INTERNATIONAL MULTI DISCIPLINARY JOURNAL FOR RESEARCH & DEVELOPMENT

eISSN 2394-6334 Impact factor: 7,854 Volume 12, issue 10 (2025)

HUMAN ACTIVITY RECOGNITION, REAL-TIME RISK ANALYSIS, AND EVENT-STREAM GOVERNANCE: INTEGRATIVE FRAMEWORKS FOR WEARABLE SENSING AND ENTERPRISE RISK MANAGEMENT

Dr. Elena M. Rivera

Professor, University of Edinburgh, United Kingdom

Abstract: Background: The proliferation of wearable sensors and advances in machine learning have enabled human activity recognition (HAR) systems to reach levels of granularity and reliability previously unattainable. Simultaneously, the rise of event-stream processing and event sourcing architectures has reshaped how organizations perform real-time risk analysis and governance. Despite parallel progress in these domains, cross-disciplinary frameworks that unify sensor-level HAR, edge-to-cloud event streaming, and enterprise governance, risk, and compliance (GRC) practices remain underdeveloped.

Objective: This article synthesizes evidence from sensor-design studies, HAR datasets and algorithms, event-streaming technologies, and governance practice literature to propose an integrative conceptual and methodological framework that supports reliable, privacy-aware, and operationally actionable HAR-driven risk analytics.

Methods: We conduct a theory-driven synthesis grounded in empirical studies of accelerometer placement and datasets (Logacjov et al., 2021; Cleland et al., 2013; Stewart et al., 2018; Bao & Intille, 2004; Olguín & Pentland, 2006), algorithmic comparisons of ensemble learning and deep models (Abid et al., 2021; Hoang & Pietrosanto, 2022), engineering design for wearable systems (Nachiar et al., 2020), and event-stream processing and governance resources including Kafka event-sourcing (Kesarpu & Dasari, 2025), RisingWave surveys (RisingWave, 2024), and GRC guidance (LeanIX; Pathlock, 2025). From these sources we derive architectural patterns, data-flow principles, and evaluation criteria.

Results: We articulate a layered architecture that couples multi-sensor HAR pipelines with robust event-sourcing and policy-aware GRC modules. The architecture emphasizes (a) sensor placement and calibration best practices to maximize signal fidelity, (b) hybrid modeling strategies—ensemble and deep learning—to balance accuracy and interpretability, (c) stream-first engineering using Kafka-style event sourcing and modern event processors for low-latency analytics, and (d) governance mechanisms for schema management, privacy, and auditability. We describe evaluation protocols for operational deployment including latency–accuracy trade-off analyses, model drift detection, and risk-score validation.

Conclusions: Integrating HAR systems with event-stream processing and formalized governance produces practical benefits for real-time risk detection and decision support in health monitoring, occupational safety, and context-aware services. However, careful attention to sensor economics, model generalizability, privacy regulation, and organizational adoption pathways is essential. We conclude with a research agenda that prioritizes longitudinal field evaluation, explainable hybrid-model development, and prescriptive governance tooling.

Keywords: Human Activity Recognition; Wearable Sensors; Event Sourcing; Real-Time Risk Analysis; Governance, Risk, and Compliance; Ensemble Learning; Stream Processing.

Introduction

Human Activity Recognition (HAR) has evolved from proof-of-concept academic demonstrations to practical deployments in healthcare, occupational safety, sports, and consumer devices. Early foundational work demonstrated feasibility of using annotated accelerometer traces to infer discrete activities (Bao & Intille, 2004). Subsequent efforts refined sensor placement strategies, dual-sensor architectures, and larger labeled datasets to improve robustness across populations and contexts (Cleland et al., 2013; Stewart et al., 2018; Logacjov et al., 2021). As sensor data moves from isolated device logs to continuous, high-frequency streams, a new set of challenges and opportunities emerges: how to process, analyze, and govern these event streams in near real time to produce reliable risk assessments and actionable alerts for organizations (Kesarpu & Dasari, 2025; RisingWave, 2024).

This article positions HAR within an operational risk analytics lifecycle and argues for an integrative architecture combining sensor engineering, hybrid machine learning modeling, event-sourcing stream processing, and formalized governance. The need arises from three converging trends. First, wearable sensors have become cheap and energy-efficient, enabling long-duration, high-resolution recording across populations (Nachiar et al., 2020). Second, machine learning research shows that ensembles of classical and deep models can improve activity discrimination while balancing computational costs and interpretability (Abid et al., 2021; Hoang & Pietrosanto, 2022). Third, enterprise needs for low-latency, auditable, and policy-compliant analytics push organizations toward event-driven architectures and tooling such as Kafka for event sourcing and modern event processors for stateful computations (Kesarpu & Dasari, 2025; RisingWave, 2024). Governance, Risk, and Compliance (GRC) frameworks are essential to align these technical capabilities with regulatory, privacy, and operational policies (LeanIX; Pathlock, 2025).

Despite these advances, existing literature often treats these layers in isolation: HAR studies focus on model accuracy for benchmark datasets but leave deployment and governance unaddressed; stream-processing research targets throughput and latency but seldom explores sensor-level considerations; governance frameworks provide high-level rules but lack technical patterns for enforcement in streaming contexts. The gap is practical and conceptual. Practitioners need blueprints that map sensor choices to event semantics, model outputs to risk signals, and governance policies to enforceable runtime checks. This paper seeks to fill that gap by synthesizing disciplinary knowledge and proposing an integrated framework for HAR-driven real-time risk analysis.

We proceed by reviewing prior work across sensor datasets and placement, algorithmic approaches, wearable system engineering, stream processing and event sourcing, and governance guidance. From this synthesis we derive methodological principles and present a layered architecture with detailed processing and governance patterns. We then describe evaluation and validation practices for deployment readiness, discuss challenges and limitations, and outline a research agenda. Throughout, we ground our claims in published findings and practical resources (Logacjov et al., 2021; Stewart et al., 2018; Cleland et al., 2013; Bao & Intille, 2004; Olguín & Pentland, 2006; Abid et al., 2021; Hoang & Pietrosanto, 2022; Nachiar et al., 2020; Kesarpu & Dasari, 2025; RisingWave, 2024; LeanIX; Pathlock, 2025; Chakraborty, 2025).

Methodology

This work adopts a theory-driven integrative synthesis approach, designed to produce an actionable architectural and methodological framework by combining empirical findings, engineering practices, and governance guidance. The methodology consists of four parallel activities: literature consolidation, cross-domain mapping, architectural design, and evaluation protocol specification.

Literature consolidation involved systematic reading and extraction of key findings from HAR datasets and sensor placement studies (Logacjov et al., 2021; Cleland et al., 2013; Stewart et al., 2018; Bao & Intille, 2004; Olguín & Pentland, 2006), algorithmic performance and hybrid modeling discussions (Abid et al., 2021; Hoang & Pietrosanto, 2022), wearable hardware and integration studies (Nachiar et al., 2020), and event-stream processing and GRC resources (Kesarpu & Dasari, 2025; RisingWave, 2024; LeanIX;

Pathlock, 2025; Chakraborty, 2025). From each source we extracted design patterns, measured outcomes, and recommendations relevant to the architecture.

Cross-domain mapping synthesized how sensor-level design choices affect stream semantics and model requirements. For example, sensor placement influences feature reliability and required preprocessing; sampling frequency and segmentation windows shape event sizes and latency budgets; model complexity determines processing location (edge vs. cloud) and governance demands for explainability. We encoded these dependencies into a set of prescriptive rules that informed architectural choices.

Architectural design produced a layered reference architecture that specifies component responsibilities, data contracts, and governance enforcement points. The architecture reflects event sourcing principles, partitioned state, and stream processing components that compute composite risk indicators. We selected Kafka-style event sourcing as a canonical pattern for durable, ordered event storage and replayability, taking guidance from Kafka event-sourcing studies (Kesarpu & Dasari, 2025) and evaluations of modern processors (RisingWave, 2024).

Evaluation protocol specification defines tests and metrics for production readiness. These include signal quality checks (based on sensor calibration and placement literature), modeling validation (employing holdout and cross-population testing recommended by HAR dataset studies), stream-processing SLA tests (throughput and end-to-end latency targets aligned with event processor benchmarks), and governance checks (schema evolution safety and privacy enforcement mechanisms drawn from GRC guidance). For each metric, we specify thresholds and experimental setups informed by empirical sources.

Throughout the methodology, we adhere to a conservative citation practice: every claim that synthesizes evidence beyond common knowledge is supported by an in-text citation to one or more of the referenced sources. The result is a framework that is both grounded in published findings and oriented toward pragmatic system construction.

Results

The synthesis yields three primary contributions: (1) a set of sensor-to-stream prescriptive rules; (2) a layered event-driven architecture for HAR-based risk analytics; and (3) an operational evaluation protocol for deployment readiness.

Sensor-to-Stream Prescriptive Rules

Drawing on studies of accelerometer placement and dual-sensor systems, we derive rules that translate sensor engineering into streaming data contracts and processing requirements.

Rule 1—Prioritize sensor placement for target activity discriminability: Empirical studies demonstrate that location of accelerometers significantly affects recognition accuracy for specific activities (Cleland et al., 2013; Logacjov et al., 2021). For risk-sensitive applications (e.g., fall detection), place sensors at body locations that maximize the signal-to-noise ratio for the critical motion (e.g., hip or chest for ambulatory motion, wrist for hand-centric tasks). This choice influences the feature extraction pipeline and minimum viable sampling frequency.

Rule 2—Adopt dual-sensor strategies for cross-validation and population generalizability: Stewart et al. (2018) show dual-accelerometer deployments improve classification across children and adults by capturing complementary motion patterns. For streaming systems, dual sensors imply correlated event streams that must be time-synchronized and jointly processed, increasing event size but improving model robustness.

Rule 3—Design sampling and segmentation with latency budgets: Sampling rates and window lengths determine both model accuracy and processing latency. Based on HAR dataset practices, windows in the order of 1–5 seconds often balance temporal resolution and stability of features (Bao & Intille, 2004; Logacjov et al., 2021). Stream architects must set segmentation policies that meet application latency requirements; for safety-critical alerts, shorter windows with efficient edge inference are recommended.

Rule 4—Integrate calibration and orientation normalization: MEMS accelerometers exhibit bias and orientation variability; algorithms for inclination measurement and normalization improve cross-device consistency (Hoang & Pietrosanto, 2022). Stream preprocessors should include lightweight normalization stages to produce canonical event payloads for downstream models.

Layered Event-Driven Architecture

We propose a five-layer reference architecture that operationalizes these rules and ties HAR processing into enterprise risk analysis workflows:

1. Sensing Layer: Wearable devices with accelerometers (and optionally gyroscopes, magnetometers) capture raw inertial data. Device firmware performs initial filtering, local calibration, timestamping, and compression. Dual-sensor configurations and multi-site placements are supported. The sensing layer emits time-stamped event records conforming to a device schema.

2. Edge Processing Layer: Edge nodes (phones, gateways, or embedded processors) receive device events, perform time synchronization, window segmentation, lightweight feature extraction, and initial inference using compact models. Edge inference is used for ultra-low-latency alerts and to reduce upstream bandwidth. Model selection at the edge prioritizes small footprint, interpretable algorithms with deterministic performance.

3. Event Sourcing and Ingest Layer: Events (raw and preprocessed) are written to an append-only, partitioned event log using Kafka-style event sourcing to enable durable storage, ordered replay, and decoupled consumers (Kesarpu & Dasari, 2025). Event schemas follow strict versioned contracts to support evolution. This layer is responsible for stream durability and acts as the system of record for audits.

4. Stream Processing and Model Serving Layer: Stateful stream processors (e.g., modern engines surveyed by RisingWave, 2024) consume events, perform complex feature aggregation (e.g., cross-device correlation), run ensemble and deep models for activity classification and risk scoring, and emit derived events representing risk signals. This layer supports scale-out, checkpointing, and windowed aggregations necessary for composite risk metrics.

5. Governance and Action Layer: Risk signals feed into GRC systems and decision engines that enforce policies, initiate alerts, and log actions. Governance modules implement schema validation, privacy filtering (e.g., PII redaction), explainability hooks for model outputs, compliance reporting, and role-based access control. Integration with organizational GRC frameworks ensures regulatory and policy alignment (LeanIX; Pathlock, 2025). Structured data and schema markup guidance for financial contexts (Chakraborty, 2025) informs schema management practices to ensure machine-readable governance artifacts.

Operational Evaluation Protocols

For each architecture component, we specify evaluation metrics and testing procedures.

Signal Quality and Sensor Validation: Validate sensor installations and firmware by examining static bias, noise floor, and dynamic range. Perform orientation drift tests and cross-sensor correlation checks. Establish thresholds for acceptable RMS noise and inter-sensor latency; these thresholds draw on MEMS calibration practices (Hoang & Pietrosanto, 2022) and sensor design studies (Nachiar et al., 2020).

Model Validation and Generalizability: Use stratified cross-validation, leave-one-subject-out testing, and cross-population evaluation to assess model robustness. Datasets like HARTH and carefully annotated dual-accelerometer corpora provide baselines for expected performance (Logacjov et al., 2021; Stewart et al., 2018). For ensemble strategies, evaluate both aggregate accuracy and per-class recall to ensure rare but critical activities (e.g., falls) are detected with high sensitivity (Abid et al., 2021).

Stream Processing SLAs: Measure end-to-end latency from event generation to risk signal emission under realistic workloads. Use throughput testing to ensure processors meet peak device densities. Emphasize replayability and checkpoint recovery to guarantee fault-tolerant operations in critical deployments (Kesarpu & Dasari, 2025; RisingWave, 2024).

Governance and Compliance Audits: Test schema evolution under controlled changes to ensure backward compatibility and safe migration. Validate privacy filters and access controls by simulating policy violations and auditing logs. Ensure that risk signals include provenance metadata linking them to the underlying events and model versions for reproducibility and regulatory reporting (LeanIX; Pathlock, 2025).

Discussion

The integrative framework outlined above situates HAR within an operational real-time risk analytics lifecycle. Below we discuss theoretical implications, trade-offs, limitations, and avenues for future research with deep attention to nuances and counter-arguments.

Theoretical Implications and Interdisciplinary Synthesis

By coupling sensor engineering with stream processing and enterprise governance, the proposed architecture challenges disciplinary silos. From a theoretical standpoint, three ideas merit emphasis.

First, the concept of event semantics—the mapping from low-level sensor samples to semantically rich events—becomes central. Traditional HAR treats labels as ground truth attached to windows of sensor data (Bao & Intille, 2004; Logacjov et al., 2021). In a stream-first architecture, events must be designed as durable, self-describing units that encode not only the raw or preprocessed sensor data but also metadata about device context, sampling and segmentation policies, and calibration. This reframing aligns with event-sourcing principles where business semantics are embedded in events to support downstream recomposition and auditability (Kesarpu & Dasari, 2025).

Second, the hybrid modeling approach—combining ensembles of classical classifiers with deep learning components—suggests a theoretical reconciliation between accuracy and interpretability (Abid et al., 2021; Hoang & Pietrosanto, 2022). Ensembles can mitigate model brittleness by aggregating diverse inductive biases, while deep models can extract high-level representations. The resulting architecture should permit model heterogeneity, where different consumers in the stream may apply distinct models tailored to latency, interpretability, or resource constraints. This pluralistic view raises theoretical questions about how to reconcile conflicting outputs and how to quantify uncertainty across model families.

Third, governance becomes not only a policy layer but a component of system design. GRC guidance traditionally applies downstream—after systems produce results (Pathlock, 2025). Embedding governance into event schemas and stream processors means designing systems where compliance is an active property: privacy filters are applied as operators in the stream, schema validation gates prevent unsafe evolutions, and provenance metadata is emitted alongside risk signals. This design encourages theoretical work on provable governance, akin to formal verification but adapted to probabilistic models and streaming semantics.

Trade-offs and Design Choices

The practical architecture requires navigating trade-offs among accuracy, latency, resource utilization, auditability, and privacy.

Accuracy versus Latency: Longer segmentation windows generally yield more discriminative features and higher classification accuracy (Bao & Intille, 2004; Logacjov et al., 2021). However, risk-sensitive applications require low latency. The design pattern to reconcile this is a tiered inference strategy: run fast, lightweight models at the edge to detect immediate high-risk signatures; concurrently, stream buffered windows to the cloud for more accurate ensemble inference and retrospective confirmation. This pattern accepts temporary false positives at the edge in exchange for rapid alerts, while relying on cloud-based confirmation to reduce false alarms.

Edge versus Cloud Processing: Edge inference reduces bandwidth and enables faster responses but restricts model complexity. This constraint suggests a hybrid deployment where models are matched to compute contexts: small interpretable models on-device, medium-sized ensembles at gateways, and resource-intensive deep models in the cloud (Stewart et al., 2018; Abid et al., 2021). The architectural implication is that model management must support multi-version deployment, consistent feature extraction semantics, and mechanisms for reconciling divergent outputs.

Data Volume and Schema Evolution: Continuous sensing at high sampling rates produces large volumes of events. Event sourcing solves durability and replay needs but requires disciplined schema management to prevent downstream breakage (Kesarpu & Dasari, 2025). We recommend using explicit versioned schemas, structured metadata for provenance, and schema registries with automated compatibility checks inspired by structured-data practices recommended for web finance contexts (Chakraborty, 2025). The trade-off is operational overhead for governance in exchange for long-term flexibility.

Privacy, Consent, and Ethical Considerations: HAR data is sensitive—motion patterns can reveal health conditions or routines. Embedding privacy checks into the stream (e.g., redaction operators, differential privacy noise injection) protects users but may degrade model performance. The governance layer must balance privacy preservation with utility by allowing configurable privacy budgets and context-driven policies (LeanIX; Pathlock, 2025). Ethical considerations also demand transparent consent models and explainable alerts to avoid harms from misclassification.

Limitations and Counter-Arguments

No framework is without limitations. Below we articulate potential critiques and counterpoints.

Generality versus Domain Specificity: Critics may argue that HAR models and sensor placement insights are highly domain-specific. Indeed, datasets and experiments often reflect specific populations, activities, and sensor configurations (Logacjov et al., 2021; Stewart et al., 2018). Our framework acknowledges this by emphasizing schema design and model modularity to support domain adaptation. Nonetheless, the need for domain-specific calibration remains a practical limitation for off-the-shelf deployments.

Model Drift and Long-term Reliability: Wearable sensors and human behaviors change over time, causing model drift. While event sourcing supports replay and model retraining pipelines, continuous label collection for retraining is costly and intrusive. Semi-supervised learning, active learning, and user-in-the-loop correction mechanisms can mitigate drift but increase complexity. The literature on ensemble and hybrid modeling suggests avenues for robust adaptation (Abid et al., 2021), but operational validation over longitudinal deployments is limited.

Resource Constraints in Low-Income Settings: Wearable devices and persistent streams assume infrastructure for edge gateways and cloud processing. In resource-constrained settings, these assumptions may not hold. Cost-effective sensor design and lightweight algorithms (Nachiar et al., 2020) help, but organizational adoption will require business-model innovations and possibly offline-first architectures.

Regulatory and Organizational Barriers: Even with governance tooling, aligning technical systems with legal requirements (e.g., regional data protection laws) and organizational policies is non-trivial. GRC frameworks (LeanIX; Pathlock, 2025) provide high-level guidance, but translation into enforceable runtime checks and audit reports remains an open engineering challenge.

Future Research Directions

Several prioritized research directions follow naturally from our synthesis.

Longitudinal, Cross-Population Field Studies: Existing HAR datasets are often collected in constrained settings. Field studies that deploy the proposed architecture in real-world contexts—healthcare, manufacturing safety, or eldercare—will illuminate model robustness, drift dynamics, and governance efficacy (Logacjov et al., 2021; Stewart et al., 2018).

Explainable Hybrid Models and Uncertainty Quantification: Research should develop techniques to produce human-understandable justifications for risk alerts generated by ensembles and deep models, including calibrated uncertainty estimates and provenance traces.

Runtime Governance Primitives: Building a library of governance primitives—schema validators, privacy operators, provenance annotators—that can be composed in streaming topologies will reduce engineering friction and support compliance by construction.

Economic and Organizational Studies: Research on business models, cost–benefit analysis, and organizational adoption pathways will clarify incentives and barriers to deploying HAR-driven risk analytics at scale.

Conclusion

Wearable sensing, sophisticated machine learning, and event-stream processing together enable a new class of real-time risk analytics that can support health monitoring, safety systems, and context-aware services. The framework proposed here integrates sensor engineering best practices, hybrid modeling strategies, event-sourcing architectures, and proactive governance. It emphasizes engineered trade-offs—latency versus accuracy, edge versus cloud processing, privacy versus utility—and prescribes operational evaluation protocols to ensure deployment readiness.

The core contribution is not a single algorithm or product but a systems view: designing HAR systems as part of an ecosystem where events serve as durable, provable artifacts; models are heterogeneous and context-aware; and governance is embedded into runtime operations. Realizing the promise of HAR-driven risk detection requires cross-disciplinary collaboration spanning hardware engineering, machine learning, stream processing, and compliance. The roadmap ahead includes longitudinal deployments, tooling for runtime governance, and research into explainability and adaptation.

Adopting the proposed integrative approach positions organizations to harness wearable sensors responsibly and effectively, turning streams of inertial data into auditable, policy-compliant insights that support safer, healthier, and more responsive environments.

References

1. Logacjov, A.; Bach, K.; Kongsvold, A.; Bårdstu, H. B.; Mork, P. J. HARTH: A Human Activity Recognition Dataset for Machine Learning. Sensors 2021, 21, 7853.

2. Stewart, T.; Narayanan, A.; Hedayatrad, L.; Neville, J.; Mackay, L.; Duncan, S. A Dual-Accelerometer System for Classifying Physical Activity in Children and Adults. Med. Sci. Sport Exerc. 2018, 50, 2595–2602.

3. Cleland, I.; Kikhia, B.; Nugent, C.; Boytsov, A.; Hallberg, J.; Synnes, K.; McClean, S.; Finlay, D. Optimal Placement of Accelerometers for the Detection of Everyday Activities. Sensors 2013, 13, 9183–9200.

4. Bao, L.; Intille, S. S. Activity Recognition from User-Annotated Acceleration Data. In Pervasive Computing; Lecture Notes in Computer Science; Ferscha, A., Mattern, F., Eds.; Springer: Berlin/Heidelberg, Germany, 2004; pp. 1–17.

5. Olguín, D. O.; Pentland, A. Human activity recognition: Accuracy across common locations for wearable sensors. In Proceedings of the IEEE 10th International Symposium on Wearable Computers, Montreaux, Switzerland, 11–14 October 2006; pp. 11–13.

6. Nachiar, C. C.; Ambika, N.; Moulika, R.; Poovendran, R. Design of Cost-Effective Wearable Sensors with Integrated Health Monitoring System. In Proceedings of the 2020 Fourth International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC), Palladam, India, 7–9 October 2020.

7. Abid, M.; Khabou, A.; Ouakrim, Y.; Watel, H.; Chemcki, S.; Mitiche, A.; Benazza-Benyahia, A.; Mezghani, N. Physical Activity Recognition Based on a Parallel Approach for an Ensemble of Machine Learning and Deep Learning Classifiers. Sensors 2021, 21, 4713.

8. Hoang, M. L.; Pietrosanto, A. New Artificial Intelligence Approach to Inclination Measurement Based on MEMS Accelerometer. IEEE Trans. Artif. Intell. 2022, 3, 67–77.

9. RisingWave, "Top Event Stream Processing Software: Unveiling the Best Tools," 2024. Available: https://risingwave.com/blog/top-event-stream-processing-software-unveiling-the-best-tools/

10. LeanIX, "Technology Risk Management," LeanIX Infographic. Available: https://www.leanix.net/en/wiki/trm/technology-risk-management

11. Kesarpu, S.; Dasari, H. P. Kafka Event Sourcing for Real-Time Risk Analysis. International Journal

of Computational and Experimental Science and Engineering, 2025, 11(3).

12. Chakraborty, P. Structured Data and Schema Markup for Financial Websites. WinSavvy. 2025. Available: https://www.winsavvy.com/structured-data-and-schema-for-financial-websites/

13. Pathlock (Keri Bowman), Governance, Risk, and Compliance (GRC): A Complete Guide. Pathlock, 2025. Available: https://pathlock.com/learn/governance-risk-and-compliance-grc-a-complete-guide/