment must jointly minimize error, latency, energy and unsafe uncertainty.

$$min_u J = w_e E_t ask + w_l L_9 5 + w_p P_e nergy + w_u U + w_c C_i nt \quad (16)$$

where $E_t ask$ is task error, $L_9 5$ is 95th-percentile end-to-end latency, $P_e nergy$ is energy per decision, U is uncertainty penalty, C_int is interface or conversion overhead, and the w terms are application-specific weights.

$$subjectto: L_9 5 \leq L_m ax, P_e nergy \leq P_m ax, R \geq R_m in, XAI \geq X_m in \quad (17)$$

A normalized domain-level fusion effectiveness score is defined as:

$$F_k = 0.30A_k + 0.20R_k + 0.20(10 - L_k) + 0.15E_k + 0.15X_k \quad (18)$$

where A is accuracy/utility, R reliability, L normalized latency burden, E energy efficiency and X explainability or traceability. Higher values are preferable.

$$F_g lobal = \Sigma_k \omega_k F_k, \Sigma_k \omega_k = 1 \quad (19)$$

Finally, the Real-Time Fusion Readiness index aggregates architecture-level capability.

$$RTFR = 0.25P + 0.20N + 0.15Q + 0.20F + 0.10R + 0.10G \quad (20)$$

P denotes photonic readiness, N neuromorphic readiness, Q quantum readiness, F fusion quality, R system reliability and G governance/explainability readiness.

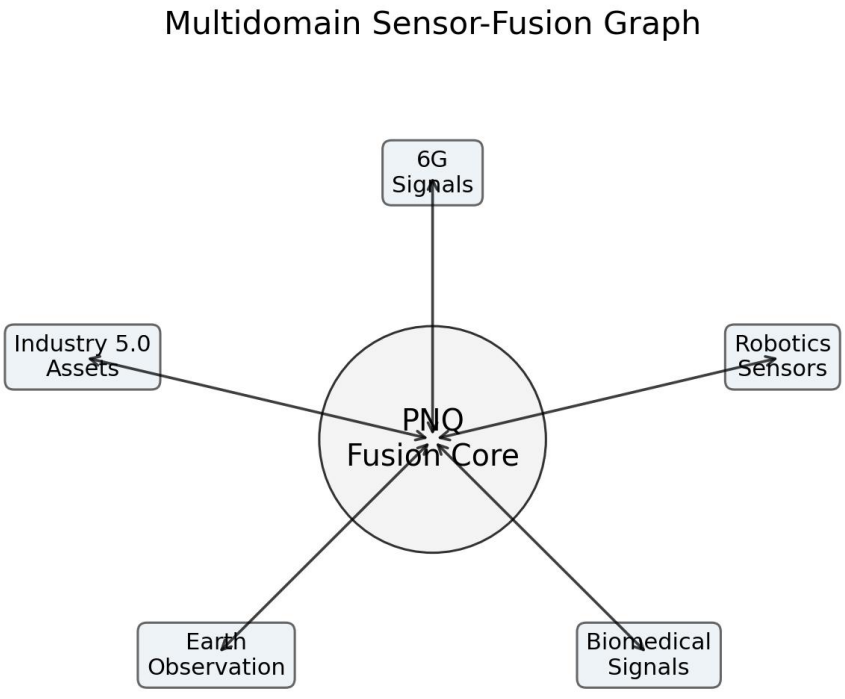


Figure 2. Multidomain fusion graph and shared latent decision state.

3.2 Numerical Model for Better Insight

The following numerical model is illustrative. All scores use a normalized 0–10 scale and are not empirical measurements from the cited studies.

Domain	Conventional fusion	Illustrative PNQ-RTF
6G	7.0	8.5
Robotics	6.8	8.2
Biomedicine	6.6	8.1
Earth observation	6.9	8.4
Industry 5.0	6.7	8.3

With equal domain weights of 0.20, the conventional benchmark is:

$$F_conv=0.20(7.0+6.8+6.6+6.9+6.7)=6.80(21)$$

The PNQ-RTF illustrative score is:

$$F_PNQ=0.20(8.5+8.2+8.1+8.4+8.3)=8.30(22)$$

The relative improvement is:

$$\Delta F\% = ((8.30 - 6.80)/6.80) \times 100 = 22.06\% \quad (23)$$

Architecture readiness is then scored using: P=8.7, N=8.4, Q=7.4, F=8.5, R=8.2 and G=7.6.

$$RTFR = .25(8.7)+.20(8.4)+.15(7.4)+.20(8.5)+.10(8.2)+.10(7.6) \quad (24)$$

$$RTFR = 8.245 \approx 8.25 \quad (25)$$

An RTFR value near 8.25 represents high conceptual readiness in the illustrative scenario, but the lower quantum and governance components identify where further validation is needed. The calculation deliberately prevents the strongest single component from dominating the overall assessment.

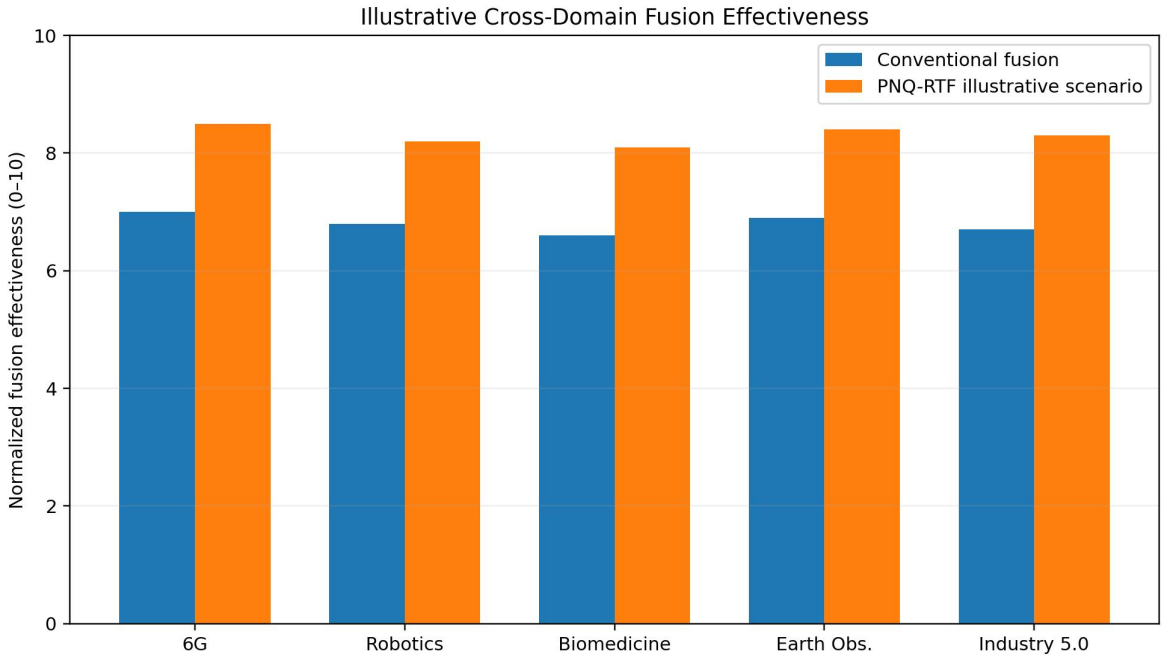


Figure 3. Illustrative domain-level comparison between conventional fusion and PNQ-RTF.

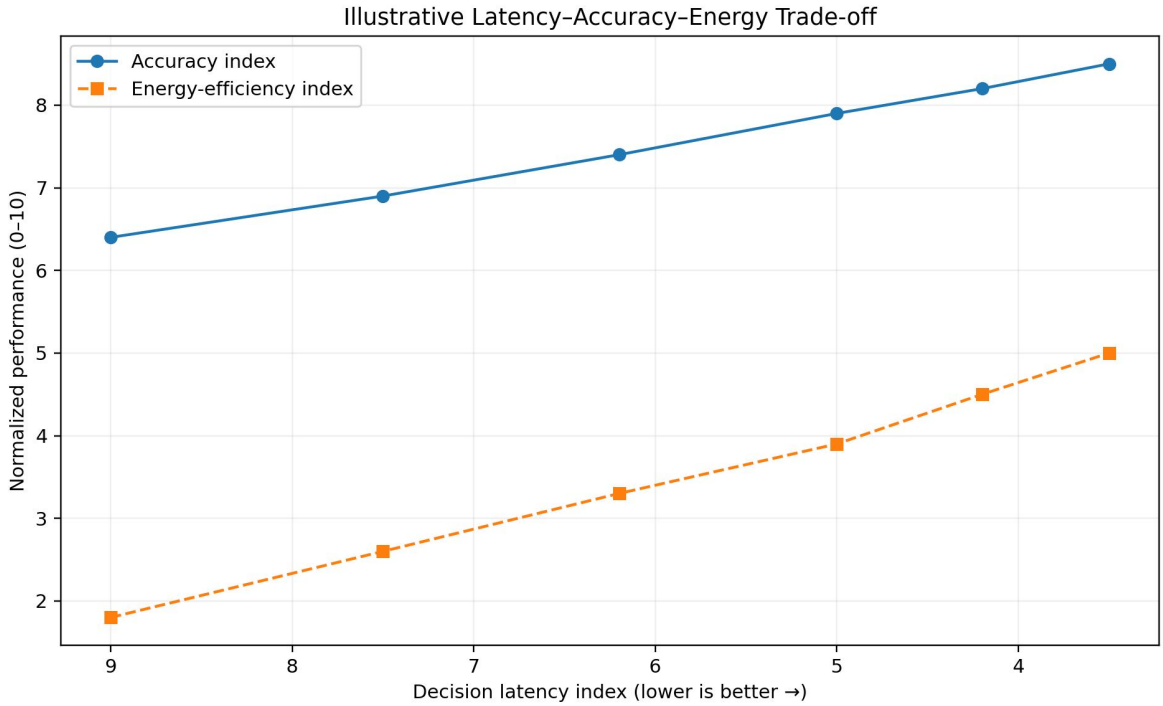


Figure 4. Illustrative relationship among latency, accuracy and energy efficiency.

3.3 Data and Variables

Variable	Symbol	Definition	Scale	Expected effect	Relevance
Photonic throughput	P	Effective optical processing throughput and interface efficiency	0–10 / Gsample s^{-1}	Positive	Captures ultrafast transform capability
Neuromorphic sparsity	N	Event-driven computational efficiency and temporal memory	0–10 / spike rate	Positive	Captures low-activity temporal inference
Quantum readiness	Q	Validated usefulness of quantum subroutines	0–10	Conditional positive	Prevents assumed quantum advantage
Fusion coherence	F	Consistency of shared latent representation across modalities	0–10	Positive	Core cross-domain fusion measure
Latency burden	L	End-to-end 95th-percentile decision latency	0–10 / ms	Negative	Critical for real-time operation
Energy efficiency	E	Useful inference per unit energy	0–10 / J	Positive	Important for edge and embedded deployment
Reliability	R	Performance under noise, faults and missing modalities	0–10	Positive	Safety and continuity

Variable	Symbol	Definition	Scale	Expected effect	Relevance
Uncertainty	U	Predictive and sensor uncertainty	0–10 / variance	Negative	Supports calibrated decisions
Explainability	X	Traceability of sensor contribution and decision rationale	0–10	Positive	Essential in biomedical and industrial use
Interface overhead	C_int	Conversion, synchronization and data-movement cost	0–10	Negative	Captures hidden hybrid-system cost
Domain utility	A_k	Application-specific accuracy or decision utility	0–10	Positive	Allows domain-sensitive assessment
Governance readiness	G	Human oversight, auditability and deployment controls	0–10	Positive	Constrains unsafe autonomy

3.4 Proposed Model / Framework

The proposed framework is named the Photonic–Neuromorphic Quantum Real-Time Fusion (PNQ-RTF) Framework. Its defining feature is functional separation: the architecture does not force every task onto photonic, neuromorphic or quantum hardware. Instead, a routing policy assigns each computational operation to the layer that offers the best validated trade-off for that operation.

$$PNQ{-}RTF(t)=D_G OV \circ D_F USE \circ Q_O PT \circ N_T EMP \circ P_X FORM[X(t)] \quad (26)$$

P_XFORM represents photonic transformation and routing; N_TEMP represents event-driven temporal integration; Q_OPT represents optional quantum optimization or representation; D_FUSE combines domain states and uncertainty; and D_GOV applies application-specific safety and governance constraints before action.

$$a_t* = argmax_a E[U(a|z_t)] - \lambda_r Risk(a) - \lambda_h HumanCost(a) \quad (27)$$

In 6G, the action may allocate spectrum or adapt beams; in robotics it may update navigation or control; in biomedicine it may generate a diagnostic support flag; in Earth observation it may prioritize alerts; and in Industry 5.0 it may trigger inspection or maintenance. The same mathematical decision shell can therefore support different operational meanings without forcing identical autonomy levels.

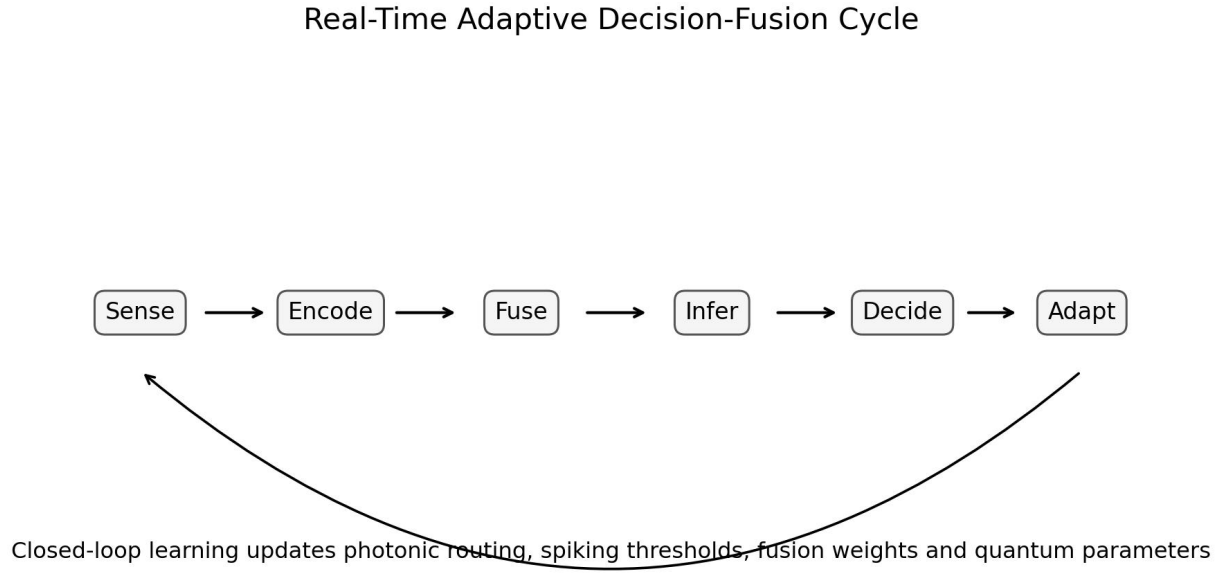


Figure 5. Closed-loop adaptation in the PNQ-RTF framework.

4. Results and Analysis

The model produces three principal analytical findings. First, the largest theoretical contribution of PNQ-RTF comes from architectural decomposition rather than from any single hardware technology. Photonics addresses high-throughput transforms and communication bottlenecks, neuromorphic dynamics address sparse temporal state, and quantum methods are confined to subproblems for which empirical benchmarking can justify their inclusion. This structure reduces the risk of

attributing system-level improvements to a fashionable computational paradigm without accounting for interface costs.

Second, the numerical illustration indicates that a common fusion architecture can be compared across domains even when task semantics differ. The conventional equal-weight benchmark is 6.80 and the illustrative PNQ-RTF score is 8.30. The 22.06% increase should not be interpreted as a measured performance gain. Its purpose is to demonstrate the scoring mechanism and to show how cross-domain benchmarking can be made transparent.

Third, the RTFR calculation reveals a more realistic readiness picture than raw performance alone. The illustrative value of 8.25 is constrained by quantum readiness and governance rather than by photonic throughput. This is desirable because deployment maturity depends on validation, robustness, auditability and institutional control as well as speed.

Domain-level interpretation also differs. In 6G and Earth observation, photonic acceleration is likely to have direct value because data-rate and transform bottlenecks are severe. In robotics and Industry 5.0, neuromorphic temporal processing may be particularly useful when events are sparse and decisions must be made at the edge. In biomedicine, the primary value of the architecture is not extreme autonomy but synchronized multimodal integration with calibrated uncertainty and explainable evidence.

The proposed framework therefore predicts that the optimal PNQ configuration will be heterogeneous. A 6G base-station implementation may allocate a larger weight to photonic acceleration; a wearable medical device may prioritize energy-efficient neuromorphic processing and defer quantum routines to remote infrastructure; an industrial digital twin may combine continuous conventional analytics with event-triggered neuromorphic fault detection and periodic quantum-assisted optimization.

Metric	**Conventional**	**PNQ-RTF illustrative**	**Interpretation**
Cross-domain fusion effectiveness	6.8	8.3	Higher indicates better integrated utility
Relative improvement	—	22.06%	Demonstration only; not empirical

Metric	Conventional	PNQ-RTF illustrative	Interpretation
Photonic readiness	—	8.7	High transform and throughput potential
Neuromorphic readiness	—	8.4	High event-driven temporal potential
Quantum readiness	—	7.4	Promising but validation-dependent
Governance readiness	—	7.6	Requires application-specific controls
Real-Time Fusion Readiness	—	8.25	Composite architecture score

4.1 Discussion

The theoretical contribution is a shift from multimodal model fusion to computational-layer fusion. Most sensor-fusion systems determine how multiple observations should be combined; PNQ-RTF additionally asks where each mathematical operation should be executed and how the choice of computational substrate affects end-to-end decision quality. This creates a bridge between algorithm design and hardware–software co-design.

For 6G, the supplied literature supports both the necessity of fast channel processing and the plausibility of photonic acceleration (Beghriche et al. 2026). Non-stationary channel models and V2X beam-space learning further imply that latency-sensitive adaptation will remain central (Sun et al. 2020; Tagliaferri et al. 2021). The framework therefore prioritizes streaming transforms, anomaly awareness and dynamic channel-state fusion rather than static network optimization.

For robotics, PNQ-RTF formalizes what recent sensor-fusion research already suggests: modality reliability varies over time. Differentiable attention weighting for LiDAR-camera fusion and landmark-augmented localization demonstrate adaptive alternatives to fixed fusion rules (Zhang, Lu & Xu 2026; Kocsis et al. 2026). A neuromorphic layer can extend this concept by treating unexpected events as first-

class computational signals rather than forcing every sensor into uniform frame-based processing.

For biomedicine, the ethical interpretation is more restrictive. Multimodal EEG/ECG models and explainable ECG systems indicate that richer signal fusion can improve decision support, but the framework does not equate faster inference with clinical autonomy (Xue et al. 2026; Santamónica et al. 2026). Human review, calibrated uncertainty and provenance of modality contributions should remain mandatory where outputs can affect diagnosis or treatment.

For Earth observation, photonic and quantum concepts should be evaluated against highly optimized classical spectral-processing baselines. Hyperspectral anomaly detection and multivariable Earth-observation models already achieve sophisticated representations (Huo et al. 2026; Miao et al. 2026). The research question is therefore whether PNQ layers reduce total decision time or energy at equivalent scientific accuracy, not whether they merely perform mathematically novel transformations.

For Industry 5.0, the framework aligns naturally with digital twins. Predictive-maintenance research increasingly combines AI, remaining-useful-life estimation and operational decision rules (Mehdipour et al. 2026; Wang et al. 2026). PNQ-RTF can operate as a real-time front end to the twin, transforming heterogeneous sensor streams into uncertainty-aware event states before maintenance optimization.

Technologically, one of the greatest risks is hidden interface overhead. Hybrid systems can lose theoretical gains during analog-to-digital conversion, synchronization, memory transfer or error correction. Equation 8 therefore treats interface latency explicitly. Future experiments should report end-to-end latency and energy from sensor input to actionable output, rather than reporting accelerator kernel speed alone.

Policy and governance requirements vary sharply by application. Spectrum adaptation and industrial maintenance can permit more automated execution than medical diagnosis. Robotics requires fail-safe control and bounded actions. Earth-observation alerts may influence emergency decisions and therefore require provenance and uncertainty communication. A universal architecture must

consequently include configurable governance constraints rather than a universal autonomy level.

5. Conclusion

This paper proposes the Photonic–Neuromorphic Quantum Real-Time Fusion framework as a unified architecture for signal, sensor and decision intelligence across 6G communications, robotics, biomedicine, Earth observation and Industry 5.0. The framework is motivated by a recurring cross-sector problem: heterogeneous sensors generate more data than conventional sequential pipelines can always transform, fuse and interpret within strict latency and energy budgets.

PNQ-RTF addresses this problem by assigning distinct roles to emerging computational paradigms. Photonic processing is used for high-throughput transforms and routing, neuromorphic processing for event-driven temporal integration, and quantum computing for selected representation or optimization tasks where benchmarking justifies its use. Uncertainty-aware fusion and governance constraints convert these computational states into domain-specific actions.

The mathematical framework includes latent-space matrix fusion, photonic transforms, spiking differential dynamics, variational quantum representations, Bayesian uncertainty, constrained multi-objective optimization and composite readiness scoring. The numerical demonstration yields a conventional fusion score of 6.80, an illustrative PNQ-RTF score of 8.30 and a Real-Time Fusion Readiness index of approximately 8.25. These values are explanatory rather than empirical.

The principal future value of the framework lies in benchmarking. It provides a common language for asking whether a hybrid architecture improves end-to-end real-time decision quality after conversion overhead, uncertainty, energy, reliability and governance are counted. Such evidence is essential before photonic, neuromorphic and quantum technologies can move from isolated demonstrations toward credible interoperable intelligence infrastructures.

5.1 Limitations

- No empirical dataset or hardware benchmark was supplied; all numerical values are illustrative.

- The 0–10 normalization simplifies heterogeneous physical metrics such as latency, energy, throughput and calibration error.
- Photonic, neuromorphic and quantum components are represented functionally rather than through a specific fabrication stack or device platform.
- Quantum benefit is not assumed and may disappear when error, data loading or interface overhead are included.
- Cross-domain comparison can conceal application-specific safety thresholds and should not replace domain validation.
- The proposed fusion weights require empirical learning, expert calibration or multi-criteria validation before deployment.

5.2 Future Research Directions

11. Build synchronized cross-domain benchmark suites with end-to-end latency, energy and uncertainty metadata.
12. Prototype photonic front ends for 6G channel inference and hyperspectral transforms and compare full-stack performance with GPUs and edge accelerators.
13. Develop photonic-neuromorphic co-designs in which optical memory or recurrent structures support event-driven temporal inference.
14. Benchmark variational quantum and quantum-photonic subroutines against strong classical optimization and kernel methods under equal data-loading assumptions.
15. Extend PNQ-RTF to federated and privacy-preserving biomedical/industrial settings where raw data cannot be centralized.
16. Develop modality-level explainability that identifies which sensor, time window and computational layer contributed to each decision.
17. Integrate formal safety constraints and runtime assurance for robotic and industrial autonomous actions.
18. Establish lifecycle sustainability metrics for hybrid hardware, including embodied energy, cooling, photonic conversion and quantum-system overhead.

References

Beghriche, A, Bouaita, B, Benmessahel, B, Berrhail, F & Seghir, F 2026, 'Photonic AI Accelerators for Ultra-Fast Wireless Signal Processing in 6G Networks', Traitement du Signal, vol. 43, no. 1, pp. 27–51, accessed <http://dx.doi.org/10.18280/ts.430103>.

Boras, I, Ewaniuk, J & Rotenberg, N 2026, 'Quantum photonic neural networks in time', Optica Quantum, accessed <http://dx.doi.org/10.1364/opticaq.603710>.

Daneshzad, S & Parsanasab, G-M 2026, 'Photonic Quantum Hardware for AI Optimization Tasks', 2026 International Interdisciplinary Conference on Artificial Intelligence: Engineering, Health, Finance and Humanities (IICAI), pp. 1–6, accessed <http://dx.doi.org/10.1109/iicai70155.2026.11620969>.

Gençtoy, S 2026, 'A PyCaret-Assisted Comparative Study of AI/ML Models for Anomaly Detection in 5G and 6G Networks', 2026 34th Signal Processing and Communications Applications Conference (SIU), pp. 1–4, accessed <http://dx.doi.org/10.1109/siu71813.2026.11636811>.

Hanik, N, Kernetzky, T, Höfler, U, Jia, Y, Schubert, C, Karvar, M, Sackey, I, Freund, R, Ronniger, G & Zimmermann, L 2026, 'Ultra-Broadband Photonic Signal-Processor', Electronic-Photonic Integrated Systems for Ultrafast Signal Processing, pp. 1–50, accessed <http://dx.doi.org/10.1007/978-3-032-08340-1_1>.

Haq, HU, Qureshi, A & Khan, AQ 2026, 'Kalman Filter–Based Sensor Fusion for Navigation of Holonomic Unmanned Ground Vehicles', Journal of Field Robotics, vol. 43, no. 4, pp. 2715–2724, accessed <http://dx.doi.org/10.1002/rob.70166>.

Huo, Y, Wang, S, Wang, C, Zhang, M & Wang, H 2026, 'Dual-stream background modeling network with anomaly suppression for hyperspectral anomaly detection', International Journal of Applied Earth Observation and Geoinformation, vol. 148, p. 105233, accessed <http://dx.doi.org/10.1016/j.jag.2026.105233>.

Jahannia, B, Hayden, B, Yu, I & Dalir, H 2026, 'Mixed-signal DAC/ADC co-design and photonic system integration', H Dalir & A Adibi (eds.), Reconfigurable Photonic Computing and Heterogeneous Architectures for AI and RF Defense Systems, p. 15, accessed <http://dx.doi.org/10.1117/12.3109126>.

Kocsis, M, Knorz, O, Sautter, M & Zöllner, R 2026, 'Landmark-Augmented Multi-Sensor Fusion for Autonomous Shuttle Localization', 2026 IEEE Intelligent Vehicles Symposium (IV), pp. 1982–1989, accessed <http://dx.doi.org/10.1109/iv66570.2026.11624075>.

Mehdipour, S, Habibzadeh, A, Esmi, N, Heravi, P, Mirsaraee, M, Shahbahrami, A & Javanmardi, Z 2026, 'Integrating Digital Twins, Data Management and Artificial Intelligence for Predictive Maintenance in Aviation: A Comprehensive Review', Digital Twins and Applications, vol. 3, no. 1, accessed <http://dx.doi.org/10.1049/dgt2.70029>.

Miao, Z, Li, C, Su, T & Wu, Y 2026, 'Global assessment of climate-driven lag and cumulative effects on vegetation using Earth Observation data and multivariable machine learning models', International Journal of Applied Earth Observation and Geoinformation, vol. 152, p. 105403, accessed <http://dx.doi.org/10.1016/j.jag.2026.105403>.

Mondal, SS, Das, A, Huq, S, Golder, A & Hossain, MA 2026, 'QuANTiM: Quantum-Inspired Adaptive Neural Networks for Photonic IoT Edge Intelligence', 2026 IEEE 2nd International Conference on Quantum Photonics, Artificial Intelligence & Networking (QPAIN), pp. 1–5, accessed <http://dx.doi.org/10.1109/qpain69676.2026.11545553>.

Nguyen, NH, Behrman, EC & Steck, JE 2020, 'Quantum learning with noise and decoherence: a robust quantum neural network', Quantum Machine Intelligence, vol. 2, no. 1, accessed <http://dx.doi.org/10.1007/s42484-020-00013-x>.

Oliveira, M, Castro, A, Madeira, T, Pedrosa, E, Dias, P & Santos, V 2020, 'A ROS framework for the extrinsic calibration of intelligent vehicles: A multi-sensor, multi-modal approach', Robotics and Autonomous Systems, vol. 131, p. 103558, accessed <http://dx.doi.org/10.1016/j.robot.2020.103558>.

Omoloja, YO, Ogunwa, T, Marian, RM & Chahl, J 2026, 'A Critical Analysis and Roadmap for the Development and Implementation of Predictive Maintenance Digital Twins', Applied System Innovation, vol. 9, no. 6, p. 123, accessed <http://dx.doi.org/10.3390/asi9060123>.

Randel, S, Mahmud, MdS, Schindler, A, Runge, P & Schell, M 2026, 'Optic-Electronic-Optic Interferometers for Ultrabroadband Arbitrary Digital Signal Processing', Electronic-Photonic Integrated Systems for Ultrafast Signal Processing, pp. 51–82, accessed <http://dx.doi.org/10.1007/978-3-032-08340-1_2>.

Saffer, O, Marines Cabello, JH, Slinkov, G, Becker, S, Geilen, A & Stiller, B 2026, 'Brillouin-based memory and signal processing for photonic neuromorphic computing', M Notomi & T Zhou (eds.), AI and Optical Data Sciences VII, p. 61, accessed <http://dx.doi.org/10.1117/12.3077887>.

Santamónica, A, Hortal, E, Larriba, Y & Rueda, C 2026, 'Explainable AI for automatic heart disease diagnosis using 3DFMMecg features: A novel ECG-based approach', Biomedical Signal Processing and Control, vol. 113, p. 109085, accessed <http://dx.doi.org/10.1016/j.bspc.2025.109085>.

Sapkale, P 2026, 'A Quantum Computing Perspective in 6G Networks', Quantum Computing and Machine Learning for 6G, pp. 311–325, accessed <http://dx.doi.org/10.1002/9781394238118.ch14>.

Shao, J, Lan, W, He, S & Guo, D 2026, 'An accurate and efficient neuromorphic computing framework with UKF-assisted spiking recurrent neural network for fault detection in UAV', Digital Signal Processing, vol. 178, p. 106160, accessed <http://dx.doi.org/10.1016/j.dsp.2026.106160>.

Skolik, A, McClean, JR, Mohseni, M, van der Smagt, P & Leib, M 2021, 'Layerwise learning for quantum neural networks', Quantum Machine Intelligence, vol. 3, no. 1, accessed <http://dx.doi.org/10.1007/s42484-020-00036-4>.

Stupina, AA, Kukartseva, SV & Stepantsevich, MN 2026, 'Application of Graph Neural Networks for Structural Dependency Analysis in Industrial Digital Twins Using the UCI AI4I 2020 Predictive Maintenance Dataset', 2026 International Conference on Industrial Engineering, Applications and Manufacturing (ICIEAM), pp. 1155–1159, accessed <http://dx.doi.org/10.1109/icieam69213.2026.11549674>.

Sun, Y, Wang, C-X, Huang, J & Wang, J 2020, 'A 3D Non-Stationary Channel Model for 6G Wireless Systems Employing Intelligent Reflecting Surface', 2020 International Conference on Wireless Communications and Signal Processing (WCSP), pp. 19–25, accessed <http://dx.doi.org/10.1109/wcsp49889.2020.9299848>.

Tagliaferri, D, Mizmizi, M, Badini, D, Mazzucco, C & Spagnolini, U 2021, 'Multi-Vehicular Beam Space Learning for Channel Estimation in 6G V2X Systems', 2021 IEEE 22nd International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), pp. 46–50, accessed <http://dx.doi.org/10.1109/spawc51858.2021.9593123>.

Tang, S, Chen, Z & Zhang, M 2021, 'Spectral Quality Evaluation of Reconstructed Hyperspectral Images', 2021 11th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS), pp. 1–5, accessed <http://dx.doi.org/10.1109/whispers52202.2021.9483970>.

Vincent, F & Besson, O 2021, 'Target detection in hyperspectral imaging combining replacement and additive models', Signal Processing, vol. 188, p. 108212, accessed <http://dx.doi.org/10.1016/j.sigpro.2021.108212>.

Wang, L, Chen, Y, Zhu, Z & Zhao, X 2026, 'Dynamic predictive maintenance framework for mechanical systems via uncertainty-aware RUL estimation', Mechanical Systems and Signal Processing, vol. 247, p. 113977, accessed <http://dx.doi.org/10.1016/j.ymssp.2026.113977>.

Wang, Y, Wang, Z, Cao, N & Cao, Y 2026, 'AI-driven multimodal sensor fusion for quantitative rehabilitation assessment and remote health monitoring: A Baduanjin

Case study', Biomedical Signal Processing and Control, vol. 118, p. 109693, accessed <http://dx.doi.org/10.1016/j.bspc.2026.109693>.

Wei, G, Zhang, C, Li, D, Lou, J, Jin, R, Chen, Z & Zheng, Z 2026, 'MECKA: A Multimodal ECG Classification Framework via Knowledge Augmentation', Biomedical Signal Processing and Control, vol. 125, p. 110738, accessed <http://dx.doi.org/10.1016/j.bspc.2026.110738>.

Wu, F, Lin, Z, Yin, B, Zhang, L & Chen, B 2026, 'EEG-HMLearner: Hybrid learning for visual stimulus decoding by extracting multimodal information from EEG signals', Biomedical Signal Processing and Control, vol. 122, p. 110378, accessed <http://dx.doi.org/10.1016/j.bspc.2026.110378>.

Xue, Z, Zhong, W, Liu, S & An, X 2026, 'A cross-attention multimodal fusion transformer for ECG-guided EEG mental workload classification', Biomedical Signal Processing and Control, vol. 124, p. 110666, accessed <http://dx.doi.org/10.1016/j.bspc.2026.110666>.

Xu, X 2026, 'Design of Autonomous Navigation System for Mobile Robots Based on Multi-Sensor Fusion and Deep Reinforcement Learning', 2026 6th International Conference on Machine Learning and Intelligent Systems Engineering (MLISE), pp. 374–377, accessed <http://dx.doi.org/10.1109/mlise70044.2026.11607485>.

Yao, H, Liu, X, Jing, M, Li, G & Wang, X 2026, 'LCQNN: linear combination of quantum neural networks', Quantum Machine Intelligence, vol. 8, no. 1, accessed <http://dx.doi.org/10.1007/s42484-026-00364-x>.

Zeng, W, Shan, L, Chen, Y, Ma, Y & Du, S 2026, 'Interpretable deep learning framework for multi-label ECG classification: Enhancing cardiac diagnostics through feature attention and explainability', Biomedical Signal Processing and Control, vol. 112, p. 108447, accessed <http://dx.doi.org/10.1016/j.bspc.2025.108447>.

Zhang, J, Lu, Y & Xu, X 2026, 'PentaFusion: Differentiable attention weighting for real-time LiDAR-camera fusion in edge autonomous vehicles', Robotics and Autonomous Systems, vol. 196, p. 105233, accessed <http://dx.doi.org/10.1016/j.robot.2025.105233>.

Source-control note: All references in this paper were selected from the bibliography supplied by the user; no outside references were added.